Simple constructions of linear-depth t-designs and pseudorandom unitaries

Tony Metger[1], Alexander Poremba[2], Makrand Sinha[3], and Henry Yuen[4]

[1]ETH Zurich
[2]Massachusetts Institute of Technology
[3]University of Illinois, Urbana-Champaign
[4]Columbia University

Abstract

Uniformly random unitaries, i.e. unitaries drawn from the Haar measure, have many useful properties, but cannot be implemented efficiently. This has motivated a long line of research into random unitaries that "look" sufficiently Haar random while also being efficient to implement. Two different notions of derandomisation have emerged: t-designs are random unitaries that information-theoretically reproduce the first t moments of the Haar measure, and pseudorandom unitaries (PRUs) are random unitaries that are computationally indistinguishable from Haar random.

In this work, we take a unified approach to constructing t-designs and PRUs. For this, we introduce and analyse the "PFC ensemble", the product of a random computational basis permutation P, a random binary phase operator F, and a random Clifford unitary C. We show that this ensemble reproduces exponentially high moments of the Haar measure. We can then derandomise the PFC ensemble to show the following:

- **Linear-depth t-designs.** We give the first construction of a (diamond-error) approximate t-design with circuit depth linear in t. This follows from the PFC ensemble by replacing the random phase and permutation operators with their $2t$-wise independent counterparts.
- **Non-adaptive PRUs.** We give the first construction of PRUs with non-adaptive security, i.e. we construct unitaries that are indistinguishable from Haar random to polynomial-time distinguishers that query the unitary in parallel on an arbitary state. This follows from the PFC ensemble by replacing the random phase and permutation operators with their pseudorandom counterparts.
- **Adaptive pseudorandom isometries.** We show that if one considers isometries (rather than unitaries) from n to $n + \omega(\log n)$ qubits, a small modification of our PRU construction achieves adaptive security, i.e. even a distinguisher that can query the isometry adaptively in sequence cannot distinguish it from Haar random isometries. This gives the first construction of adaptive pseudorandom isometries. Under an additional conjecture, this proof also extends to adaptive PRUs.

1 Introduction

The Haar measure formalises the notion of a uniformly random unitary and plays an important role in quantum information and computation. However, Haar random unitaries suffer from a significant practical drawback: they cannot be efficiently implemented. This has prompted research into ensembles of unitaries that "look" sufficiently Haar random for applications while also being efficient to implement.

The same situation arises in classical computer science: here, random functions have many useful properties, but cannot be efficiently implemented. There exist two different approaches to derandomising a fully random function: t-wise independent functions and pseudorandom functions.

t-wise independent functions are functions that information-theoretically reproduce the behaviour of uniformly random functions if one only evaluates the function t times. In other words, even with arbitrary computational power one cannot distinguish t queries to a random t-wise independent function from t queries to a uniformly random function. The drawback of this approach is that one needs to know the number of queries t ahead of time when constructing the functions: if a t-wise independent function is evaluated $t+1$ times, its output may look very different from that of a uniformly random function.

In contrast, pseudorandom functions are indistinguishable from uniformly random functions to any *polynomial-time*[1] distinguisher that can make an arbitrary polynomial number of queries to the function. The advantage over t-wise independent functions is that there is no a priori bound on the number of allowed queries; a pseudorandom function is secure against any polynomial number of queries. The drawback is that pseudorandom functions are only indistinguishable to computationally bounded distinguishers, whereas t-wise independent functions are information-theoretically secure.[2]

One encounters the same situation when trying to derandomise Haar random unitaries. The analogue to t-wise independent functions are called t-designs, which are random unitaries that reproduce the first t moments of the Haar measure [DCEL09, GAE07, AE07]. The analogue to pseudorandom functions are pseudorandom unitaries (PRUs), which are indistinguishable from Haar random unitaries to any polynomial-time quantum distinguisher with query access to the unitary [JLS18].

This paper makes three main contributions to the derandomisation of the Haar measure:

- **Linear-depth t-designs.** It has been a long-standing open question to construct optimal t-designs, i.e. t-designs with the smallest possible circuit complexity. We give the first construction of a (diamond-error) approximate t-design that only requires circuit depth linear in t (and polynomial in the number of qubits, which can be made quasilinear using upcoming work [CHH+24]). The best prior constructions required depth that scaled quadratically in t [CDX+24, HLT24].

- **Non-adaptive PRUs.** We give the first construction of PRUs with non-adaptive security, i.e. we construct unitaries that are computationally indistinguishable from Haar random to distinguishers that query the unitary *in parallel* on an arbitrary input state. Prior work on PRUs imposed severe limitations on the distinguisher, e.g. only allowing it to query the unitary on i.i.d. tensor product states [LQS+23].

- **Adaptive Pseudorandom Isometries (PRIs).** We show that if one considers isometries (rather than unitaries) from n to $n+\omega(\log n)$ qubits, a small modification of our PRU construction achieves adaptive security, i.e. even a distinguisher that can query the isometry adaptively in sequence cannot distinguish it from Haar random isometries. This gives the first construction of adaptive PRIs. Since the number of ancilla qubits introduced by the isometry is only $\omega(\log n)$, this construction may already be a good substitute for adaptive PRUs. Furthermore, we show that under an additional concrete conjecture our proof also extends to full adaptive PRUs.

Our constructions for t-designs and PRUs/PRIs are very simple and follow from a more general result: we show that the product of a random Clifford unitary, a random binary phase operator, and a random

[1] Whenever we refer to polynomial-time or polynomial numbers of queries, we mean polynomial in the input length n of the function.

[2] Note that one cannot efficiently construct t-wise independent functions where t is superpolynomial in n.

permutation approximates the Haar measure up to extremely high moments. The advantage of this random unitary over the Haar measure itself is that it can easily be derandomised: by replacing the random phase and random permutation with their t-wise independent counterparts, we get a linear-depth t-design; and by replacing them with their pseudorandom counterparts, we get our PRU/PRI construction.

We now give more detailed background on prior results on t-designs and pseudorandom unitaries, before returning to our construction and security proof in Section 1.1.

t-designs. A unitary t-design is an ensemble of unitaries that (approximately) reproduces the first t moments of the Haar measure [DCEL09, GAE07, AE07]. In contrast to Haar random unitaries, t-designs on n qubits can be implemented in polynomial time (in n and t), making them useful for applications ranging from randomised benchmarking [MGE12, KLR$^+$08] to quantum complexity growth [BCHJ$^+$21] and black hole physics [HP07].

Unitary t-design constructions are typically approximate, i.e., they only reproduce the moments of the Haar measure approximately. Various notions of approximation have been introduced (see e.g. [Mel23] for an overview). For our discussion, we need to consider two different notions of approximation: diamond-error and relative-error. We refer to Definitions 2.5 and 2.6 for the formal definitions. For our discussion here, it suffices to know that diamond-error corresponds to non-adaptive security (i.e. the t-design looks like a Haar random ensemble to a distinguisher making t parallel queries) and relative-error corresponds to (and is in fact stronger than) adaptive security (i.e. the t-design still looks Haar random even if the distinguisher is allowed to make t adaptive queries). This means that relative-error is a stronger notion, and converting from diamond-error to relative-error incurs a multiplicative $2^{O(nt)}$ blow-up in the error parameter. For our discussion below, we always require the approximation error (in either notion of approximation) to be negligible in n for any $t = \mathrm{poly}(n)$ and do not write it explicitly.

It is a long-standing open question to construct t-designs as efficiently as possible; by "efficient" we mean with a minimal number of quantum gates or circuit depth, but randomness-efficient constructions have also been studied [OSP23]. In a seminal result, Brandao, Harrow and Horodecki [BHH16] showed that random 2-local quantum circuits on n qubits form relative-error t-designs in depth $O(n^2 t^{10})$, which was improved to depth $O(nt^{5+o(1)})$ by Haferkamp [Haf22]. In addition to these results on random circuits, very recently two independent works used more structured circuits to achieve more efficient t-designs. Chen, Docter, Xu, Bouland, and Hayden [CDX$^+$24] achieve diamond-error t-designs in depth $\tilde{O}(n^2 t^2)$ using an intricate analysis of exponentiated Gaussian Unitary Ensembles. Haah, Liu, and Tan [HLT24] use random Pauli rotations to achieve relative-error t-designs in depth $\tilde{O}(nt^2)$.

Our construction is very simple and is the first to achieve linear scaling in t: we construct diamond-error t-designs with circuit depth $O(\mathrm{poly}(n)t)$. We can also amplify our construction to achieve *relative-error* t-designs with circuit depth $O(\mathrm{poly}(n)t^2)$. The exact polynomial dependence in n depends on the details of a construction of $O(t)$-wise independent permutations due to Kassabov [Kas07], which is analysed explicitly in an upcoming work [CHH$^+$24], showing that the depth can also be made quasilinear in n.

Pseudorandom unitaries. Pseudorandom unitaries (PRUs) are ensembles[3] $\{U_k\}_{k\in\mathcal{K}_n}$ of unitaries that are efficient to implement, but that look indistinguishable from Haar random unitaries to any polynomial-time distinguisher. This means that no polynomial-time distinguisher with oracle access to either a Haar random unitary or a unitary chosen uniformly from the PRU ensemble $\{U_k\}$ can tell the two cases apart. This makes PRUs the natural quantum analogue to pseudorandom functions (PRFs).

The concept of PRUs was introduced by Ji, Liu, and Song [JLS18]. Their paper gave a conjectured construction of PRUs, but only proved security (assuming quantum-secure one-way functions) for a much weaker primitive called *pseudorandom states* (PRSs). A pseudorandom state ensemble is a set of states (rather than unitaries) that look indistinguishable from Haar random states (even with access to polynomially many copies of the state). Since their introduction, PRSs have become an influential concept with applications in

[3] Strictly speaking, a PRU ensemble is an infinite sequence $\mathcal{U} = \{\mathcal{U}_n\}_{n\in\mathbb{N}}$ of n-qubit unitary ensembles $\mathcal{U}_n = \{U_k\}_{k\in\mathcal{K}}$, where $n \in \mathbb{N}$ serves as the security parameter (see Definition 5.1 for a formal statement).

quantum cryptography [AQY22, MY22], lower bounds in quantum learning theory [HBC+22], and even connections to quantum gravity [ABF+24]. However, proving security for a pseudorandom *unitary* construction has remained an open problem, and [HBK23] even explores restrictions on possible PRU constructions.

Given the difficulty of proving security for PRUs, recent work has considered intermediate steps between PRSs and PRUs. Lu, Qin, Song, Yao and Zhao [LQS+23] introduced the notion of a pseudorandom state scrambler (PRSS). A PRSS ensemble is a set of unitaries $\{U_k\}_{k\in\mathcal{K}}$ such that for all states $|\phi\rangle$, the state family $\{U_k|\phi\rangle\}_{k\in\mathcal{K}}$ is a PRS. In other words, a PRSS is a PRU that is only secure when queried on i.i.d. product states $|\phi\rangle^{\otimes\mathrm{poly}(n)}$. If we restrict the state $|\phi\rangle$ to the all-0 state $|0\rangle$, then we recover PRSs as a special case. [LQS+23] showed how to construct a PRSS ensemble (assuming quantum-secure one-way functions) by an intricate analysis of Kac's random walk. Ananth, Gulati, Kaleoglu, and Lin [AGKL23] define the notion of pseudorandom isometries (PRIs), which are like PRUs except that the operation can introduce extra qubits, i.e. it is an isometry, not a unitary. Towards constructing a PRI ensemble, [AGKL23] show that a candidate construction is computationally indistinguishable from Haar random isometries when applied to certain types of tensor product states (including tensor powers of the same state and i.i.d. Haar random states). [AGKL23] also explore the cryptographic applications of PRIs, which we briefly discuss in Section 1.3.

In this work, we give the first construction of PRUs with non-adaptive security and PRIs with adaptive security. For the non-adaptive PRUs, we allow the polynomial-time distinguisher to query the PRU polynomially many times *in parallel* on an *arbitrary entangled state* (rather than the restricted classes of input states considered in prior work). For the adaptive PRIs, we show that appending a $\omega(\log n)$-qubit auxiliary state to the input (which is what makes this construction an isometry) and applying our PRU construction achieves security even against adaptive adversaries, i.e. the ensemble of isometries looks computationally indistinguishable from Haar random isometries even to distinguishers that can perform arbitrary sequential adaptive queries to the isometry.

1.1 A unified approach to t-designs and PRUs: the PFC ensemble

t-designs and PRUs are intimately related and our work takes a unified approach to constructing both. The key insight of our proofs is that while the Haar measure is difficult to derandomise directly, we can construct a different ensemble of unitaries, which we call the PFC ensemble, that matches exponential-order moments of the Haar measure. In other words, this ensemble is a (diamond-error) t-design even when t is exponential in the number of qubits.

The PFC ensemble is not efficient to implement either, but it has a key advantage over the Haar measure: it is easy to derandomise, both in an information-theoretic manner to get t-designs and in a computational manner to get PRUs. As such, a significant portion of our paper is concerned with analysing the properties of the PFC ensemble itself, and the results on t-designs and (non-adaptive) PRUs follow straightforwardly from classical results on t-wise independence and pseudorandom functions.

The PFC ensemble is the following random ensemble of unitary matrices.

Definition 1.1 (PFC ensemble)**.** The n-qubit (or 2^n-dimensional) PFC ensemble is given by the product PFC where P is a uniformly random permutation matrix on n-qubit computational basis states, $F : |x\rangle \to (-1)^{f(x)}|x\rangle$ is a diagonal unitary with a uniformly random function $f : \{0,1\}^n \to \{0,1\}$, and C is a uniformly random n-qubit Clifford, all sampled independently.

Our main technical result says that the PFC ensemble is a t-design for exponential t with negligible diamond distance error. We refer to Section 1.2.1 for a proof sketch and to Theorem 3.1 for the formal statement and proof.

Theorem (Informal)**.** *For $d = 2^n$, the d-dimensional PFC ensemble is a diamond ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$. This means that for all t and all states $|\psi\rangle_{\mathsf{AE}}$, where A is an nt-qubit register (on which the unitary acts) and E is an arbitrary ancilla register,*

$$\left\| \mathop{\mathbb{E}}_{U\sim\mathrm{Haar}} U_{\mathsf{A}}^{\otimes t} |\psi\rangle\langle\psi| U_{\mathsf{A}}^{\otimes t,\dagger} - \mathop{\mathbb{E}}_{PFC} (PFC)_{\mathsf{A}}^{\otimes t} |\psi\rangle\langle\psi| (PFC)_{\mathsf{A}}^{\otimes t,\dagger} \right\|_1 \leq O(t/\sqrt{d})\,.$$

This means that even for exponential t, e.g. $t = 2^{n/4}$, the PFC ensemble is a t-design with exponentially small (in n) diamond distance error. These parameters turn out to be strong enough to give a unified construction of efficient t-designs and non-adaptive PRUs, as we will see next. Before doing so, we remark that the only property of a random Clifford C that we require is that it is an exact 2-design. As such, we can take C to be any exact 2-design (or even an approximate one with small enough error) in the construction given in Definition 1.1.

Information-theoretic derandomisation: t-designs with linear depth. While Section 1.1 shows that the PFC ensemble is a t-design with error $O(t/\sqrt{d})$, this result is not immediately useful as it is not possible to efficiently implement a uniformly random phase operator F or a uniformly random permutation P. However, we can replace the (classical) uniform functions and permutations underlying the phase and permutation operators by their $O(t)$-wise independent variants (for a formal definition of t-wise independence, see Section 4) and show that the resulting random unitary remains a t-design (see Theorem 4.4).

Theorem (Efficient unitary t-design, informal)**.** *Let $d = 2^n$ and let ν be the n-qubit PFC ensemble where F is instantiated with a random $2t$-wise independent Boolean function, P is instantiated with a random (approximately) t-wise independent permutation, and C is a random n-qubit Clifford. Then, ν is a diamond error ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$.*

By using efficient constructions of $2t$-wise independent functions and approximate t-wise independent permutations [Kas07, AL13, CK23, CHH$^+$24], we obtain t-designs with circuit size and depth $O(t\,\mathrm{poly}(n))$. By repeating the construction multiple times in sequence, we can make the error ϵ decay exponentially and also obtain relative-error t-designs with circuit size and depth $O(t^2 \cdot \mathrm{poly}(n))$. More precisely, we get the following (see Corollaries 4.9 and 4.12 for the formal statement):

Corollary (informal)**.** *For any number of qubits n, tolerated error $\epsilon > 0$, and $t \leq 2^{n/4}$, there exists a diamond ϵ-approximate t-design with circuit size and depth $O(t \cdot \mathrm{poly}(n) + t \log 1/\epsilon)$ and a relative-error ϵ-approximate t-design with circuit size and depth $O(t^2 \cdot \mathrm{poly}(n) + t^2 \log 1/\epsilon)$.*

In fact, as mentioned before, the upcoming work [CHH$^+$24] gives explicit circuits for the $O(t)$-wise independent permutations from [Kas07] that also achieve quasilinear depth in n, which implies t-designs in depth $\tilde{O}(tn)$ (for diamond distance error $\epsilon = e^{-\Omega(n)}$).

Computational derandomisation: non-adaptive pseudorandom unitaries. Replacing the uniformly random function and permutation in the PFC ensemble with their pseudorandom counterparts, we immediately get a PRU ensemble with non-adaptive security. More concretely, we require:

- An ensemble of (quantum-secure) pseudorandom permutations (PRPs) [Zha16]. Broadly speaking, this is a family $\{\pi_{k_1} : \{0,1\}^n \to \{0,1\}^n\}_{k_1 \in \mathcal{K}_1}$ of permutations with the property that, for a randomly chosen key $k_1 \sim \mathcal{K}_1$, the permutation π_{k_1} is computationally indistinguishable from a perfectly random permutation. For a given π_{k_1}, we let P_{k_1} be the corresponding n-qubit permutation matrix.
- An ensemble of (quantum-secure) pseudorandom functions (PRFs) [Zha21]. More formally, this is a family $\{f_{k_2} : \{0,1\}^n \to \{0,1\}\}_{k_2 \in \mathcal{K}_2}$ with the property that, for a randomly chosen key $k_2 \sim \mathcal{K}_2$, the function f_{k_2} is computationally indistinguishable from a perfectly random Boolean function. For a given f_{k_2}, we let $F_{k_2} : |x\rangle \to (-1)^{f_{k_2}(x)} |x\rangle$ be the n-qubit phase oracle implementing f_{k_2}.

Then, indexing the n-qubit Clifford group as $\{C_{k_3}\}_{k_3 \in \mathcal{K}_3}$, our PRU is the n-qubit ensemble $\{U_k\}_{k \in \mathcal{K}}$ which is specified by a key $k = (k_1, k_2, k_3) \in \mathcal{K}_1 \times \mathcal{K}_3 \times \mathcal{K}_3 =: \mathcal{K}$, where

$$U_k = P_{k_1} F_{k_2} C_{k_3}. \tag{1.1}$$

Note that, the Clifford group on n-qubits has size $2^{O(n^2)}$. Thus assuming the PRP and PRF scheme both have a key space consisting of $\mathcal{K}_1 = \mathcal{K}_2 = \{0,1\}^n$, our PRU ensemble has a key length[4] of $|k| = n + n + O(n^2) =$

[4] We note that one can apply a classical pseudorandom random generator to the key to reduce the key length.

$O(n^2)$, where $n \in \mathbb{N}$ is the security parameter. Because an ensemble of (quantum-secure) PRFs and PRPs can efficiently be constructed from (quantum-secure) one-way functions [Zha16, Zha21], the security of our PRU ensemble also only relies on the existence of quantum-secure one-way functions. Furthermore, since random Cliffords can be efficiently sampled [vdB21], there is a polynomial-time quantum algorithm that implements the PRU.

From Section 1.1, it is not hard to show that the PRU ensemble is secure against any distinguisher that makes polynomially many *parallel* queries to the PRU, i.e. the distinguisher is allowed a single query to $U^{\otimes \mathrm{poly}(n)}$ on *any input state*, rather than just the restricted classes of input states allowed in PRSs and PRSSs. We give the the formal statement and proof in Section 5.

Theorem (Non-adaptive PRUs, informal)**.** *Assuming the existence of quantum-secure one-way functions, the ensemble described in Equation (1.1) satisfies non-adaptive PRU security.*

Pseudorandom isometries with adaptive security. We conjecture that our PRU construction is also adaptively secure, but so far we are not able to prove adaptive security for the case of unitaries. However, we can show adaptive security under a small relaxation: if instead of PRUs, we consider PRIs that map n qubits to $n + \omega(\log n)$ qubits, then we can show adaptive security. Our PRI construction is the same as our PRU construction, with two minor differences: we fix part of the input state to the $|+\rangle$-state (which is what makes this an isometry) and we do not require the random Clifford.

Definition 1.2 (PRI construction)**.** Let $\{P_{k_1}\}_{k_1 \in \mathcal{K}_1}$ and $\{F_{k_2}\}_{k_2 \in \mathcal{K}_2}$ be the pseudorandom permutation and binary phase operators on n qubits from Eq. (1.1). Choose any function $s(n) = \omega(\log n)$. We then define an ensemble of isometries $\{V_k\}_{k \in \mathcal{K}_1 \times \mathcal{K}_2}$ from $n - s(n)$ to n qubits given by

$$V_{k_1,k_2} |\psi\rangle = P_{k_1} F_{k_2} (|\psi\rangle \otimes |+\rangle^{\otimes s(n)}) \,.$$

We show that no polynomial-time quantum algorithm can distinguish such isometries from a Haar random isometry (see Section 6.1 for a definition). We prove this formally in Section 6 and sketch the proof in Section 1.2.2.

Theorem (Adaptive PRIs, informal)**.** *Assuming the existence of quantum-secure one-way functions, the isometries defined in Definition 1.2 are pseudorandom isometries with adaptive security.*

Since the number of additional output qubits $\omega(\log n)$ is quite small, it seems likely that this PRI construction already suffices in many situations where one would like to use adaptively secure PRUs. Furthermore, we can also extend the proof of PRI security to full adaptive PRU security assuming a concrete conjecture (see Section 6.1).

Note that when written as matrices, the isometries V_k from Section 1.1 have only real entries. In contrast, it is known that PRUs with only real entries cannot exist [HBK23]. The fact that (even adaptively-secure) PRIs with real entries are possible may seem surprising at first sight. However, we note that already in our PRU construction, the only source of complex numbers were the random Cliffords, and their only purpose in our construction is to ensure that the initial input state is sufficiently scrambled so that it has a large overlap with a certain subspace that we call the *distinct subspace* (see Section 1.2). Our analysis for Theorem 5.2 in fact shows that on the distinct subspace, the (real-valued) PF-ensemble (with pseudorandom permutations and functions) forms a non-adaptive PRU. This is also closely related to the very recent work of Brakerski and Magrafta [BM24], who construct real-valued unitaries that look Haar random on any polynomial-sized set of orthogonal input states.

1.2 Proof Overview

We break this section into two parts: the first focuses on analyzing the PFC ensemble and the second focuses on adaptive security for PRIs.

1.2.1 t-design property of the PFC ensemble

As mentioned before, showing that the PFC ensemble is a diamond distance t-design corresponds to showing that any non-adaptive algorithm (possibly inefficient) that makes t queries cannot distinguish the PFC ensemble from a Haar random ensemble. Such an algorithm starts with an initial state $|\psi\rangle_{\mathsf{A}_1\cdots\mathsf{A}_t\mathsf{B}}$ where registers $\mathsf{A}_1,\ldots,\mathsf{A}_t$ are each on n qubits, and B is an arbitrary workspace register. The algorithm then applies the unitary $U^{\otimes t}$ on the registers $\mathsf{A}_1,\cdots,\mathsf{A}_t$ (where U is either PFC or a Haar random matrix) and performs a measurement afterwards. In order to show that no such algorithm can distinguish the PFC ensemble from a Haar random ensemble, it suffices to show that for all initial states $|\psi\rangle_{\mathsf{A}_1\cdots\mathsf{A}_t\mathsf{B}}$, the following two density matrices are close in trace distance:

$$\mathop{\mathbb{E}}_{PFC}((PFC)^{\otimes t}\otimes\mathbb{1})\,|\psi\rangle\!\langle\psi|\,((PFC)^{\otimes t}\otimes\mathbb{1})^\dagger \approx \mathop{\mathbb{E}}_{U\sim\mathrm{Haar}}(U^{\otimes t}\otimes\mathbb{1})\,|\psi\rangle\!\langle\psi|\,(U^{\otimes t}\otimes\mathbb{1})^\dagger\ . \tag{1.2}$$

Here, $|\psi\rangle$ is a completely arbitrary state: the registers $\mathsf{A}_1,\ldots,\mathsf{A}_t$ of $|\psi\rangle$ may all be entangled with each other, the reduced states on each register may be different from each other, and finally the distinguisher is allowed access to the purification of the input state to the unitaries.

For this proof overview, we assume that our algorithm has no workspace, i.e. the initial state is $|\psi\rangle_{\mathsf{A}_1\cdots\mathsf{A}_t}$. This is merely to simplify the notation, the proof works exactly in the same way even in the case of an entangled workspace. It then suffices to show that the following two density matrices are close in trace distance:

$$\mathop{\mathbb{E}}_{PFC}(PFC)^{\otimes t}\,|\psi\rangle\!\langle\psi|\,((PFC)^{\otimes t})^\dagger \approx \mathop{\mathbb{E}}_{U\sim\mathrm{Haar}}U^{\otimes t}\,|\psi\rangle\!\langle\psi|\,(U^{\otimes t})^\dagger\ . \tag{1.3}$$

For this, we leverage Schur-Weyl duality to compute what these two density matrices explicitly look like. Let $d=2^n$ be the dimension. According to Schur-Weyl duality:

(1) the space $(\mathbb{C}^d)^{\otimes t}$ can be decomposed as a direct sum[5] $\bigoplus_{\lambda\vdash t}P_\lambda$ where $P_\lambda=W_\lambda\otimes V_\lambda$ is a tensor product of two spaces, and

(2) any unitary $U^{\otimes t}$ only acts non-trivially on the subspaces W_λ and any unitary R_π that permutes the t subsystems according to a permutation $\pi\in S_t$ acts non-trivially only on the subspaces V_λ.

Using this, we show that applying a *t-wise Haar twirl* to $|\psi\rangle$ to obtain the state on the right hand side of Eq. (1.3) results in the following state:

$$\mathop{\mathbb{E}}_{U\sim\mathrm{Haar}}U^{\otimes t}\,|\psi\rangle\!\langle\psi|\,(U^{\otimes t})^\dagger=\sum_{\lambda\vdash t}\frac{\mathbb{1}_{W_\lambda}}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}\otimes\mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}\,|\psi\rangle\!\langle\psi|\,\mathbb{1}_{P_\lambda}]\,,$$

where $\mathbb{1}_{W_\lambda}$ is the identity on the subspace W_λ, i.e. the orthogonal projection onto that subspace. In particular, the state is a direct sum of tensor product states where the state on the subspaces W_λ is maximally mixed.

Next, we would like to show that applying a *t-wise PFC twirl* to $|\psi\rangle$ to obtain the left hand side state in Eq. (1.3) results in a state that is close to the above. For technical reasons, the computations here are easier if the state $|\psi\rangle$ is supported only on the subspace of distinct computational basis states in the registers $\mathsf{A}_1,\ldots\mathsf{A}_t$, i.e., the subspace spanned by $|x_1,\ldots,x_t\rangle$ where $x_1,\ldots,x_t\in\{0,1\}^n$ are all distinct. We show that applying a random Clifford ensures this by showing that

$$\mathrm{Tr}\left[\Lambda\mathop{\mathbb{E}}_{C}C^{\otimes t}\,|\psi\rangle\!\langle\psi|\,C^{\otimes t,\dagger}\right]\geq 1-O(t^2/d)\,,$$

where Λ denotes the projector on the *distinct subspace*. In other words, after applying the random Clifford on the arbitrary input state $|\psi\rangle$, the resulting state has so little weight on the non-distinct subspace that we only need to deal with the distinct subspace.

[5] Here $\lambda\vdash t$ means that λ is a partition of $[t]$.

For states $|\psi\rangle$ in the distinct subspace, we show that a t-wise PF twirl results in

$$\mathop{\mathbb{E}}_{PF}(PF)^{\otimes t}|\psi\rangle\!\langle\psi|\left((PF)^{\otimes t}\right)^{\dagger} = \sum_{\lambda\vdash t}\frac{\mathbb{1}_{\Lambda_\lambda}}{\mathrm{Tr}[\mathbb{1}_{\Lambda_\lambda}]}\otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}|\psi\rangle\!\langle\psi|\,\mathbb{1}_{P_\lambda}]\,,$$

where Λ_λ is a subspace of W_λ that comes from decomposing the distinct subspace projector Λ in terms of the Schur-Weyl subspaces. We show that Λ_λ fills most of W_λ, i.e. the dimension of the subspace Λ_λ is close to the dimension of W_λ:

$$\frac{\mathrm{Tr}[\mathbb{1}_{\Lambda_\lambda}]}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]} = 1 - O\left(\frac{t^2}{d}\right).$$

This implies that the mixed state on Λ_λ is close in trace distance to the maximally mixed state on W_λ:

$$\left\|\frac{\mathbb{1}_{\Lambda_\lambda}}{\mathrm{Tr}[\mathbb{1}_{\Lambda_\lambda}]} - \frac{\mathbb{1}_{W_\lambda}}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}\right\|_1 = O\Big(\frac{t^2}{d}\Big).$$

Since this is true for each λ, putting all the above together, we get that the left and right hand sides in Eq. (1.3) have trace distance[6] at most $\sqrt{t}/d$, which is exponentially small since $t = \mathrm{poly}(n)$ and $d = 2^n$.

1.2.2 Pseudorandom isometries with adaptive security

Before explaining how we can construct adaptively secure PRIs, let us briefly discuss the obstacles to extending the proof of Theorem 3.1 to adaptive security. If we were able to prove that the PFC ensemble is a *relative-error* design, adaptive security for our PRU construction would follow using known techniques (see e.g. [Kre21, Section 3]). In particular, one can convert an adaptive algorithm that makes t-queries to a non-adaptive t-query algorithm with post-selection. Relative error designs allow one to obtain multiplicative approximations to the acceptance probabilities, allowing the whole argument to go through.

However, we can only prove that the PFC ensemble is a *diamond-error* t-design, and diamond error only gives additive approximations which are too weak to make this post-selection argument work. As outlined in first part of the proof overview, to prove that the PFC ensemble is a superpolynomial design, we first use the random Clifford unitary to restrict our attention to the distinct-string subspace, and then analyse the PF twirl on the distinct subspace. It turns out that a small modification of our proof of Theorem 5.2 can show that the PF twirl behaves like a (one-sided) *relative-error* design *on the distinct subspace* (Lemma 6.3). The problem is that the first step, the reduction to the distinct subspace, does not seem to work in the relative-error setting. If one were able to perform the reduction to the distinct subspace for adaptive algorithms as conjectured in Conjecture 6.5, adaptive PRU security would follow.

Fortunately the relative-error design property of the PF ensemble on the distinct subspace is already useful by itself. In the usual conversion from relative-error designs to adaptive security, one starts from the resource state $U^{\otimes t}|\Omega\rangle\!\langle\Omega|\,U^{\otimes t,\dagger}$, where $|\Omega\rangle$ is the maximally entangled state between $(\mathbb{C}^d)^{\otimes t}\otimes(\mathbb{C}^d)^{\otimes t}$, and uses gate teleportation on this resource state to simulate adaptive queries to the unitary U. This gate teleportation trick essentially reduces adaptive to non-adaptive security, since we now only need to analyse the application of U in parallel on the resource state. However, the gate teleportation trick introduces a large "post-selection factor", which can only be handled with relative-error designs.

Since we have a relative-error design (the PF ensemble) on the distinct subspace, a natural idea is to perform the same gate teleportation trick, except using $U^{\otimes t}|\Omega_\Lambda\rangle\!\langle\Omega_\Lambda|\,U^{\otimes t,\dagger}$ as the resource state. Here, $|\Omega_\Lambda\rangle$ is the unnormalised maximally entangled state on the distinct subspace of $(\mathbb{C}^d)^{\otimes t}$, i.e.

$$|\Omega_\Lambda\rangle = \sum_{x_1,\dots,x_t\in[d]\text{ distinct}} |x_1,\dots,x_t\rangle\,|x_1,\dots,x_t\rangle\,.$$

Of course $U^{\otimes t}|\Omega_\Lambda\rangle\!\langle\Omega_\Lambda|\,U^{\otimes t,\dagger}$ is not the "correct" resource state for performing gate teleportation; consequently, the resulting state is not simply the state with U applied adaptively in the right places (which

[6] We lose a square-root factor in the analysis because of an application of the gentle measurement lemma.

is what we get when using the correct resource state $U^{\otimes t} |\Omega\rangle\langle\Omega| U^{\otimes t,\dagger}$). Instead, we get a slightly different state.

The challenge then is to show that the outputs of gate teleportation with the "correct" and "distinct" resource states are close. This is the step that we are not able to prove for the general case of unitaries (although it seems very plausible that it can be proven using similar techniques). However, if we fix $\omega(\log n)$ of the n qubits of the input to the unitary, then we are able to show that the outputs of the gate teleportation for the two different resource states $U^{\otimes t} |\Omega\rangle\langle\Omega| U^{\otimes t,\dagger}$ and $U^{\otimes t} |\Omega_\Lambda\rangle\langle\Omega_\Lambda| U^{\otimes t,\dagger}$ are indeed close. Fixing parts of the input to a unitary (to some universal state independent from the other input – in our case the fixed input qubits are simply in a $|+\rangle$-state) turns the unitary into an isometry. This is how we achieve PRIs with adaptive security.

Seeing how fixing part of the input to the unitary is helpful requires an explicit expression for the output state of the gate teleportation, so we defer this discussion to Section 6. We also remark that it seems likely that this proof strategy can be extended to adaptive PRUs: the only reason we need to relax our construction to PRIs is in order to ensure that the adaptive queries are in some adaptive version of the distinct "subspace". If one could instead use query complexity arguments to prove this property for some unitary ensemble, one could simply replace our random Clifford unitary by this ensemble and achieve adaptively secure PRUs. We formalise this idea in Section 6.1.

1.3 Discussion and future directions

In this work we have taken a unified approach to constructing t-designs and PRUs by introducing the PFC ensemble and showing that it mimicks the Haar measure extremely well. Derandomising the PFC ensemble with existing classical primitives such as t-wise independent functions or pseudorandom functions yields surprisingly simple constructions of diamond-error t-designs with circuit depth linear in t and PRUs with non-adaptive security.

There are several ways in which one would like to improve upon these results and we hope that our techniques can be extended accordingly. Our t-design construction only achieves small diamond error (corresponding to non-adaptive security) with linear scaling in t, but requires quadratic scaling in t to achieve small relative error. Constructing a relative-error t-design with linear scaling in t remains an interesting open problem. It may well be that the PFC ensemble is a relative-error t-design with linear scaling in t, but our current analysis does not show this because we do not know how to analyse the PFC ensemble on the non-distinct subspace. We can extend the analysis to a subspace with a constant number of collisions, but going beyond that seems to require new ideas. One can also optimise the dependence in the number of qubits n; for this, we refer to [CHH+24].

Similarly, our PRU construction only achieves non-adaptive security, and in contrast to t-designs we cannot "amplify" the PRU construction to relative/adaptive security because that amplification requires t to be known ahead of time. Looking forward, there are three directions in which one would like to extend the security guarantee of our PRU construction: allowing adaptive queries, allowing inverse queries, and allowing controlled queries to the unitary. Our construction plausibly has adaptive security and we discuss an approach to extending our current proof to the adaptive setting in Section 6.1. In contrast, the PFC ensemble is *not* indistinguishable from Haar random unitaries with inverse queries: this is because applying $C^\dagger F^\dagger P^\dagger$ to $|0\rangle$ always yields a stabiliser state, but applying a Haar random unitary to $|0\rangle$ does not, and this difference can be tested efficiently. However, if one simply adds another independent Clifford at the end (i.e. considers $C'PFC$), the construction is plausibly secure against inverse queries, but we do not know how to analyse this. Proving security with controlled access to the unitary is similarly unclear.

Aside from constructing PRUs (or PRIs), their applications are also largely unexplored. Given the utility of PRFs in classical computer science, PRUs appear to be a fundamental primitive for quantum computer science, but relatively few concrete applications have been proposed. Below we briefly discuss some applications that have been mentioned in the literature and suggest some new ones.

One natural area of application is quantum cryptography. For example, [LQS+23, AGKL23] showed that PRSSs and PRIs can be useful for multi-copy quantum cryptography: in most encryption schemes for quantum messages (e.g., the quantum one-time pad and its variants), if we wish to encrypt multiple identical copies of the same quantum state, we need to sample fresh keys for each new ciphertext. [LQS$^+$23, AGKL23] observed that Haar (pseudo)randomness allows one to perform such a task in a compact manner with only a single key (for an arbitrary polynomial amount of identical copies). If the state to be encrypted is guaranteed to be unentangled with the environment, the PRSSs and PRIs from [LQS$^+$23, AGKL23] suffice; in the general case which allows for entanglement with an auxiliary system, non-adaptively PRUs as constructed in our work seem necessary. We remark that [AGKL23] also use PRIs (rather than PRUs) to construct succinct quantum state commitments, many-copy unforgeable quantum message authentication schemes, as well as to increase the length of pseudorandom quantum states. As another example, PRUs might be useful in the context of unclonable cryptography. Many constructions, such as those for unclonable encryption [BL20] or quantum copy-protection [CMP20, CLLZ21], make use of either Wiesner states or subspace coset states—both of which are completely broken once identical copies become available. It seems plausible that one could use PRUs to construct multi-copy secure unclonable encryption schemes, and even multi-copy secure quantum copy-protection schemes.

Another area of application concerns the time evolution of chaotic quantum systems. In the past few years, a series of works has proposed using Haar random unitaries as "perfect scramblers" [HP07, BF12, Sus16] to model such dynamics. However, a more recent line of work has instead shifted towards quantum pseudorandomness in order to model such phenomena in terms of *efficient* processes. For example, Kim and Preskill [KP23] use PRUs to model the internal dynamics of a black hole, whereas Engelhardt et al. [EFL$^+$24] use PRUs to model the time evolution operator of a holographic conformal field theory. Due to their scrambling properties, it is conceivable that PRUs will find more applications in theoretical physics.

Related independent work. After the initial announcement of our results on pseudorandom unitaries on the arXiv [MPSY24] (which this present paper supersedes), we were made aware of independent work by Chen, Bouland, Brandao, Docter, Hayden, and Xu [CBB$^+$24], who achieve similar results using a different ensemble of unitaries. In another related independent work, Brakerski and Magrafta [BM24] construct real-valued unitaries that look Haar random on any polynomial-sized set of orthogonal input states; they use an ensemble that consists of the product of a Hadamard, a random permutation, and a binary phase operator.

Acknowledgments. We thank Prabhanjan Ananth, Adam Bouland, Chi-Fang Chen, Tudor Giurgica-Tiron, Jonas Haferkamp, Robert Huang, Isaac Kim, Fermi Ma, Xinyu Tan, and John Wright for helpful discussions. We also thank John for suggesting a simplified proof of Lemma 3.9 and Jonas for pointing out a proof of Lemma 4.11. We thank the Simons Institute for the Theory of Computing, where some of this work was conducted. TM acknowledges support from the ETH Zurich Quantum Center, the SNSF QuantERA project (grant 20QT21_187724), the AFOSR grant FA9550-19-1-0202, and an ETH Doc.Mobility Fellowship. AP is supported by the National Science Foundation (NSF) under Grant No. CCF-1729369. MS acknowledges support from the NSF award QCIS-FF: Quantum Computing & Information Science Faculty Fellow at the University of Illinois Urbana-Champaign (NSF 1955032). HY is supported by AFOSR awards FA9550-21-1-0040 and FA9550-23-1-0363, NSF CAREER award CCF-2144219, NSF award CCF-2329939, and the Sloan Foundation.

2 Preliminaries

2.1 Notation and basic definitions

For $N \in \mathbb{N}$, we use $[N] = \{1, 2, \ldots, N\}$ to denote the set of integers up to N. We oftentimes identify elements $x \in [N]$ with bit strings $x \in \{0,1\}^n$ via their binary representation whenever $N = 2^n$ and $n \in \mathbb{N}$. We use $\delta_{i,j}$ to denote the delta function which is 1 iff $i = j$ and 0 otherwise. The notation $x \sim X$, for a set X, describes that an element x is drawn uniformly at random from X. Similarly, if $\mathcal{D}$ is a distribution, we let $x \sim \mathcal{D}$ denote that x is sampled according to $\mathcal{D}$. The statistical distance between two distributions $\mathcal{D}$ and $\mathcal{D}'$ over a set X is defined as $\|\mathcal{D} - \mathcal{D}'\|_1 = \sum_{x \in X} |\mathcal{D}(x) - \mathcal{D}'(x)|$.

Linear maps. For a Hilbert space $\mathcal{H}$, we denote by $\mathrm{L}(\mathcal{H})$ linear operators on $\mathcal{H}$. A map $\mathcal{M} : \mathrm{L}(\mathcal{H}) \to \mathrm{L}(\mathcal{H})$ is called a quantum channel if it is completely positive and trace-preserving. If $\mathcal{H}'$ is an additional Hilbert space and $O \in \mathrm{L}(\mathcal{H} \otimes \mathcal{H}')$ is an operator on a larger space, we write $\mathcal{M}(O)$ to mean $\mathcal{M}$ applied to the $\mathcal{H}$-subsystem, i.e. $\mathcal{M}(O)$ is shorthand for $(\mathcal{M} \otimes \mathbb{1}_{\mathcal{H}'})(O)$ where $\mathbb{1}_{\mathcal{H}'} : \mathrm{L}(\mathcal{H}') \to \mathrm{L}(\mathcal{H}')$ denotes the identity map. For a subspace W of a vector space V, we also denote by $\mathbb{1}_W$ the orthogonal projection onto W (i.e. the identity on the subspace W). We use the notation $\mathrm{U}(d)$ and $\mathrm{C}(d)$ to denote the d-dimensional unitary group and Clifford group, respectively.

Distance measures. For a linear operator $A \in \mathrm{L}(\mathcal{H})$, we use $\|A\|_1 = \mathrm{Tr}\big[(A^\dagger A)^{1/2}\big]$ to denote the Schatten 1-norm (also called trace norm) and $\|A\|_\infty$ to denote the Schatten ∞-norm (also called operator norm), which is defined to be the largest singular value of A. The trace distance between two operators $A, B \in \mathrm{L}(\mathcal{H})$ is defined as $\|A - B\|_1$.

For two quantum channels $\mathcal{M}, \mathcal{N} : \mathrm{L}(\mathcal{H}) \to \mathrm{L}(\mathcal{H}')$, we define the diamond distance

$$\|\mathcal{M} - \mathcal{N}\|_\diamondsuit = \max_{|\psi\rangle \in \mathcal{H} \otimes \mathcal{H}} \|(\mathcal{M} \otimes \mathbb{1}_{\mathcal{H}})(|\psi\rangle\langle\psi|) - (\mathcal{N} \otimes \mathbb{1}_{\mathcal{H}})(|\psi\rangle\langle\psi|)\|_1 ,$$

where $\mathbb{1}_{\mathcal{H}} : \mathrm{L}(\mathcal{H}) \to \mathrm{L}(\mathcal{H})$ denotes the identity channel. The diamond norm is submultiplicative, i.e. for two non necessarily completely positive superoperators $\mathcal{E}, \mathcal{F}$ (e.g. $\mathcal{E} = \mathcal{F} = \mathcal{M} - \mathcal{N}$ could be the difference between channels $\mathcal{M}$ and $\mathcal{N}$), it holds that

$$\|\mathcal{E} \circ \mathcal{F}\|_\diamondsuit \leq \|\mathcal{E}\|_\diamondsuit \cdot \|\mathcal{F}\|_\diamondsuit . \tag{2.1}$$

Distinct tuples. We use bold faced fonts to denote tuples. For a tuple $\boldsymbol{x} = (x_1, \ldots, x_t) \in [d]^t$ and a permutation $\sigma \in S_t$, we write $\boldsymbol{x}_\sigma = (x_{\sigma(1)}, \cdots, x_{\sigma(t)})$ for the tuple where the indices are permuted according to σ. We call a tuple $\boldsymbol{x} \in [d]^t$ distinct if $x_i \neq x_j$ for all $i \neq j$. We denote the set of distinct tuples in $[d]^t$ by $\mathrm{distinct}(d, t)$. We also define the projector onto the subspace of distinct tuples

$$\Lambda_{d,t} = \sum_{\boldsymbol{x} \in \mathrm{distinct}(d,t)} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| . \tag{2.2}$$

We will frequently drop the indices d, t and just write Λ, with d and t being clear from context.

In the proof of adaptive security for our PRI construction, we will also make use of the maximally entangled state between two copies of the distict subspace, which we denote by

$$|\Omega_{\Lambda_{d,t}}\rangle = \frac{1}{\sqrt{|\mathrm{distinct}(d,t)|}} \sum_{\boldsymbol{x} \in \mathrm{distinct}(d,t)} |\boldsymbol{x}\rangle |\boldsymbol{x}\rangle .$$

As for the projector onto the distinct string subspace, we will frequently drop the indices d, t and simply write $|\Omega_\Lambda\rangle$

Permutation operators. We define the following permutation operator on $\mathbb{C}^d$.

Definition 2.1 (Permutation operator on $\mathbb{C}^d$)**.** Define the permutation operator P_π on $\mathbb{C}^d$ for $\pi \in S_d$ to be the linear map

$$P_\pi : |x\rangle \mapsto |\pi(x)\rangle . \tag{2.3}$$

We will frequently consider uniformly random permutation operators on $\mathbb{C}^d$. We will suppress the dependence on π and write the random operator as P.

Symmetric group and representations. Unitary representations of a group allow us to represent the elements of the group as unitary matrices over a vector space in a way that the group operation is represented by matrix multiplication. We consider the following representation of the symmetric group which permutes the tensor factors.

Lemma 2.2 (Representation of S_t on tensor product spaces)**.** *For any fixed d, define the permutation operator R_π on $(\mathbb{C}^d)^{\otimes t}$ for $\pi \in S_t$ to be the map*

$$R_\pi : |\boldsymbol{a}\rangle \mapsto |\boldsymbol{a}_{\pi^{-1}}\rangle \ .$$

Then $((\mathbb{C}^d)^{\otimes t}, R_\pi$ forms a unitary representation of S_t. Note that we leave the dependence of R_π on the choice of d implicit.

We note that the representation $((\mathbb{C}^d)^{\otimes t}, R_{(\cdot)})$ can be decomposed into (isotypic) copies of the irreducible representations (or irreps) of the symmetric group S_t, which we denote by $\{(V_\lambda, R^\lambda_{(\cdot)})\}_\lambda$, where $\lambda \vdash t$ is a partition of t (or Young diagram with at most t boxes) and V_λ are vector spaces called Specht modules.

Binary phase operators. For a function $f : [d] \to \{0, 1\}$, we define the binary phase operator

$$F_f : |x\rangle \mapsto (-1)^{f(x)} |x\rangle \, . \tag{2.4}$$

We will frequently consider uniformly random binary phase operators, which we will just write as F.

2.2 Haar measure and unitary designs

We recall the definition of the Haar measure and t-wise twirl.

Definition 2.3 (Haar measure)**.** The Haar measure is the unique left- and right-invariant probability measure on the unitary group $\mathrm{U}(d)$. Throughout this paper, we denote sampling from the Haar measure over $\mathrm{U}(d)$ by $U \sim \mathrm{Haar}(d)$. If the dimension d is clear from the context, we simply write $U \sim \mathrm{Haar}$.

Definition 2.4 (t-wise twirl)**.** Let ν be an ensemble of unitary operators in $\mathrm{U}(d)$. Then, the t-wise twirl (also called t-th moment operator) with respect to ν is defined as the operator

$$\mathcal{M}^{(t)}_\nu(\cdot) = \mathop{\mathbb{E}}_{R\sim\nu} R^{\otimes t}(\cdot) R^{\otimes t,\dagger} \, .$$

If R is a random unitary matrix whose distribution is clear in context, we oftentimes use the shorthand notation $\mathcal{M}^{(t)}_R(\cdot)$. For example, we frequently consider the t-wise $\mathcal{M}^{(t)}_{PF}(\cdot)$ twirl, where $R = PF$ is a product of a random permutation operator P and a binary phase operator F, which are sampled independently. We call this the permutation-phase twirl. Similarly, we use $\mathcal{M}^{(t)}_{PFC}(\cdot)$ to denote the "PFC-twirl", where $R = PFC$ for P and F as before and C a uniformly random Clifford unitary. Finally, we denote by $\mathcal{M}^{(t)}_{\mathrm{Haar}}(\cdot)$ the t-wise twirl over unitaries which are sampled according to the Haar measure.

Just like t-wise independent functions "simulate" uniformly random functions if one only considers t-th moments, there is a notion of simulating the Haar measure up to the t-th moment, called a *t-design.* Various notions of approximation exist for t-designs and we refer to [Mel23] for an overview. Here, we will use the following two notions of approximate t-designs.

Definition 2.5 (Diamond ϵ-approximate t-design)**.** An ensemble ν of unitary operators in $\mathrm{U}(d)$ is called a diamond ϵ-approximate t-design, if

$$\left\| \mathcal{M}^{(t)}_\nu - \mathcal{M}^{(t)}_{\mathrm{Haar}} \right\|_\diamond \leq \epsilon.$$

Definition 2.6 (Relative-error ϵ-approximate t-design)**.** An ensemble ν of unitaries acting on a Hilbert space $\mathcal{H}$ is a relative-error ϵ-approximate t-design if, for all states $|\psi\rangle \in \mathcal{H} \otimes \mathcal{H}$, the following operator inequality holds

$$(1-\epsilon)\mathcal{M}_{\nu}^{(t)}(|\psi\rangle\!\langle\psi|) \leq \mathcal{M}_{\mathrm{Haar}}^{(t)}(|\psi\rangle\!\langle\psi|) \leq (1+\epsilon)\mathcal{M}_{\nu}^{(t)}(|\psi\rangle\!\langle\psi|)\,.$$

We call a random unitary a left-one-sided or right-one-sided relative-error ϵ-approximate t-design if only the left or right side of these inequalities holds.

Every diamond error t-design is also a relative-error t-design, but this conversion incurs a large blowup in the error.

Lemma 2.7 ([BHH16, Lemma 3])**.** *Suppose a random unitary R in d dimensions is a diamond ϵ-approximate t-design. Then R is also a relative-error $(\epsilon \cdot d^{2t})$-approximate t-design.*

2.3 Schur-Weyl duality

Consider the representation R_π of the symmetric group S_t and the representation $U^{\otimes t}$ of the unitary group $\mathrm{U}(d)$ both on the vector space $(\mathbb{C}^d)^{\otimes t}$. Schur-Weyl duality says that the irreducible subrepresentations of both these representations fit together nicely.

Lemma 2.8 (Schur-Weyl duality, see e.g. [Chr06, Theorem 1.10])**.** *The tensor product space $(\mathbb{C}^d)^{\otimes t}$ can be decomposed as*

$$(\mathbb{C}^d)^{\otimes t} \cong \bigoplus_{\lambda \vdash t} P_\lambda \quad \textit{with } P_\lambda = W_\lambda \otimes V_\lambda$$

where $\lambda \vdash t$ indexes partitions of $\{1,\ldots,t\}$, which are commonly represented by Young diagrams.[7]

The Weyl modules W_λ *are irreducible subspaces for the unitary group* $\mathrm{U}(d)$ *and the* Specht modules V_λ *are irreducible subspaces for the symmetric group S_t. Consequently, the action of the product group $S_t \times U_d$ on $(\mathbb{C}^d)^{\otimes t}$ decomposes as*

$$R_\pi = \sum_{\lambda \vdash t} \mathbb{1}_{W_\lambda} \otimes R_\pi^{(\lambda)} \quad \textit{and} \quad U^{\otimes t} = \sum_{\lambda \vdash t} U^{(\lambda)} \otimes \mathbb{1}_{V_\lambda}\ ,$$

where $(W_\lambda, U^{(\lambda)})$ and $(V_\lambda, R_\pi^{(\lambda)})$ are irreducible representations of the unitary group $\mathrm{U}(d)$ and the symmetric group S_t, respectively.

We denote the specific basis which block-diagonalizes all the above operations as the Schur-Weyl basis.

Definition 2.9 (Schur-Weyl basis)**.** Let $\{|w_{\lambda,i}\rangle\}_i$ and $\{|v_{\lambda,j}\rangle\}_j$ be orthonormal bases of W_λ and V_λ, respectively. Then we call

$$\{|w_{\lambda,i}\rangle \otimes |v_{\lambda,j}\rangle\}_{\lambda,i,j}$$

a Schur-Weyl basis of $(\mathbb{C}^d)^{\otimes t}$ (where we interpret each vector $|w_{\lambda,i}\rangle \otimes |v_{\lambda,j}\rangle \in P_\lambda$ as a $(\mathbb{C}^d)^{\otimes t}$-vector in the natural way).

The following decomposition of the distinct subspace projector Λ easily follows from Schur-Weyl duality: since Λ is invariant under permutation of the tensor factors (i.e. $R_\pi \Lambda = \Lambda R_\pi$ for all $\pi \in S_t$), it acts as an identity on the Specht modules V_λ by Schur's lemma. The same decomposition also holds for any permutation-invariant operator but we shall only need the following.

[7] We note that throughout this paper $d \gg t$, otherwise the Young diagrams need to be restricted to d rows.

Lemma 2.10 (Decomposition of the distinct subspace projector)**.** *Let $\Lambda \in \mathrm{L}((\mathbb{C}^d)^{\otimes t})$ be the projector on the tuples of distinct strings defined in Equation (2.2). Then,*

$$\Lambda = \sum_{\lambda \vdash t} \Lambda^{(\lambda)}_{W_\lambda} \otimes \mathbb{1}_{V_\lambda} \,,$$

with each $\Lambda^{(\lambda)}_{W_\lambda}$ a projector on a subspace of W_λ.

We will also need the following relation between the dimensions of the Weyl and Specht modules.

Lemma 2.11 ([Chr06], Theorem 1.16)**.** *The dimensions of the Weyl and Specht modules W_λ and V_λ satisfy*

$$\dim(W_\lambda) = \frac{\dim(V_\lambda)}{t!} \prod_{(i,j)\in\lambda} (d + j - i) \,,$$

where (i, j) denotes the row and column number of a box in the Young diagram corresponding to λ.

The following lemma follows from the standard formula for computing the projector on isotypical copies of an irreducible subspace (see [FH13, Section 2.4]). Here we apply it to the representation of the symmetric group S_t over $(\mathbb{C}^d)^{\otimes t}$, where Schur-Weyl duality implies that the subspace of all isotypic copies of the Specht modules V_λ is exactly $P_\lambda = W_\lambda \otimes V_\lambda$.

Lemma 2.12. *The projection onto the subspace P_λ is given by*

$$\mathbb{1}_{P_\lambda} = \frac{\dim(V_\lambda)}{t!} \sum_{\pi\in S_t} \chi_\lambda(\pi^{-1}) R_\pi \,, \tag{2.5}$$

where $\chi_\lambda(\cdot) = \mathrm{Tr}\left[R^\lambda_{(\cdot)}\right]$ is the character corresponding to the irrep $(V_\lambda, R^{(\lambda)}_{(\cdot)})$.

We will also need (a special case of) the standard Schur orthogonality relations for matrix coefficients (see [Bum13], Theorems 2.3 and 2.4). These relations say that if we express unitary irreducible representations of a group in any basis, then the different matrix entries are orthogonal under an inner product obtained by averaging over the group. Here we specialize the above to the Schur-Weyl basis and the unitary irreducible subrepresentations $R^\lambda_{(\cdot)}$ of $R_{(\cdot)}$ as given in Lemma 2.8.

Lemma 2.13 (Schur orthogonality relations)**.** *Let $(V_\lambda, R^\lambda_{(\cdot)}), (V_{\lambda'}, R^{\lambda'}_{(\cdot)})$ be two irreducible representations of the symmetric group S_t. Then,*

$$\mathop{\mathbb{E}}_{\pi\in S_t} \left[\langle v_{\lambda,i}| \, R^\lambda_\pi \, |v_{\lambda,j}\rangle \, \overline{\langle v_{\lambda',k}| \, R^{\lambda'}_\pi \, |v_{\lambda',\ell}\rangle} \right] = \frac{1}{\dim(V_\lambda)} \delta_{\lambda,\lambda'} \delta_{i,k} \delta_{j,\ell} \,.$$

3 Analysis of the PFC ensemble

Our main technical result is to introduce a new ensemble of random unitaries, the concatenation of a random Clifford, a random binary phase, and a random basis permutation, and show that this ensemble is a t-design even for superpolynomially large t. More formally, we prove the following.

Theorem 3.1. *Let $t \in \mathbb{N}$. For $d = 2^n$, the d-dimensional PFC random unitary ensemble (see Definition 1.1) is a diamond ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$.*

Proof. Let P denote a uniformly random permutation operator on n-qubits as defined in Equation (2.3), let F be a uniformly random binary phase operator as in Equation (2.4) and let C be a uniformly random n-qubit Clifford, all sampled independently. Let A be an n-qubit quantum register, E another quantum register isomorphic to A, and consider an arbitrary state $|\psi\rangle_{\mathsf{AE}}$. By Definition 2.5, it suffices to argue that the density matrices

$$\rho := \mathop{\mathbb{E}}_{P,F,C} (PFC)^{\otimes t}_{\mathsf{A}} \, |\psi\rangle\!\langle\psi|_{\mathsf{AE}} \, (PFC)^{\otimes t,\dagger}_{\mathsf{A}} \quad \text{and} \quad \rho^{\mathrm{hr}} := \mathop{\mathbb{E}}_{U\sim\mathrm{Haar}} U^{\otimes t}_{\mathsf{A}} \, |\psi\rangle\!\langle\psi|_{\mathsf{AE}} \, U^{\otimes t,\dagger}_{\mathsf{A}}$$

are $O(t/\sqrt{d})$-close in trace distance. It will be convenient to define

$$\xi_{\mathsf{AE}} := \mathop{\mathbb{E}}_{C} C^{\otimes t}_{\mathsf{A}} \, |\psi\rangle\!\langle\psi|_{\mathsf{AE}} \, C^{\otimes t,\dagger}_{\mathsf{A}} \, .$$

Recalling the definition of t-wise R-twirl operator (acting on the register A, which we shall omit from the notation henceforth), our goal is to bound the following trace distance:

$$\|\rho - \rho^{\mathrm{hr}}\|_1 = \left\| \mathcal{M}^{(t)}_{PF}\left(\xi_{\mathsf{AE}}\right) - \mathcal{M}^{(t)}_{\mathrm{Haar}}\left(\xi_{\mathsf{AE}}\right) \right\|_1 , \tag{3.1}$$

where we used the fact that the product unitary UC is also Haar distributed by the invariance of the Haar measure, so $\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(\xi_{\mathsf{AE}}\right) = \mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\psi\rangle\!\langle\psi|_{\mathsf{AE}}\right)$.

Step 1: Reduction to distinct subspace. To show the above, the calculations are easier if we restrict attention to the subspace of $(\mathbb{C}^d)^{\otimes t}$ that consists of distinct strings, i.e. basis states of the form $|\boldsymbol{x}\rangle$ where $\boldsymbol{x} = (x_1, \ldots, x_t) \in [d]^t$ is a tuple of distinct strings. We show in Section 3.1 that applying a t-wise Clifford twirl ensures that the input state has a large overlap with this subspace.

Lemma 3.2 (Clifford twirl and distinct subspace)**.** *Let Λ be the projector on the distinct subspace defined in Equation (2.2). Then, for any state ρ on the register A, we have*

$$\mathrm{Tr}\left[\Lambda \mathop{\mathbb{E}}_{C} C^{\otimes t} \rho \, C^{\otimes t,\dagger}\right] \geq 1 - O(t^2/d) \, .$$

Let ϕ_{AE} be the mixed state obtained by normalizing the (positive semi-definite) matrix $\Lambda_{\mathsf{A}}\xi_{\mathsf{AE}}\Lambda_{\mathsf{A}}$, where ξ_{AE} is defined in Equation (3.1). Then the above together with the Gentle Measurement lemma (see [Win99, Lemma 9]) implies that $\|\phi - \xi\|_1 \leq O(t/\sqrt{d})$. Consider a purification $|\phi\rangle_{\mathsf{A\tilde{E}}}$ of the state ϕ_{AE} which satisfies $\Lambda_{\mathsf{A}} |\phi\rangle_{\mathsf{A\tilde{E}}} = |\phi\rangle_{\mathsf{A\tilde{E}}}$; such a purification exists by construction of ϕ_{AE}. Then,

$$\begin{aligned} \left\|\rho - \rho^{\mathrm{hr}}\right\|_1 &\leq \left\| \mathcal{M}^{(t)}_{PF}\left(\phi_{\mathsf{AE}}\right) - \mathcal{M}^{(t)}_{\mathrm{Haar}}\left(\phi_{\mathsf{AE}}\right) \right\|_1 + O(t/\sqrt{d}) \\ &\leq \left\| \mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}}\right) - \mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}}\right) \right\|_1 + O(t/\sqrt{d}) \, , \end{aligned} \tag{3.2}$$

where we use that a t-wise twirl is a quantum channel and the 1-norm can only decrease under partial trace. Thus, we may assume for the rest of the proof that the input is supported over the distinct subspace of register A.

In order to bound Equation (3.2), we find explicit expressions for both states $\mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}}\right)$ and $\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}}\right)$ in the Schur-Weyl basis.

Step 2: Action of Haar twirl. Recall that we consider states $|\phi\rangle_{\mathsf{A\tilde{E}}}$ where the register A is supported over the distinct subspace, i.e. states satisfying $\Lambda_{\mathsf{A}} |\phi\rangle_{\mathsf{A\tilde{E}}} = |\phi\rangle_{\mathsf{A\tilde{E}}}$. To apply Schur-Weyl duality, we decompose $\mathsf{A} \cong (\mathbb{C}^d)^{\otimes t} \cong \bigoplus_{\lambda\vdash t} P_\lambda$ (where $P_\lambda = W_\lambda \otimes V_\lambda$) in terms of the Weyl and Specht modules W_λ and V_λ, respectively. We first analyse the output of the Haar twirl in the Schur-Weyl basis, proving the following decomposition.

Lemma 3.3 (Action of Haar twirl)**.** *Let $\rho_\lambda = \dfrac{\mathbb{1}_{W_\lambda}}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}$ be the maximally mixed state on W_λ. Then*

$$\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}}\right) = \sum_{\lambda\vdash t} \rho_\lambda \otimes \mathrm{Tr}_{W_\lambda}\left[\mathbb{1}_{P_\lambda} \, |\phi\rangle\!\langle\phi|_{\mathsf{A\tilde{E}}} \, \mathbb{1}_{P_\lambda}\right] . \tag{3.3}$$

We prove the above lemma in Section 3.2. In fact, the same proof also shows that Lemma 3.3 holds for arbitrary input states, not just input states supported only on the distinct string subspace. However, we will not need this since we are only interested in bounding the trace distance in Equation (3.2), which considers states on the distinct string subspace.

Step 3: Action of permutation-phase twirl on distinct subspace. Next, recalling the decomposition of the distinct subspace projector given by Lemma 2.10, we show that the action of the permutation-phase twirl results in the following; we defer the proof to Section 3.3.

Lemma 3.4 (Action of permutation-phase twirl)**.** *Denoting by* $\sigma_\lambda = \frac{\Lambda^{(\lambda)}_{W_\lambda}}{\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right]}$ *the maximally mixed state on the subspace* $\mathrm{supp}(\Lambda^{(\lambda)}_{W_\lambda}) \cap W_\lambda$*, we have*

$$\mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \sum_{\lambda \vdash t} \sigma_\lambda \otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}\, |\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\, \mathbb{1}_{P_\lambda}]\,.$$

Step 4: Haar twirl vs permutation-phase twirl. We now have explicit expressions for both of the twirling operations in Equation (3.2). We will use these expressions to show that

$$\left\|\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) - \mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right)\right\|_1 \leq O(t^2/d)\,. \tag{3.4}$$

Plugging this into Equation (3.2) completes the proof of the theorem.

To prove Equation (3.4), first note that all the subspaces P_λ are orthogonal. Therefore, inserting the expressions from Lemma 3.3 and Lemma 3.4 we have

$$\begin{aligned}\left\|\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) - \mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right)\right\|_1 &= \sum_\lambda \left\|(\rho_\lambda - \sigma_\lambda) \otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}\, |\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\, \mathbb{1}_{P_\lambda}]\right\|_1 \\ &= \sum_\lambda \|\rho_\lambda - \sigma_\lambda\|_1 \cdot \left\|\mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}\, |\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\, \mathbb{1}_{P_\lambda}]\right\|_1 \\ &\leq \max_\lambda \|\rho_\lambda - \sigma_\lambda\|_1,\end{aligned} \tag{3.5}$$

where the second equality used that the 1-norm (trace norm) is multiplicative under tensor products, while the inequality follows from the fact that $\sum_\lambda \|\,\mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda}\, |\phi\rangle\!\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\, \mathbb{1}_{P_\lambda}]\| = 1$, which can be seen by taking the trace on both sides of Equation (3.3).

Note that in general, given a vector space $A = B \oplus B^\perp$, we have that

$$\left\|\frac{\mathbb{1}_A}{\dim(A)} - \frac{\mathbb{1}_B}{\dim(B)}\right\|_1 = \left\|\frac{\mathbb{1}_B}{\dim(A)} - \frac{\mathbb{1}_B}{\dim(B)}\right\|_1 + \left\|\frac{\mathbb{1}_{B^\perp}}{\dim(A)}\right\|_1 = 2 - 2\frac{\dim(B)}{\dim(A)}\,.$$

Since $\Lambda^{(\lambda)}_{W_\lambda}$ is a projector on a subspace of W_λ, we obtain the following for any $\lambda \vdash t$:

$$\|\rho_\lambda - \sigma_\lambda\|_1 = \left\|\frac{\mathbb{1}_{W_\lambda}}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]} - \frac{\Lambda^{(\lambda)}_{W_\lambda}}{\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right]}\right\|_1 = 2 - 2\frac{\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right]}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}\,. \tag{3.6}$$

We first compute the trace in the numerator of the right hand side.

Claim 3.5. $\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right] = \frac{\dim(V_\lambda)}{t!} \cdot \mathrm{Tr}[\Lambda]$.

Proof. Lemma 2.10 implies that $\Lambda = \sum_{\lambda \vdash t} \Lambda^{(\lambda)}_{W_\lambda} \otimes \mathbb{1}_{V_\lambda}$. Since $P_\lambda = W_\lambda \otimes V_\lambda$, it follows that

$$\mathrm{Tr}\Big[\Lambda^{(\lambda)}_{W_\lambda}\Big] = \frac{1}{\dim V_\lambda} \mathrm{Tr}[\Lambda \mathbb{1}_{P_\lambda}]\,. \tag{3.7}$$

Plugging in the expression for the projector $\mathbb{1}_{P_\lambda}$ from Lemma 2.12, we get that

$$\mathrm{Tr}[\Lambda \mathbb{1}_{P_\lambda}] = \frac{\dim(V_\lambda)}{t!} \sum_\pi \chi_\lambda(\pi^{-1}) \,\mathrm{Tr}[\Lambda R_\pi] = \frac{\dim(V_\lambda)^2}{t!} \mathrm{Tr}[\Lambda]\,,$$

where we used the fact that $\mathrm{Tr}(R_\pi \Lambda) = 0$ unless $\pi = e$, and that $\chi_\lambda(e) = \mathrm{Tr}[\mathbb{1}_{V_\lambda}] = \dim(V_\lambda)$. This yields the desired result after insertion into Equation (3.7). □

Given the above, we can now compute the quantity in Equation (3.6) by using the dimension bounds for Weyl and Specht modules.

Claim 3.6. $1 - \dfrac{\mathrm{Tr}\Big[\Lambda^{(\lambda)}_{W_\lambda}\Big]}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]} \leq O(t^2/d)$.

Proof. Using Claim 3.5 and Lemma 2.11, we have

$$1 - \frac{\mathrm{Tr}\Big[\Lambda^{(\lambda)}_{W_\lambda}\Big]}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]} = 1 - \frac{\dim(V_\lambda)\,\mathrm{Tr}[\Lambda]}{t!\,\dim(W_\lambda)} = 1 - \frac{\mathrm{Tr}[\Lambda]}{\Pi_{(i,j)\in\lambda}(d+j-i)}\,.$$

Note that $\mathrm{Tr}[\Lambda] = \frac{d!}{(d-t)!} \geq (d-t)^t$ and $\Pi_{(i,j)\in\lambda}(d+j-i) \leq (d+t)^t$, since there are at most t boxes in the Young diagram corresponding to λ, and the coordinates i, j range from 1 to t. Thus (assuming $t^2 < d$, as otherwise the claim becomes trivial),

$$1 - \frac{\mathrm{Tr}\Big[\Lambda^{(\lambda)}_{W_\lambda}\Big]}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]} \leq 1 - \left(\frac{d-t}{d+t}\right)^t \leq 1 - \left(\frac{1-t/d}{1+t/d}\right)^t \leq O(t^2/d)\,.$$

□

Plugging the bound from Claim 3.6 into Equation (3.6) shows Equation (3.4). This completes the proof of Theorem 3.1. □

3.1 Clifford twirl and distinct subspace (proof of Lemma 3.2)

We will omit the registers $\mathsf{A} = \mathsf{A}_1 \ldots \mathsf{A}_t$ from the notation unless needed. Our goal is to show that applying a t-wise Clifford twirl on any input state ρ produces a state that has a large overlap with the distinct subspace. The proof only relies on the standard fact that the Clifford group forms a 2-design. We consider the projector on the orthogonal complement of the distinct subspace

$$\bar{\Lambda} = \mathbb{1} - \Lambda = \sum_{\boldsymbol{x} \in [d]^t \setminus \mathrm{distinct}(d,t)} |\boldsymbol{x}\rangle\langle\boldsymbol{x}|\,,$$

and decompose the projector into $O(t^2)$ sub-projectors according to which elements collide: since any tuple of non-distinct strings must have at least two equal entries, we have

$$\bar{\Lambda} \leq \sum_{1 \leq i < j \leq t} \Pi_{ij} \otimes \mathbb{1}_{[n]\setminus\{ij\}}\,, \qquad \text{where } \Pi_{ij} = \sum_{x \in [d]} |x\rangle\langle x|_i \otimes |x\rangle\langle x|_j\,,$$

where the subscript i denotes the register A_i. We shall omit the identity from the notation below.

Using cyclicity of trace, we have

$$\operatorname{Tr}\left[\bar{\Lambda}\mathop{\mathbb{E}}_{C} C^{\otimes t}\rho C^{\otimes t}\right] \leq \sum_{i<j}\operatorname{Tr}\left[\mathop{\mathbb{E}}_{C} C^{\otimes t,\dagger}\Pi_{ij}C^{\otimes t}\rho\right] = \sum_{i<j}\operatorname{Tr}\left[\mathop{\mathbb{E}}_{C}(C_i^\dagger \otimes C_j^\dagger)\Pi_{ij}(C_i\otimes C_j)\rho\right],$$

where for the second equality, we cancelled the C-unitaries on all systems except i and j, with C_i denoting application of C on system i.

Using the standard fact that the Clifford group forms a 2-design, we can replace the average over Clifford unitaries with an average over the Haar measure in the above expression since we only use the second moment. Thus,

$$\operatorname{Tr}\left[\bar{\Lambda}\mathop{\mathbb{E}}_{C} C^{\otimes t}\rho C^{\otimes t}\right] \leq \sum_{i<j}\operatorname{Tr}\left[\mathop{\mathbb{E}}_{U\sim\text{Haar}}(U_i^\dagger\otimes U_j^\dagger)\Pi_{ij}(U_i\otimes U_j)\rho\right] = \sum_{i<j}\operatorname{Tr}\left[\mathop{\mathbb{E}}_{U\sim\text{Haar}}(U^\dagger\otimes U^\dagger)\Pi_{ij}(U\otimes U)\rho_{ij}\right],$$

where for the last equality, we performed the partial trace over all systems except i and j, with ρ_{ij} denoting the reduced state on these systems. Since ρ_{ij} is a quantum state, we can bound each of the $O(t^2)$ trace terms by the corresponding operator norms to obtain

$$\operatorname{tr}\left[\bar{\Lambda}\mathop{\mathbb{E}}_{C} C^{\otimes t}\rho C^{\otimes t}\right] \leq O(t^2 d)\left\|\mathop{\mathbb{E}}_{U\sim\text{Haar}}(U\otimes U)^\dagger\left(\frac{\Pi}{\operatorname{Tr}[\Pi]}\right)(U\otimes U)\right\|_\infty, \tag{3.8}$$

where $\Pi = \sum_{x\in[d]}|x\rangle\langle x|\otimes|x\rangle\langle x|$ with $\operatorname{Tr}[\Pi] = d$.

Since applying a Haar random unitary on any state gives a Haar random state $|\psi\rangle \in \mathbb{C}^d$, by linearity of expectation the expression inside the operator norm is

$$\mathop{\mathbb{E}}_{|\psi\rangle\sim\text{Haar}}[|\psi\rangle\langle\psi|\otimes|\psi\rangle\langle\psi|].$$

It is well known that this is the maximally mixed state on the symmetric subspace $\operatorname{Sym}^2(d)$ of $\mathbb{C}^d\otimes\mathbb{C}^d$, which has dimension $\frac{d(d+1)}{2}$ (see [Har13], Proposition 6). Thus, the operator norm is $\frac{2}{d(d+1)}$, and inserting this into Equation (3.8) yields the claimed result.

3.2 Action of the t-wise Haar twirl (proof of Lemma 3.3)

In order to derive the expression given by Lemma 3.3, we first compute the result of applying the t-wise Haar twirl on Schur-Weyl basis states.

Lemma 3.7. *Let $\mathsf{A} \cong (\mathbb{C}^d)^{\otimes t}$ and let $|\alpha\rangle = |w_{\lambda,i}\rangle\otimes|v_{\lambda,j}\rangle$ and $|\beta\rangle = |w_{\lambda',i'}\rangle\otimes|v_{\lambda',j'}\rangle$ be Schur-Weyl basis states on A. Then*

$$\mathcal{M}^{(t)}_{\text{Haar}}(|\alpha\rangle\langle\beta|) = \begin{cases}\frac{\mathbb{1}_{W_\lambda}}{\dim W_\lambda}\otimes|v_{\lambda,j}\rangle\langle v_{\lambda,j'}|\,, & \textit{if } \lambda=\lambda' \textit{ and } i=i' \\ 0\,, & \textit{otherwise}\end{cases}.$$

Proof. By Schur-Weyl duality (Lemma 2.8), we have that $U^{\otimes t} = \sum_{\lambda_1} U^{(\lambda_1)}\otimes\mathbb{1}_{V_{\lambda_1}}$ where $U^{(\lambda_1)}$ only acts on W_{λ_1}. Thus,

$$\mathcal{M}^{(t)}_{\text{Haar}}(|\alpha\rangle\langle\beta|) = \sum_{\lambda_1,\lambda_2}\left(\mathop{\mathbb{E}}_{U\sim\text{Haar}}U^{(\lambda_1)}|w_{\lambda,i}\rangle\langle w_{\lambda',i'}|U^{(\lambda_2),\dagger}\right)\otimes\mathbb{1}_{V_{\lambda_1}}|v_{\lambda,j}\rangle\langle v_{\lambda',j'}|\mathbb{1}_{V_{\lambda_2}}.$$

We may assume that $\lambda=\lambda'$, since the above is zero otherwise. Then, we have that

$$\mathcal{M}^{(t)}_{\text{Haar}}(|\alpha\rangle\langle\beta|) = \left(\mathop{\mathbb{E}}_{U\sim\text{Haar}}U^{(\lambda)}|w_{\lambda,i}\rangle\langle w_{\lambda,i'}|U^{(\lambda),\dagger}\right)\otimes|v_{\lambda,j}\rangle\langle v_{\lambda,j'}|\,.$$

We claim that the expression inside the paranthesis above is zero unless $i = i'$, in which case it is the maximally mixed state on the subspace W_λ. This follows from a standard fact in representation theory called Schur's Lemma (see [FH13], Lemma 1.7), which says that if (μ, H) is an irreducible representation of a group G and $T : H \to H$ is a linear map such that $T \circ \mu = \mu \circ T$ (such a map is called an *intertwiner*), then $T = \gamma \cdot \mathbb{1}_H$ for some scalar $\gamma \in \mathbb{C}$. Applying it to our setting, we see that the operator

$$T = \mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} U^{(\lambda)} |w_{\lambda,i}\rangle\langle w_{\lambda,i'}| U^{(\lambda),\dagger}$$

is an intertwiner for the irrep $(U^{(\lambda)}, W_\lambda)$. This fact is a simple consequence of Haar invariance, since

$$\begin{aligned} T\tilde{U}^{(\lambda)} &= \left(\mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} U^{(\lambda)} |w_{\lambda,i}\rangle\langle w_{\lambda,i'}| U^{(\lambda),\dagger} \right) \tilde{U}^{(\lambda)} \\ &= \mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} U^{(\lambda)} |w_{\lambda,i}\rangle\langle w_{\lambda,i'}| \left(\tilde{U}^{(\lambda),\dagger} U^{(\lambda)} \right)^\dagger \\ &= \mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} \left(\tilde{U}^{(\lambda)} U^{(\lambda)} \right) |w_{\lambda,i}\rangle\langle w_{\lambda,i'}| U^{(\lambda),\dagger} = \tilde{U}^{(\lambda)} T. \end{aligned}$$

Therefore, by Schur's Lemma,

$$\mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} U^{(\lambda)} |w_{\lambda,i}\rangle\langle w_{\lambda,i'}| U^{(\lambda),\dagger} = \gamma_{\lambda,i,i'} \mathbb{1}_{W_\lambda}$$

for some scalars $\gamma_{\lambda,i,i'}$. Since the operator on the left is traceless if $i \neq i'$, we have that $\gamma_{\lambda,i,i'} = 0$ unless $i = i'$. For $i = i'$, the normalisation follows because the operator on the left is a mixed state, i.e. it has unit trace. This completes the proof. □

We can now compute the result of applying a t-wise Haar twirl to a general state.

Proof of Lemma 3.3. Expanding in the Schur-Weyl basis,

$$|\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}} = \sum_{\lambda,i,j} (|w_{\lambda,i}\rangle |v_{\lambda,j}\rangle)_{\mathsf{A}} \otimes |e_{\lambda,i,j}\rangle_{\tilde{\mathsf{E}}}$$

for not necessarily normalised vectors $|e_{\lambda,i,j}\rangle_{\tilde{\mathsf{E}}}$. It follows from Lemma 3.7 and linearity that

$$\begin{aligned} \mathcal{M}^{(t)}_{\mathrm{Haar}} \left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \right) &= \sum_{\lambda,i,j,j'} \left(\frac{\mathbb{1}_{W_\lambda}}{\dim W_\lambda} \otimes |v_{\lambda,j}\rangle\langle v_{\lambda,j'}| \right)_{\mathsf{A}} \otimes |e_{\lambda,i,j}\rangle\langle e_{\lambda,i,j'}|_{\tilde{\mathsf{E}}} \\ &= \sum_{\lambda \vdash t} \rho_\lambda \otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda} |\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \mathbb{1}_{P_\lambda}], \end{aligned}$$

where ρ_λ is the maximally mixed state on the subspace W_λ. □

3.3 Permutation-phase twirl on distinct subspace (proof of Lemma 3.4)

In order to derive an exact expression for the permutation-phase twirl on the distinct subspace with projector Λ, we start with the following lemma.

Lemma 3.8. *Let* $\boldsymbol{x}, \boldsymbol{y} \in \mathrm{distinct}(d, t)$. *Then*

$$\mathcal{M}^{(t)}_{PF} (|\boldsymbol{x}\rangle\langle\boldsymbol{y}|) = \begin{cases} \dfrac{\Lambda R_\sigma}{\mathrm{Tr}[\Lambda]} & \text{if } \boldsymbol{y} = \boldsymbol{x}_\sigma \text{ for } \sigma \in S_t, \\ 0 & \text{otherwise.} \end{cases}$$

Note that since $\boldsymbol{x}, \boldsymbol{y} \in \mathrm{distinct}(d, t)$, there exists at most one permutation for which $\boldsymbol{y} = \boldsymbol{x}_\sigma$.

Proof. Applying the t-wise F-twirl first, we get that

$$\mathop{\mathbb{E}}_{F} F^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{y}| F^{\otimes t,\dagger} = \left(\mathop{\mathbb{E}}_{f} (-1)^{\sum_i f(x_i)+f(y_i)}\right) |\boldsymbol{x}\rangle\langle\boldsymbol{y}| \ .$$

Here, the expectation $\mathbb{E}_f$ is over a uniformly random function $f : [d] \to \{0,1\}$. Because $\boldsymbol{x}$ and $\boldsymbol{y}$ are both tuples of distinct strings, it is easy to see that

$$\mathop{\mathbb{E}}_{f} \left[(-1)^{\sum_i f(x_i)+f(y_i)}\right] = \begin{cases} 1, & \text{if } \boldsymbol{y} = \boldsymbol{x}_\sigma, \text{ for some } \sigma \in S_t, \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$

Next applying the t-wise P-twirl,

$$\mathop{\mathbb{E}}_{P} P^{\otimes t} \left(\mathop{\mathbb{E}}_{F} F^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| F^{\otimes t,\dagger}\right) P^{\otimes t,\dagger} = \mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| R_\sigma P^{\otimes t,\dagger} = \mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| P^{\otimes t,\dagger} R_\sigma \ ,$$

where we used that $P^{\otimes t}$ commutes with R_σ. To conclude, we note that for any tuple of distinct strings $\boldsymbol{x}$,

$$\mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| P^{\otimes t,\dagger} = \frac{\Lambda}{\mathrm{Tr}[\Lambda]} .$$

□

We will also need the following technical result, which follows from the Schur orthogonality relations (Lemma 2.13).

Lemma 3.9. *Let $|\alpha\rangle = |w_{\lambda,i}\rangle \otimes |v_{\lambda,j}\rangle$ and $|\beta\rangle = |w_{\lambda',i'}\rangle \otimes |v_{\lambda',j'}\rangle$ be Schur-Weyl basis states. Then*

$$\sum_{\sigma \in S_t} \langle\beta| R_\sigma^\dagger |\alpha\rangle R_\sigma = \delta_{\lambda,\lambda'}\delta_{i,i'} \cdot \frac{t!}{\dim V_\lambda} \cdot \left(\mathbb{1}_{W_\lambda} \otimes |v_{\lambda,j}\rangle\langle v_{\lambda,j'}|\right) .$$

Proof. We compute the matrix elements of $R_\sigma^\dagger$ in the Schur-Weyl basis. Recall that Lemma 2.8 implies that $R_\sigma = \sum_\lambda \mathbb{1}_{W_\lambda} \otimes R_\sigma^{(\lambda)}$, where $R_\sigma^{(\lambda)}$ is the irreducible sub-representation of R_σ on the Specht module V_λ. Thus,

$$\begin{aligned}
(\langle w_{\lambda',i'}| \langle v_{\lambda',j'}|) R_\sigma^\dagger (|w_{\lambda,i}\rangle |v_{\lambda,j}\rangle) &= (\langle w_{\lambda',i'}| \langle v_{\lambda',j'}|) \left(\sum_\lambda \mathbb{1}_{W_\lambda} \otimes R_\sigma^{(\lambda),\dagger}\right) (|w_{\lambda,i}\rangle |v_{\lambda,j}\rangle) \\
&= \delta_{\lambda,\lambda'}\delta_{i,i'} \cdot \langle v_{\lambda,j'}| R_\sigma^{(\lambda),\dagger} |v_{\lambda,j}\rangle \\
&= \delta_{\lambda,\lambda'}\delta_{i,i'} \cdot \overline{\langle v_{\lambda,j}| R_\sigma^{(\lambda)} |v_{\lambda,j'}\rangle} \ .
\end{aligned} \tag{3.9}$$

Therefore, for the rest of the proof, we consider $|\alpha\rangle$ and $|\beta\rangle$ with $\lambda = \lambda'$ and $i = i'$.

Again using the decomposition of R_σ in terms of its irreducible sub-representations together with Equation (3.9), we can write

$$\sum_{\sigma \in S_t} ((\langle w_{\lambda,i}| \langle v_{\lambda,j'}|) R_\sigma^\dagger (|w_{\lambda,i}\rangle |v_{\lambda,j}\rangle)) R_\sigma = \sum_{\lambda_1 \vdash t} \mathbb{1}_{W_{\lambda_1}} \otimes \left(\sum_{\sigma \in S_t} \overline{\langle v_{\lambda,j}| R_\sigma^{(\lambda)} |v_{\lambda,j'}\rangle} R_\sigma^{(\lambda_1)}\right)$$

Schur's orthogonality relations (Lemma 2.13) now imply that the operator in the paranthesis is zero unless $\lambda_1 = \lambda$, in which case it equals

$$\sum_{\sigma \in S_t} \overline{\langle v_{\lambda,j}| R_\sigma^{(\lambda)} |v_{\lambda,j'}\rangle} R_\sigma^{(\lambda)} = \frac{t!}{\dim(V_\lambda)} |v_{\lambda,j}\rangle\langle v_{\lambda,j'}| \, .$$

Plugging this in gives the desired result. □

We are now in a position to prove Lemma 3.4 which expresses the result of applying the permutation-phase twirl to any state that is only supported on distinct strings on A in terms of the Schur-Weyl subspaces.

Proof of Lemma 3.4. We first expand $|\phi\rangle$ in the standard basis on A:

$$|\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}} = \sum_{\boldsymbol{x}\in\mathrm{distinct}(d,t)} |\boldsymbol{x}\rangle_{\mathsf{A}} |\tilde{e}_{\boldsymbol{x}}\rangle_{\tilde{\mathsf{E}}} \,, \tag{3.10}$$

where $|\tilde{e}_{\boldsymbol{x}}\rangle_{\mathsf{E}}$'s are unnormalized and not necessarily orthogonal vectors and we used that $|\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}}$ is supported over the distinct subspace in the register A.

Applying the permutation-phase twirl to the register A, we have by linearity,

$$\mathcal{M}_{PF}^{(t)}\left(|\phi\rangle\langle\phi|_{\mathsf{AE}}\right) = \sum_{\boldsymbol{x},\boldsymbol{y}\in\mathrm{distinct}(d,t)} \left(\mathcal{M}_{PF}^{(t)}\left(|\boldsymbol{x}\rangle\langle\boldsymbol{y}|_{\mathsf{A}}\right)\right) \otimes |\tilde{e}_{\boldsymbol{x}}\rangle\langle\tilde{e}_{\boldsymbol{y}}|_{\tilde{\mathsf{E}}} \,.$$

Lemma 3.8 implies that the term in the parantheses is non-zero only when $\boldsymbol{y} = \boldsymbol{x}_\sigma$ for some $\sigma \in S_t$. Since we sum over all possible $\boldsymbol{x}, \boldsymbol{y} \in \mathrm{distinct}(d,t)$ and there is at most one such σ for each pair of tuples $\boldsymbol{x}$ and $\boldsymbol{y}$, it follows that

$$\begin{aligned}
\mathcal{M}_{PF}^{(t)}\left(|\phi\rangle\langle\phi|_{\mathsf{AE}}\right) &= \sum_{\substack{\boldsymbol{x}\in\mathrm{distinct}(d,t)\\ \sigma\in S_t}} \frac{(\Lambda R_\sigma)_{\mathsf{A}}}{\mathrm{Tr}[\Lambda]} \otimes |\tilde{e}_{\boldsymbol{x}}\rangle\langle\tilde{e}_{\boldsymbol{x}_\sigma}|_{\tilde{\mathsf{E}}} \\
&= \frac{\Lambda_{\mathsf{A}}}{\mathrm{Tr}[\Lambda]} \sum_{\sigma\in S_t} (R_\sigma)_{\mathsf{A}} \otimes \left(\sum_{\boldsymbol{x}\in\mathrm{distinct}(d,t)} |\tilde{e}_{\boldsymbol{x}}\rangle\langle\tilde{e}_{\boldsymbol{x}_\sigma}|_{\tilde{\mathsf{E}}} \right) \\
&= \frac{\Lambda_{\mathsf{A}}}{\mathrm{Tr}[\Lambda]} \sum_{\sigma\in S_t} (R_\sigma)_{\mathsf{A}} \otimes \mathrm{Tr}_{\mathsf{A}}\left[\left(R_\sigma^\dagger \otimes \mathbb{1}_{\tilde{\mathsf{E}}}\right)|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right] \,.
\end{aligned} \tag{3.11}$$

For the second equality we simply rearranged sums, and for the last equality we wrote the expression in parentheses in the second line more compactly as a partial trace, which follows directly from the expansion in Equation (3.10).

We now rewrite the above partial trace in the Schur-Weyl basis on A. Let

$$|\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}} = \sum_{\lambda,i,j} (|w_{\lambda,i}\rangle\, |v_{\lambda,j}\rangle)_{\mathsf{A}}\, |e_{\lambda,i,j}\rangle_{\tilde{\mathsf{E}}} \tag{3.12}$$

be the state in the Schur-Weyl basis, where $|e_{\lambda,i,j}\rangle_{\tilde{\mathsf{E}}}$ are unnormalized and not necessarily orthogonal vectors. Then,

$$\mathrm{Tr}_{\mathsf{A}}\left[\left(R_\sigma^\dagger \otimes \mathbb{1}_{\tilde{\mathsf{E}}}\right)|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right] = \sum_{\substack{\lambda,i,j\\ \lambda',i',j'}} \left(\left(\langle w_{\lambda',i'}|\, \langle v_{\lambda',j'}|\right) R_\sigma^\dagger \left(|w_{\lambda,i}\rangle\, |v_{\lambda,j}\rangle\right)\right) |e_{\lambda,i,j}\rangle\langle e_{\lambda',i',j'}|_{\tilde{\mathsf{E}}} \,.$$

Plugging the above into Equation (3.11),

$$\mathcal{M}_{PF}^{(t)}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \frac{\Lambda_{\mathsf{A}}}{\mathrm{Tr}[\Lambda]} \sum_{\substack{\lambda,i,j\\ \lambda',i',j'}} \left(\sum_{\sigma} \left(\left(\langle w_{\lambda',i'}|\, \langle v_{\lambda',j'}|\right) R_\sigma^\dagger \left(|w_{\lambda,i}\rangle\, |v_{\lambda,j}\rangle\right)\right) R_\sigma \right)_{\mathsf{A}} \otimes |e_{\lambda,i,j}\rangle\langle e_{\lambda',i',j'}|_{\tilde{\mathsf{E}}} \,.$$

Applying Lemma 3.9 to the term in parentheses, we can simplify this to

$$\mathcal{M}_{PF}^{(t)}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \frac{\Lambda_{\mathsf{A}}}{\mathrm{Tr}[\Lambda]} \sum_{\lambda,i,j,j'} \frac{t!}{\dim(V_\lambda)} \left(\mathbb{1}_{W_\lambda} \otimes |v_{\lambda,j}\rangle\langle v_{\lambda,j'}|\right)_{\mathsf{A}} \otimes |e_{\lambda,i,j}\rangle\langle e_{\lambda,i,j'}|_{\tilde{\mathsf{E}}} \,. \tag{3.13}$$

By Lemma 2.10, we can write $\Lambda = \bigoplus_{\lambda \vdash t} \Lambda^{(\lambda)}_{W_\lambda} \otimes \mathbb{1}_{V_\lambda}$, where $\Lambda^{(\lambda)}_{W_\lambda}$ are projectors supported on W_λ. Since $\mathbb{1}_{W_\lambda}$ is simply the identity on subspace W_λ, this implies

$$\Lambda(\mathbb{1}_{W_\lambda} \otimes |v_{\lambda,j}\rangle\langle v_{\lambda,j'}|) = \Lambda^{(\lambda)}_{W_\lambda} \otimes |v_{\lambda,j}\rangle\langle v_{\lambda,j'}| \,.$$

Plugging this into Equation (3.13) and rewriting it as a partial trace,

$$\mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \sum_{\lambda \vdash t} \frac{t!}{\dim(V_\lambda)\operatorname{Tr}[\Lambda]} \Lambda^{(\lambda)}_{W_\lambda} \otimes \operatorname{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda} |\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \, \mathbb{1}_{P_\lambda}] \,.$$

Since $\operatorname{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right] = (\dim(V_\lambda)\operatorname{Tr}[\Lambda])/t!$ as we showed in Claim 3.5, the first tensor factor is indeed the maximally mixed state on the support of the projector $\Lambda^{(\lambda)}_{W_\lambda}$. This completes the proof. □

4 Efficient unitary t-designs

As explained in Section 1.1, to turn the PFC ensemble from Theorem 3.1 into an ensemble of efficiently implementable unitaries, we can replace the random functions and permutations in the definition of the F and P random unitaries by their $O(t)$-wise independent counterparts. This will yield a very efficient t-design construction. We first recall the definitions of t-wise independent functions and permutations.

Definition 4.1 (t-wise independent functions)**.** A distribution $\mathcal{D}$ over functions $\mathcal{F} = \{f : [N] \to [M]\}$ is called t-wise independent if, for all distinct $x_1, \ldots, x_t \in [N]$ and for all $y_1, \ldots, y_t \in [M]$ it holds that

$$\Pr_{f \sim \mathcal{D}}\left[f(x_1) = y_1 \wedge \ldots \wedge f(x_t) = y_t\right] = M^{-t}.$$

Moreover, we say that an ensemble of functions $\mathcal{F}$ is t-wise independent if the uniform distribution over the set $\mathcal{F}$ is t-wise independent.

Definition 4.2 (t-wise independent permutations)**.** A distribution $\mathcal{D}$ over permutations $\mathcal{P} = \{\pi : [N] \to [N]\}$ is called t-wise independent if, for all distinct $x_1, \ldots, x_t \in [N]$ and all distinct $y_1, \ldots, y_t \in [N]$, it holds that

$$\Pr_{\pi \sim \mathcal{D}}\left[\pi(x_1) = y_1 \wedge \ldots \wedge \pi(x_t) = y_t\right] = (N \cdot (N-1) \cdots (N-t+1))^{-1} \,.$$

Moreover, we say that $\mathcal{D}$ is δ-approximate t-wise independent[8] if instead

$$\left| \Pr_{\pi \sim \mathcal{D}}\left[\pi(x_1) = y_1 \wedge \ldots \wedge \pi(x_t) = y_t\right] - (N \cdot (N-1) \cdots (N-t+1))^{-1} \right| \leq \delta.$$

We will use the following theorem due to Alon and Lovett [AL13], which says that an *approximate* t-wise independent distribution over permutations is always statistically close to an exact t-wise independent distribution over permutations—provided that the error is sufficiently small.

Theorem 4.3 (Theorem 1.1, [AL13])**.** *Let $\mathcal{D}$ be a distribution over permutations $\mathcal{P} = \{\pi : [N] \to [N]\}$ which is δ-approximate t-wise independent. Then, there exists a distribution $\mathcal{D}'$ over $\mathcal{P}$ which is exactly t-wise independent such that $\|\mathcal{D} - \mathcal{D}'\|_1 \leq O\left(\delta \cdot N^{4t}\right)$.*

Our efficient unitary t-design ensemble is a product of a random Clifford unitary, a $2t$-wise independent binary phase operator, and an *approximate* t-wise independent permutation operator. To prove this, we make use of $O(t)$-wise independence in order to switch to the "fully random" PFC ensemble from Section 3. One way to argue that this switch is justified is to invoke a theorem due to Zhandry [Zha12, Theorem 3.1] which says that the behavior of any quantum algorithm making at most t queries to a $2t$-wise independent function is identical to its behavior when querying a uniformly random function. While Zhandry's result

[8] We remark that our definition of δ-approximate t-wise independence is the same as in [AL13].

allows us to immediately switch from a $2t$-wise independent binary-phase operator to a fully random binary-phase operator, the switch from *approximate* t-wise independent permutations is more subtle, especially so because the permutation operator is applied in-place rather than as a regular reversible oracle. It seems very likely that Zhandry's general result can be extended to the case of approximately independent permutations, but since we do not need the full power of Zhandry's result, we opt for a more direct and self-contained approach: we show that we can directly work with the $O(t)$-wise independence of the underlying function and permutation in order to carry out a similar analysis as in the "fully random" case in Section 3. Using this direct approach, we show the following theorem.

Theorem 4.4 (Efficient unitary t-design)**.** *Let $\mathcal{F} = \{f : [d] \to \{0,1\}\}$ be a $2t$-wise independent function family and $\mathcal{D}$ be a distribution over permutations $\mathcal{P} = \{\pi : [d] \to [d]\}$ which is δ-approximate t-wise independent for $\delta = O(t/d^{4t+\frac{1}{2}})$. Then, the unitary ensemble ν over $\mathrm{U}(d)$ which samples $U \sim \nu$ as a product*

$$U = P_\pi \, F_f \, C \,, \qquad \textit{where} \qquad \pi \sim \mathcal{D},\, f \sim \mathcal{F},\, C \sim \mathrm{C}(d)\,,$$

is a diamond ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$ (as per Definition 2.5).

Proof. First, we argue that we can replace the δ-approximate t-wise independent distribution $\mathcal{D}$ over $\mathcal{P}$ in our ensemble ν with a distribution which is exactly t-wise independent; specifically, we show that this incurs a small error in terms of diamond distance. To see this, we first use Theorem 4.3 to argue that there exists a distribution $\mathcal{D}'$ over $\mathcal{P}$ which is exactly t-wise independent such that $\|\mathcal{D} - \mathcal{D}'\|_1 \leq O\left(\delta \cdot d^{4t}\right)$. Let ν be our ensemble of unitaries and let ν' be the same ensemble, except that we replace $\mathcal{D}$ with $\mathcal{D}'$. Then,

$$\begin{aligned}\left\|\mathcal{M}_{\nu}^{(t)} - \mathcal{M}_{\nu'}^{(t)}\right\|_{\diamond} &= \max_{|\psi\rangle_{\mathsf{AE}}} \left\| \mathop{\mathbb{E}}_{\substack{\pi\sim\mathcal{D}\\ f\sim\mathcal{F}\\ C\sim\mathrm{C}(d)}} (P_\pi F_f C)_{\mathsf{A}}^{\otimes t} \, |\psi\rangle\langle\psi|_{\mathsf{AE}} \, (P_\pi F_f C)_{\mathsf{A}}^{\otimes t,\dagger} - \mathop{\mathbb{E}}_{\substack{\pi\sim\mathcal{D}\\ f\sim\mathcal{F}\\ C\sim\mathrm{C}(d)}} (P_\pi F_f C)_{\mathsf{A}}^{\otimes t} \, |\psi\rangle\langle\psi|_{\mathsf{AE}} \, (P_\pi F_f C)_{\mathsf{A}}^{\otimes t,\dagger} \right\|_1 \\ &= \max_{|\psi\rangle_{\mathsf{AE}}} \left\| \mathop{\mathbb{E}}_{\pi\sim\mathcal{D}} \xi_{\mathsf{AE}}^{\pi,\psi} - \mathop{\mathbb{E}}_{\pi\sim\mathcal{D}'} \xi_{\mathsf{AE}}^{\pi,\psi} \right\|_1 \,,\end{aligned}$$

where we define the ensemble of density matrices $\{\xi_{\mathsf{AE}}^{\pi,\psi}\}$ as

$$\xi_{\mathsf{AE}}^{\pi,\psi} := \mathop{\mathbb{E}}_{\substack{f\sim\mathcal{F}\\ C\sim\mathrm{C}(d)}} (P_\pi F_f C)_{\mathsf{A}}^{\otimes t} \, |\psi\rangle\langle\psi|_{\mathsf{AE}} \, (P_\pi F_f C)_{\mathsf{A}}^{\otimes t,\dagger}.$$

Then, by the strong convexity of the trace distance (see [NC10, Theorem 9.3]), we have that

$$\left\|\mathcal{M}_{\nu}^{(t)} - \mathcal{M}_{\nu'}^{(t)}\right\|_{\diamond} = \max_{|\psi\rangle_{\mathsf{AE}}} \left\| \mathop{\mathbb{E}}_{\pi\sim\mathcal{D}} \xi_{\mathsf{AE}}^{\pi,\psi} - \mathop{\mathbb{E}}_{\pi\sim\mathcal{D}'} \xi_{\mathsf{AE}}^{\pi,\psi} \right\|_1 \leq \left\|\mathcal{D} - \mathcal{D}'\right\|_1 \leq O\left(\delta \cdot d^{4t}\right) = O\left(\frac{t}{\sqrt{d}}\right) \,,$$

where we used Theorem 4.3 and that $\mathcal{D}$ is δ-approximate with parameter $\delta = O(t/d^{4t+\frac{1}{2}})$. Because of the triangle inequality, it suffices to show that the exact ensemble ν' satisfies

$$\left\|\mathcal{M}_{\nu'}^{(t)} - \mathcal{M}_{\mathrm{Haar}}^{(t)}\right\|_{\diamond} \leq O\left(\frac{t}{\sqrt{d}}\right).$$

Next, we argue that we can carry out essentially the same proof as in the "fully random" ensemble. For this, notice that the only place in the proof of Theorem 3.1 where we use properties of the F and P unitaries is Lemma 3.8. Therefore, it suffices to show that Lemma 3.8 still holds for the exact ensemble ν' which uses $O(t)$-wise independent functions and permutations rather than their "fully random" counterparts. In other words, we can complete the proof of Theorem 4.4 by showing Lemma 4.5, which we prove below. □

Lemma 4.5 (Restatement of Lemma 3.8 for t-wise independence.)**.** *Consider the PF ensemble, where F is chosen from an ensemble $\mathcal{F} = \{f : [d] \to \{0,1\}\}$ of $2t$-wise independent functions and where P is* *sampled according to an exact t-wise independent distribution $\mathcal{D}$ over permutations $\mathcal{P} = \{\pi : [d] \to [d]\}$. Let $\boldsymbol{x}, \boldsymbol{y} \in \mathrm{distinct}(d,t)$ be arbitrary. Then,*

$$\mathcal{M}^{(t)}_{PF}(|\boldsymbol{x}\rangle\langle\boldsymbol{y}|) = \begin{cases} \dfrac{\Lambda R_\sigma}{\mathrm{Tr}[\Lambda]} & \text{if } \boldsymbol{y} = \boldsymbol{x}_\sigma \text{ for } \sigma \in S_t, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Applying the $2t$-wise F-twirl first, we get that

$$\mathop{\mathbb{E}}_{F} F^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{y}| F^{\otimes t,\dagger} = \left(\mathop{\mathbb{E}}_{f \sim \mathcal{F}} (-1)^{\sum_i f(x_i) + f(y_i)} \right) |\boldsymbol{x}\rangle\langle\boldsymbol{y}| \, .$$

Here, the expectation is over the $2t$-wise independent function family $\mathcal{F}$. Because $\boldsymbol{x}, \boldsymbol{y} \in \mathrm{distinct}(d,t)$, it immediately follows from the $2t$-wise independence of $\mathcal{F}$ that

$$\mathop{\mathbb{E}}_{f \sim \mathcal{F}} \left[(-1)^{\sum_i f(x_i) + f(y_i)} \right] = \mathop{\mathbb{E}}_{f} \left[(-1)^{\sum_i f(x_i) + f(y_i)} \right] = \begin{cases} 1, & \text{if } \boldsymbol{y} = \boldsymbol{x}_\sigma, \text{ for some } \sigma \in S_t, \text{ and} \\ 0, & \text{otherwise} \, , \end{cases}$$

where the second expectation is over "fully random" functions f. Next applying the t-wise P-twirl,

$$\mathop{\mathbb{E}}_{P} P^{\otimes t} \left(\mathop{\mathbb{E}}_{F} F^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| F^{\otimes t,\dagger} \right) P^{\otimes t,\dagger} = \mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| R_\sigma P^{\otimes t,\dagger} = \mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| P^{\otimes t,\dagger} R_\sigma \, ,$$

where we used that $P^{\otimes t}$ commutes with R_σ. Finally, using that $\mathcal{D}$ is a distribution over t-wise independent permutations over $\mathcal{P}$, we get that for any tuple of distinct strings $\boldsymbol{x}$,

$$\begin{aligned} \mathop{\mathbb{E}}_{P} P^{\otimes t} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| P^{\otimes t,\dagger} &= \mathop{\mathbb{E}}_{\pi \sim \mathcal{D}} |\pi(x_1), \ldots, \pi(x_t)\rangle\langle\pi(x_1), \ldots, \pi(x_t)| \\ &= \mathop{\mathbb{E}}_{\pi} |\pi(x_1), \ldots, \pi(x_t)\rangle\langle\pi(x_1), \ldots, \pi(x_t)| = \frac{\Lambda}{\mathrm{Tr}[\Lambda]} \, , \end{aligned}$$

where the last expectation is over "fully random" permutations over $[d]$. This proves the claim. □

Linear-depth t-design construction. We can instantiate the *PFC* ensemble in Theorem 4.4 with a concrete ensemble of $O(t)$-wise independent functions and permutations. Very efficient $O(t)$-wise independent functions and permutations have been constructed, allowing us to get very efficient t-designs.

Concretely, let us say that an ensemble $\mathcal{F} = \{f : [N] \to \{0,1\}\}$ of functions is (s,r)-explicit if there is a circuit of size at most s and depth at most r that evaluates any function $f \in \mathcal{F}$ on a given input. Analogously, we say that a distribution $\mathcal{D}$ over permutations $\mathcal{P} = \{\pi : [N] \to [N]\}$ is (s,r)-explicit if there is a circuit of size at most s and depth at most r that evaluates $\pi \sim \mathcal{D}$ on a given input. Then, we show the following.

Lemma 4.6 (Explicit t-design)**.** *Let $d = 2^n$. Let $\mathcal{F} = \{f : [d] \to \{0,1\}\}$ be an (s,r)-explicit family of $2t$-wise independent functions and let $\mathcal{D}$ be an (s,r)-explicit, δ-approximate t-wise independent distribution over permutations $\mathcal{P} = \{\pi : [d] \to [d]\}$ for $\delta = O(t/d^{4t+\frac{1}{2}})$. Then, the resulting n-qubit ensemble ν from Theorem 4.4 is a diamond ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$, and each $U \sim \nu$ can be implemented in size $O(n^2 + s)$ and depth $O(n + r)$.*

Proof. Note that the depth which is required to implement a unitary U from the *PFC* ensemble ν is captured by the depth of implementing a random Clifford, a binary-phase operator (with respect to $\mathcal{F}$), and a permutation operator (with respect to $\mathcal{P}$). Implementing any n-qubit Clifford unitary requires size $O(n^2)$ and depth of $O(n)$ [BM21]. Because Toffoli gates are universal for classical computation [NC10], the unitaries $U_f : |x\rangle |y\rangle \to |x\rangle |y \oplus f(x)\rangle$ and $R_\pi : |x\rangle \to |\pi(x)\rangle$ can also be implemented in size $O(s)$ and depth $O(r)$ —potentially using using ancilla qubits initialized to $|0\rangle$. Therefore, both the binary phase operator F (consisting of a layer of Hadamard gates followed by Toffoli gates) as well as the permutation operator P (consisting solely of Toffoli gates) can be implemented in size $O(s)$ and depth $O(r)$. □

We remark that since one can use any 2-design instead of a Clifford in our construction, the above parameters can be optimized even further. For instance, it is known that one can sample a 2-design with circuits of quasilinear size and logarithmic depth (see [CLLW16]). We do not try to optimize the parameters here as the dominant parameters come from the construction of t-wise independent functions and permutations.

For $O(t)$-wise independent functions, we can use the following result.

Fact 4.7 ([Jof74, ABI86])**.** *For any $1 \leq t \leq N$, there exists an ensemble $\mathcal{F} = \{f : [N] \to \{0, 1\}\}$ of t-wise independent functions which can be evaluated in size and depth $O(t \log N)$.*

Note that if the domain is the space of n-qubit computational basis states, then $N = 2^n$ and the above gives a (s, r)-explicit family of $2t$-wise independent functions where $s = O(tn)$ and $r = O(tn)$.

For approximate $O(t)$-wise independent permutations, existing constructions are more complicated and less explicit than for $O(t)$-wise independent functions. A very elegant construction of $O(t)$-wise independent elements of the alternating group (i.e. the even permutations) on $p^3 - 1$ elements (for a sufficiently large prime p) can be found in [CK23, Theorem 1.5], which can be turned into $O(t)$-independent permutations on any (sufficiently large) number of elements N using techniques from [Kas07]. This construction only involves basic finite-field arithmetic, yielding $O(t)$-wise independent permutations with circuit size $O(t \operatorname{polylog}(N))$. Alternatively, one can use a matrix-based construction from [Kas07]. For this, [AL13] claimed (informally without proof) that the permutations can be implemented in time $O(t \log N)$, but upon closer inspection it is not obvious how to achieve this without additional $\operatorname{polylog}(N)$-factors.

Fact 4.8 ([Kas07, CK23, CHH$^+$24])**.** *For any sufficiently large $N \in \mathbb{N}$ and for any constant $C > 0$, there exists a δ-approximate t-wise independent distribution $\mathcal{D}$ over permutations $\mathcal{P} = \{\pi : [N] \to [N]\}$ with $\delta = O(N^{-Ct})$ such that each $\pi \sim \mathcal{D}$ can be evaluated in size $O(t \operatorname{polylog}(N))$.*

An upcoming work [CHH$^+$24] analyses the construction from [Kas07] in detail and presents explicit circuits, showing that the circuits can be made to have size $O(t \log^2(N))$ and depth $O(t \log N \log \log N)$ (with ancillas). In particular, for the domain of n-qubit computational basis states, using the result from [CHH$^+$24] gives an (s, r)-explicit family of approximate $O(t)$-wise independent permutations with $s = O(t \operatorname{poly}(n))$ and depth $r = \tilde{O}(tn)$, where $\tilde{O}$ hides logarithmic factors in n.

Using these aforementioned constructions of $O(t)$-wise independent functions and permutations as a black box, we obtain a t-design on n-qubits with the same size and depth as well.

Corollary 4.9 (Linear-depth t-design)**.** *Let $d = 2^n$. Let $\mathcal{F} = \{f : [d] \to \{0, 1\}\}$ be the $2t$-wise independent function family from Fact 4.7 and let $\mathcal{D}$ be the δ-approximate t-wise independent distribution over permutations $\mathcal{P} = \{\pi : [d] \to [d]\}$ in Fact 4.8 for $\delta = O(t/d^{4t+\frac{1}{2}})$. Then, the resulting n-qubit ensemble ν from Theorem 4.4 is a diamond ϵ-approximate t-design for $\epsilon = O(t/\sqrt{d})$, and each $U \sim \nu$ can be implemented in depth $O(t \operatorname{poly}(n))$.*

As mentioned before, the n-dependence of the depth can be improved to quasilinear using [CHH$^+$24], yielding a depth of $\tilde{O}(tn)$.

4.1 Amplification of approximation error

So far, we have constructed n-qubit t-designs with diamond error $\epsilon = O(t/2^{n/2})$. By repeating our construction m times independently in sequence, we can make the error decay like ϵ^m. Choosing $m = \Theta(t)$, this pushes down the error far enough that we can apply Lemma 2.7 to convert our diamond-error t-designs into relative-error ones. Of course these repetitions come at the cost of increasing the size and the depth of the circuits: choosing $m = \Theta(t)$ introduces an additional factor t in size and depth, which is why we obtain relative-error t-designs with quadratic scaling in t. These amplification techniques are standard in the t-design literature (see e.g. [Mel23]), but we spell out some of the details for completeness.

We begin with a simple auxiliary lemma about the concatenation of Haar moment operators with other moment operators, which follows immediately from the invariance of the Haar measure.

Lemma 4.10. *Let $X_1, \ldots, X_m$ be a collection of random matrices. Suppose that at least one of the X_i is independent and Haar random. Then*

$$\mathcal{M}^{(t)}_{X_m} \circ \mathcal{M}^{(t)}_{X_m} \circ \cdots \circ \mathcal{M}^{(t)}_{X_2} \circ \mathcal{M}^{(t)}_{X_1}(\cdot) = \mathcal{M}^{(t)}_{\mathrm{Haar}}(\cdot)\ .$$

With this, we can show that concatenating independent samples from a t-design results in exponential decay of the error.

Lemma 4.11. *Suppose that $X \sim \mathcal{X}$ is a diamond ϵ-approximate t-design. Let $X_1, \ldots, X_m$ be independent random unitaries sampled from $\mathcal{X}$. Then $X_1 \cdots X_m$ is a diamond ϵ^m-approximate t-design.*

Proof. The moment operator for the random unitary $X_1 \cdots X_m$ can be written as $(\mathcal{M}^{(t)}_X)^{\circ m} := \underbrace{\mathcal{M}^{(t)}_X \circ \cdots \circ \mathcal{M}^{(t)}_X}_{m \text{ times}}$. Observe that

$$(\mathcal{M}^{(t)}_X - \mathcal{M}^{(t)}_{\mathrm{Haar}})^{\circ m} = (\mathcal{M}^{(t)}_X)^{\circ m} + \sum_{i=1}^{m} \binom{m}{i} (-1)^i \mathcal{M}^{(t)}_{\mathrm{Haar}} = (\mathcal{M}^{(t)}_X)^{\circ m} - \mathcal{M}^{(t)}_{\mathrm{Haar}}\,.$$

The first equality uses the fact that when expanding out the product, all terms except $(\mathcal{M}^{(t)}_X)^{\circ m}$ have at least one copy of $\mathcal{M}^{(t)}_{\mathrm{Haar}}$ in them, which allows us to apply Lemma 4.10. The lemma now follows directly from the submultiplicativity of the diamond norm (Equation (2.1)). □

Applying Lemma 4.11 and Lemma 2.7 to the t-designs from Lemma 4.6, we get the following.

Corollary 4.12. *Let ν be the ensemble of n-qubit unitaries from Lemma 4.6 (instantiated with (s, r)-explicit families of functions and permutations). Let $U \sim \nu^{\circ m}$ be the ensemble of m-fold products of unitaries from ν, i.e. $U = U_1 \cdots U_m$ for $U_i \sim \nu$. Then, each $U \sim \nu^{\circ m}$ can be implemented in size $O(m \cdot (n^2 + s))$ and depth $O(m \cdot (n + r))$ and the ensemble $\nu^{\circ m}$ is*

1. *a diamond ϵ-approximate t-design with $\epsilon = O(t^m \cdot 2^{-nm/2})$, and*
2. *a relative-error ϵ-approximate t-design with $\epsilon = O(t^m \cdot 2^{2nt - nm/2})$.*

In particular, this means that for all $\epsilon > 0$ and $t \leq 2^{-n/4}$, we have

1. *diamond ϵ-approximate t-designs in size $O(t(n^2 + s) + t \log 1/\epsilon)$ and depth $O(t(n + r) + t \log 1/\epsilon)$,*
2. *relative-error ϵ-approximate t-designs in size $O(t^2(n^2 + s) + t \log 1/\epsilon)$ and depth $O(t^2(n + r) + t^2 \log 1/\epsilon)$.*

Proof. The first two statements follow immediately from Corollary 4.9 by means of Lemma 4.11 and Lemma 2.7. The second two statements follow from the first two simply by choosing $m = \max\{1, O(\frac{\log 1/\epsilon}{n})\}$ large enough that the desired ϵ is achieved. □

Again, using the upcoming work [CHH+24] the depth can be made quasilinear in n.

5 Pseudorandom unitaries with non-adaptive security

We first give a formal definition of PRUs, as proposed by Ji, Liu, and Song [JLS18]. Then, we prove that if we simply replace the random function in the F-operator and the random permutation in the P-operator by their pseudorandom counterparts, the resulting ensemble (described in Equation (1.1)) is a non-adaptive pseudorandom unitary.

Definition 5.1 (Pseudorandom unitary)**.** Let $n \in \mathbb{N}$ be the security parameter. An infinite sequence $\mathcal{U} = \{\mathcal{U}_n\}_{n \in \mathbb{N}}$ of n-qubit unitary ensembles $\mathcal{U}_n = \{U_k\}_{k \in \mathcal{K}}$ is a pseudorandom unitary if it satisfies the following conditions.

- (Efficient computation) There exists a polynomial-time quantum algorithm $\mathcal{Q}$ such that for all keys $k \in \mathcal{K}$, where $\mathcal{K}$ denotes the key space, and any $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n}$, it holds that

$$\mathcal{Q}(k, |\psi\rangle) = U_k |\psi\rangle \ .$$

- (Pseudorandomness) The unitary U_k, for a random key $k \sim \mathcal{K}$, is computationally indistinguishable from a Haar random unitary $U \sim \mathrm{Haar}(2^n)$. In other words, for any QPT algorithm $\mathcal{A}$, it holds that

$$\left| \Pr_{k \sim \mathcal{K}} [\mathcal{A}^{U_k}(1^\lambda) = 1] - \Pr_{U \sim \mathrm{Haar}} [\mathcal{A}^{U}(1^\lambda) = 1] \right| \leq \mathrm{negl}(n)\,.$$

  We call $\mathcal{U} = \{\mathcal{U}_n\}_{n \in \mathbb{N}}$ a *non-adaptive* pseudorandom unitary if $\mathcal{A}$ is only allowed to make parallel queries to the unitary U_k (or U in the Haar random case).

Note that, whenever we write $\mathcal{U}_n = \{U_k\}_{k \in \mathcal{K}}$, it is implicit that the key space $\mathcal{K}$ depends on the security parameter $n \in \mathbb{N}$, and that the length of each key $k \in \mathcal{K}$ is polynomial in n.

The main result of this section is that the construction in Equation (1.1) is indeed a non-adaptive PRU.

Theorem 5.2. *Let $n \in \mathbb{N}$ be the security parameter. Then, the ensemble $\mathcal{U}_n = \{U_k\}_{k \in \mathcal{K}}$ of n-qubit unitary operators defined in Equation (1.1) is a non-adaptive pseudorandom unitary when instantiated with ensembles of n-bit (quantum-secure) PRFs and PRPs.*

Proof. From the construction, it is clear that a random unitary U_k from the above family can be sampled efficiently (see e.g. [vdB21] for simple way to sample a uniform Clifford unitary). To argue security against any non-adaptive algorithm making $t = \mathrm{poly}(n)$ queries, it suffices to show that for any initial state $|\psi\rangle_{\mathsf{AE}}$, where register $\mathsf{A} \cong ((\mathbb{C}^2)^{\otimes n})^{\otimes t}$ is on nt qubits, and E is an arbitrary workspace register, the density matrices

$$\rho := \mathop{\mathbb{E}}_{k \in \mathcal{K}} (U_k)_{\mathsf{A}}^{\otimes t} |\psi\rangle\langle\psi|_{\mathsf{AE}} (U_k)_{\mathsf{A}}^{\otimes t, \dagger} \quad \text{and} \quad \rho^{\mathrm{hr}} := \mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} U_{\mathsf{A}}^{\otimes t} |\psi\rangle\langle\psi|_{\mathsf{AE}} U_{\mathsf{A}}^{\otimes t, \dagger},$$

are computationally indistinguishable with at most negligible advantage, since that is the general form of a non-adaptive distinguisher. From the post-quantum security of the PRF and PRP families assumed in Equation (1.1), it follows immediately that if we replace the pseudorandom permutation and function with their fully random counterparts to obtain the "fully random" state ρ^{fr}, then ρ^{fr} is computationally indistinguishable from ρ up to negligible advantage in n. Furthermore, since the PFC ensemble is a diamond-distance t-design, using the error bounds from Theorem 3.1, it follows that $\|\rho^{\mathrm{fr}} - \rho^{\mathrm{hr}}\|_1 \leq O(t/\sqrt{2^n})$. This is also negligible in n since $t = \mathrm{poly}(n)$ and the computational indistinguishability of ρ and ρ^{hr} follows. □

6 Pseudorandom isometries with adaptive security

One drawback of our PRU construction in Section 5 is that we are only able to prove non-adaptive security. This is a consequence of the fact that the analysis of the PFC ensemble achieves diamond-error, not relative error. In contrast to the t-design amplification in Section 4.1, we cannot simply amplify our PRU construction to achieve relative error. This is because the number of iterations in the amplification in Section 4.1 depends on the number of queries t, but for our PRUs we do not have an a priori bound on t.

However, it turns that if we relax the notion of PRUs to PRIs, a simple modification of our construction is able to achieve adaptive security. We have already given a high-level of this construction in Section 1.2.2, so we proceed with the formal statements here.

The formal definition of PRIs is entirely analogous to Definition 5.1, but we spell it out again for the sake of completeness.

Definition 6.1 (Pseudorandom isometry)**.** Let $n \in \mathbb{N}$ be the security parameter and choose an integer function $s(n) \in [0, n)$. An infinite sequence $\mathcal{V} = \{\mathcal{V}_n\}_{n \in \mathbb{N}}$ of ensembles $\mathcal{V}_n = \{V_k : \mathbb{C}^{2^{n-s(n)}} \to \mathbb{C}^{2^n}\}_{k \in \mathcal{K}}$ of isometries from $n - s(n)$ qubits to n qubits is a pseudorandom isometry if it satisfies the following conditions.

- (Efficient computation) There exists a polynomial-time quantum algorithm $\mathcal{Q}$ such that for all keys $k \in \mathcal{K}$, where $\mathcal{K}$ denotes the key space, and any $|\psi\rangle \in (\mathbb{C}^2)^{\otimes n-s(n)}$, it holds that

$$\mathcal{Q}(k, |\psi\rangle) = V_k |\psi\rangle \, .$$

- (Pseudorandomness) The isometry V_k, for a random key $k \sim \mathcal{K}$, is computationally indistinguishable from a Haar random isometry $V \sim \mathrm{Haar}(2^{n-s(n)}, 2^n)$.[9] In other words, for any QPT algorithm $\mathcal{A}$, it holds that

$$\left| \Pr_{k\sim\mathcal{K}} [\mathcal{A}^{V_k}(1^\lambda) = 1] - \Pr_{V\sim\mathrm{Haar}(2^{n-s(n)},2^n)} [\mathcal{A}^{V}(1^\lambda) = 1] \right| \leq \mathrm{negl}(n) \, .$$

Note that, whenever we write $\mathcal{V}_n = \{V_k\}_{k\in\mathcal{K}}$, it is implicit that the key space $\mathcal{K}$ depends on the security parameter $n \in \mathbb{N}$, and that the length of each key $k \in \mathcal{K}$ is polynomial in n. As in Definition 5.1, we can also define a notion of non-adaptive security, but we do not do so here as we will show full adaptive security.

Before stating the main result, we recall our PRI construction from Definition 1.2, which applied the PF operator to input state after appending it with ancillas in the $|+\rangle^{s(n)}$ state where $s(n) = \omega(\log n)$. As mentioned in the proof overview earlier, we do not require the random Cliffords here (because the fixed $|+\rangle$-input already ensures that a suitable distinct string condition holds).

Theorem 6.2. *The isometries defined in Definition 1.2 are pseudorandom isometries with adaptive security.*

The proof of Theorem 6.2 proceeds by performing a reduction to the distinct subspace, then performing gate teleportation to reduce to the non-adaptive case with post-selection, which can then be analyzed with relative error t-designs. Before we execute this strategy to prove Theorem 6.2, we need two auxiliary statements. The first one of these is a simple modification of Theorem 3.1 and shows that on the distinct subspace, the PF ensemble is a one-sided relative error t-design (Definition 2.6) for superpolynomial t.

Lemma 6.3. *Let Λ be the projector onto the distinct string subspace of $\mathsf{A} \cong (\mathbb{C}^d)^{\otimes t}$. Let $\tilde{\mathsf{E}}$ be an arbitrary register and $|\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}}$ a state such that $(\Lambda_{\mathsf{A}} \otimes \mathbb{1}_{\tilde{\mathsf{E}}}) |\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}} = |\phi\rangle_{\mathsf{A}\tilde{\mathsf{E}}}$. Then*

$$\mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) \leq (1 + O(t^2/d))\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right),$$

where as usual the twirling channels only act on register A.

Proof. From Lemma 3.3 and Lemma 3.4, we have the following explicit expressions:

$$\mathcal{M}^{(t)}_{PF}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \sum_{\lambda\vdash t} \sigma_\lambda \otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda} |\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \mathbb{1}_{P_\lambda}] \, ,$$

$$\mathcal{M}^{(t)}_{\mathrm{Haar}}\left(|\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}}\right) = \sum_{\lambda\vdash t} \rho_\lambda \otimes \mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda} |\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \mathbb{1}_{P_\lambda}] \, .$$

Here, $\sigma_\lambda = \dfrac{\Lambda^{(\lambda)}_{W_\lambda}}{\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right]}$ is the maximally mixed state on the subspace $\mathrm{supp}(\Lambda^{(\lambda)}_{W_\lambda}) \cap W_\lambda$ and $\rho_\lambda = \dfrac{\mathbb{1}_{W_\lambda}}{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}$ is the maximally mixed state on W_λ.

Since $\mathrm{supp}(\Lambda^{(\lambda)}_{W_\lambda}) \cap W_\lambda$ is a subspace of W_λ, we have the operator inequality $\Lambda^{(\lambda)}_{W_\lambda} \leq \mathbb{1}_{W_\lambda}$. This implies that for every λ,

$$\sigma_\lambda \leq \frac{\mathrm{Tr}[\mathbb{1}_{W_\lambda}]}{\mathrm{Tr}\left[\Lambda^{(\lambda)}_{W_\lambda}\right]}\rho_\lambda \leq \left(1 + O(t^2/d)\right)\rho_\lambda \, .$$

[9] By a Haar random isometry $V \sim \mathrm{Haar}(2^{n-s(n)}, 2^n)$ we mean an isometry sampled as follows. First sample a Haar random unitary $U \sim \mathrm{Haar}(2^n)$ on n qubits, then set $V|\psi\rangle := U(|\psi\rangle|0\rangle^{\otimes s(n)})$. In other words, to apply a Haar random isometry to an $(n - s(n))$-qubit input state $|\psi\rangle$, first pad $|\psi\rangle$ with $s(n)$ 0-kets and then apply a Haar random n-qubit unitary. We also note that the choice of padding with $|0\rangle^{\otimes s(n)}$ is arbitrary – due to the invariance of the Haar measure we could use any fixed $s(n)$-qubit state for the padding.

The second inequality uses Claim 3.6 and the fact that $\rho_\lambda \geq 0$.

Since this holds for every λ and taking the taking the tensor product with $\mathrm{Tr}_{W_\lambda}[\mathbb{1}_{P_\lambda} |\phi\rangle\langle\phi|_{\mathsf{A}\tilde{\mathsf{E}}} \mathbb{1}_{P_\lambda}]$ is an operator-monotone operation, the lemma follows. □

The second auxiliary lemma for the proof of Theorem 6.2 shows what happens in quantum gate teleportation when we do not use the "correct" resource state. This is a straightforward calculation; the main difficulty is the notation required to distinguish the different registers from one another. Both the statement and the proof of the lemma make heavy use of the unnormalised maximally entangled state between two registers $\mathsf{A} \cong \mathsf{A}'$:

$$|\Omega\rangle_{\mathsf{AA}'} = \sum_{i=1}^{|A|} |i\rangle_{\mathsf{A}} |i\rangle_{\mathsf{A}'} \ .$$

Lemma 6.4. *For $i = 1, \dots, t$, consider quantum registers $\mathsf{C_i}$ and $\mathsf{C_i'}$, all of size d. Fix a collection of d-dimensional unitaries A_i and a d-dimensional quantum state $|\psi\rangle_{\mathsf{C_0}}$. Define the vectors*

$$|\Omega_{A_i}\rangle_{\mathsf{C_i'C_{i+1}}} = ((A_i^\dagger)_{\mathsf{C_i'}} \otimes \mathbb{1}_{\mathsf{C_{i+1}}}) |\Omega\rangle_{\mathsf{C_i'C_{i+1}}} \ .$$

Define the following superoperator:

$$\mathcal{E}(\phi_{\mathsf{C_1C_1'\dots C_tC_t'}}) = \Big(\langle\Omega|_{C_0C_1} \langle\Omega_{A_1}|_{\mathsf{C_1'C_2}} \otimes \dots \otimes \langle\Omega_{A_{t-1}}|_{\mathsf{C_{t-1}'C_t}} \otimes (A_t)_{\mathsf{C_t'}}\Big) \left(|\psi\rangle\langle\psi|_{\mathsf{C_0}} \otimes \phi_{\mathsf{C_1C_1'\dots C_tC_t'}}\right)$$
$$\Big(|\Omega\rangle_{C_0C_1} |\Omega_{A_1}\rangle_{\mathsf{C_1'C_2}} \otimes \dots \otimes |\Omega_{A_{t-1}}\rangle_{\mathsf{C_{t-1}'C_t}} \otimes (A_t^\dagger)_{\mathsf{C_t'}}\Big) \ .$$

Let $S \subseteq [d]^t$, and overloading the notation define

$$|\Omega_S\rangle_{\mathsf{C_1C_1'\dots C_tC_t'}} = \sum_{(x_1,\dots,x_t)\in S} |x_1\rangle_{\mathsf{C_1}} |x_1\rangle_{\mathsf{C_1'}} \otimes \dots \otimes |x_t\rangle_{\mathsf{C_t}} |x_t\rangle_{\mathsf{C_t'}} \ .$$

Then for any collection of d-dimensional unitaries U_i,

$$\mathcal{E}\left(\Big(\mathbb{1}_{\mathsf{C_1}} \otimes (U_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (U_t)_{\mathsf{C_t'}}\Big) |\Omega_S\rangle\langle\Omega_S| \Big(\mathbb{1}_{\mathsf{C_1}} \otimes (U_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (U_t)_{\mathsf{C_t'}}\Big)^\dagger\right)$$
$$= \mathrm{proj}\left(\sum_{(x_1,\dots,x_t)\in S} A_tU_t |x_t\rangle\langle x_t| A_{t-1}U_{t-1} |x_{t-1}\rangle\langle x_{t-1}| \dots A_1U_1 |x_1\rangle\langle x_1| |\psi\rangle_{\mathsf{C_0}}\right) \ .$$

Here, $\mathrm{proj}(|\phi\rangle)$ *is shorthand for* $|\phi\rangle\langle\phi|$.

Proof. It suffices to show that

$$\Big(\langle\Omega|_{C_0C_1} \langle\Omega_{A_1}|_{\mathsf{C_1'C_2}} \otimes \dots \otimes \langle\Omega_{A_{t-1}}|_{\mathsf{C_{t-1}'C_t}} \otimes (A_t)_{\mathsf{C_t'}}\Big) \left(|\psi\rangle_{\mathsf{C_0}} \otimes \Big(\mathbb{1}_{\mathsf{C_1}} \otimes (U_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (U_t)_{\mathsf{C_t'}}\Big) |\Omega_S\rangle_{\mathsf{C_1C_1'\dots C_tC_t'}}\right)$$
$$= \sum_{(x_1,\dots,x_t)\in S} A_tU_t |x_t\rangle\langle x_t| A_{t-1}U_{t-1} |x_{t-1}\rangle\langle x_{t-1}| \dots A_1U_1 |x_1\rangle\langle x_1| |\psi\rangle_{\mathsf{C_0}} \ .$$

Inserting the definitions of $|\Omega_{A_i}\rangle$ and $|\Omega_S\rangle$, we get that the l.h.s. of the above equation is equal to

$$
\begin{aligned}
&\sum_{\substack{(i_0,\dots,i_t)\in[d]^t\\(x_1,\dots,x_t)\in S}} \Big(\langle i_0|_{\mathsf{C}_0} \langle i_0|_{\mathsf{C}_1} \langle i_1|_{\mathsf{C}_1'} \langle i_1|_{\mathsf{C}_2} \otimes \ldots \otimes \langle i_{t-1}|_{\mathsf{C}_{t-1}'} \langle i_{t-1}|_{\mathsf{C}_t} \Big)\Big(|\psi\rangle_{\mathsf{C}_0} \otimes \mathbb{1}_{\mathsf{C}_1} \otimes (A_1U_1)_{\mathsf{C}_1'} \otimes \ldots \otimes \mathbb{1}_{\mathsf{C}_t} \otimes (A_tU_t)_{\mathsf{C}_t'} \Big) \\
&\qquad\qquad \Big(|x_1\rangle_{\mathsf{C}_1} |x_1\rangle_{\mathsf{C}_1'} \otimes \ldots \otimes |x_t\rangle_{\mathsf{C}_t} |x_t\rangle_{\mathsf{C}_t'} \Big) \\
= &\sum_{\substack{(i_0,\dots,i_t)\in[d]^t\\(x_1,\dots,x_t)\in S}} \delta_{i_0,x_1}\cdots\delta_{i_{t-1},x_t} \Big(\langle i_0|_{\mathsf{C}_0} \langle i_1|_{\mathsf{C}_1'} \otimes \ldots \otimes \langle i_{t-1}|_{\mathsf{C}_{t-1}'} \Big)\Big(|\psi\rangle_{\mathsf{C}_0} \otimes (A_1U_1)_{\mathsf{C}_1'} \otimes \ldots \otimes (A_tU_t)_{\mathsf{C}_t'} \Big) \\
&\qquad\qquad \Big(|x_1\rangle_{\mathsf{C}_1'} \otimes \ldots \otimes |x_t\rangle_{\mathsf{C}_t'} \Big) \\
= &\sum_{(x_1,\dots,x_t)\in S} \langle x_1|\psi\rangle \langle x_2|A_1U_1|x_1\rangle \langle x_3|A_2U_2|x_2\rangle \cdots \langle x_t|A_{t-1}U_{t-1}|x_{t-1}\rangle \left(A_tU_t\,|x_t\rangle\right) \\
= &\sum_{(x_1,\dots,x_t)\in S} A_tU_t\,|x_t\rangle\!\langle x_t|\,A_{t-1}U_{t-1}\,|x_{t-1}\rangle\!\langle x_{t-1}|\ldots A_1U_1\,|x_1\rangle\!\langle x_1|\,|\psi\rangle\ .
\end{aligned}
$$

For the second line, we multiplied bras and kets on the non-primed systems $\mathsf{C}_1,\dots,\mathsf{C}_t$ and used orthonormality of the basis kets. □

With the above two lemmas in hand, we can now prove Theorem 6.2.

Proof of Theorem 6.2. First note that V_k is efficient to implement by the same reasoning as in Theorem 5.2, and because the $|+\rangle^{\otimes s(n)}$-state is efficient to prepare.

We need to prove that for any QPT algorithm $\mathcal{A}$, it holds that

$$
\left| \Pr_{k\sim\mathcal{K}} [\mathcal{A}^{V_k}(1^\lambda)=1] - \Pr_{V\sim\mathrm{Haar}(2^{n-s(n)},2^n)} [\mathcal{A}^{V}(1^\lambda)=1] \right| \le \mathrm{negl}(n)\,.
$$

As in the proof of Theorem 5.2, we can replace the pseudorandom isometry V_k by a "fully random" isometry, where instead of using pseudorandom permutations and phases we sample uniformly random permutations and phases, i.e. we sample from the PF ensemble and then fix the last $s(n)$ qubits to be $|+\rangle$. We denote an isometry sampled from this ensemble as $V\sim(PF)_{s(n)}$. Since the algorithm $\mathcal{A}$ is computationally efficient and we use quantum-secure pseudorandom permutations and functions, it follows that

$$
\left| \Pr_{k\sim\mathcal{K}} [\mathcal{A}^{V_k}(1^\lambda)=1] - \Pr_{V\sim(PF)_{s(n)}} [\mathcal{A}^{V}(1^\lambda)=1] \right| \le \mathrm{negl}(n)\,.
$$

By purifying the operations of $\mathcal{A}$, for any fixed choice of isometry V we can describe the algorithm's operation as follows: first prepare the state

$$
|A(V)\rangle \coloneqq A_t(V\otimes I)A_{t-1}(V\otimes I)\cdots A_1(V\otimes I)\,|A_0\rangle \tag{6.1}
$$

(for some choice of unitaries $A_1,\dots,A_t$, which describe $\mathcal{A}$'s actions in between querying the isometry, and an arbitrary initial state $|A_0\rangle$) and then perform a binary measurement $\{M, \mathbb{1}-M\}$, with "M" denoting the "1"-outcome. Then for any distribution $\mathcal{V}$ of isometries,

$$
\Pr_{V\sim\mathcal{V}}[\mathcal{A}^V=1] = \mathrm{Tr}\left[M \mathop{\mathbb{E}}_{V\sim\mathcal{V}} |A(V)\rangle\!\langle A(V)| \right].
$$

Therefore to prove the theorem, it suffices to show that

$$
\left\| \mathop{\mathbb{E}}_{V\sim(PF)_{s(n)}} |A(V)\rangle\!\langle A(V)| - \mathop{\mathbb{E}}_{V\sim\mathrm{Haar}(2^{n-s(n)},2^n)} |A(V)\rangle\!\langle A(V)| \right\|_1 = \mathrm{negl}(n)\,. \tag{6.2}
$$

To proceed with the proof, we need to introduce some additional notation in Eq. (6.1). We can without loss of generality assume that $\mathcal{A}$'s space is made up of the following registers: an $(n - s(n))$-qubit register A, which is the input register for each application of the isometry; $s(n)$-qubit registers $\mathsf{B_1}, \dots, \mathsf{B_t}$, where $\mathsf{B_i}$ is the additional output register produced by the i-the call to the isometry (i.e. the i-call to the isometry is a map $\mathsf{A} \to \mathsf{AB_i}$); and an arbitrary workspace register R. The final state $|A(V)\rangle$ is then a state on registers $\mathsf{ARB_1 \dots B_t}$.

We can view the i-th call to the isometry $V : \mathsf{A} \to \mathsf{AB_i}$ as a call to a unitary $U : \mathsf{AB_i} \to \mathsf{AB_i}$, with the input on the $\mathsf{B_i}$ register fixed to $|+^{s(n)}\rangle_{\mathsf{B_i}}$. This holds for both our construction of V, where we have the corresponding unitary $U = FP$, and a Haar random isometry V, where the corresponding unitary U is Haar random. Then we can view the additional $|+^{s(n)}\rangle_{\mathsf{B_i}}$-states as part of the input and write

$$|A(V)\rangle \coloneqq A_t(U_{\mathsf{AB_t}} \otimes \mathbb{1}_{\mathsf{RB_{\setminus t}}})A_{t-1}(U_{\mathsf{AB_{t-1}}} \otimes \mathbb{1}_{\mathsf{RB_{\setminus t-1}}}) \cdots A_1(U_{\mathsf{AB_1}} \otimes \mathbb{1}_{\mathsf{RB_{\setminus 1}}})(|A_0\rangle_{\mathsf{AR}} \otimes |+^{s(n)}\rangle_{\mathsf{B_1}} \otimes \cdots \otimes |+^{s(n)}\rangle_{\mathsf{B_t}}) .$$

Here, $B_{\setminus \mathsf{i}}$ is shorthand for $\mathsf{B_1 \dots B_{i-1} B_{i+1} \dots B_t}$.

Reduction to distinct string inputs. As in the proof of Theorem 3.1, the next step is to restrict the inputs to the unitaries in $|A(V)\rangle$ to distinct strings. The notion of the distinct string subspace is less clear for adaptive queries; we comment on what this means in the case of PRUs in Section 6.1. For the simpler case of isometries, we can simply project the fixed inputs $|+^{s(n)}\rangle_{\mathsf{B_1}} \otimes \cdots \otimes |+^{s(n)}\rangle_{\mathsf{B_t}}$ onto the distinct string subspace.

Formally, let $\Lambda_{\mathsf{B_1 \dots B_t}}$ be the projector onto the distinct string subspace of $\mathsf{B_1 \dots B_t}$ and let $|+\rangle_{\mathsf{B_1 \dots B_t}} \coloneqq |+^{s(n)}\rangle_{\mathsf{B_1}} \otimes \cdots \otimes |+^{s(n)}\rangle_{\mathsf{B_t}}$ be the uniform superposition on registers $\mathsf{B_1 \dots B_t}$. Then, since $s(n) = \omega(\log n)$, it follows that

$$\left\| |+\rangle_{\mathsf{B_1 \dots B_t}} - \Lambda_{\mathsf{B_1 \dots B_t}} |+\rangle_{\mathsf{B_1 \dots B_t}} \right\|_2 = \mathrm{negl}(n) ,$$

because the probability of observing a collision (i.e. the probability of getting the same computational basis string in any two registers in $\mathsf{B_1 \dots B_t}$) when measuring $|+\rangle_{\mathsf{B_1 \dots B_t}}$ is negligible in n. This implies that

$$\| |A(V)\rangle - |A(V)_\Lambda\rangle \| = \mathrm{negl}(n) ,$$

where we defined

$$|A(V)_\Lambda\rangle \coloneqq A_t(U_{\mathsf{AB_t}} \otimes I)A_{t-1}(U_{\mathsf{AB_{t-1}}} \otimes I) \cdots A_1(U_{\mathsf{AB_1}} \otimes I)(|A_0\rangle_{\mathsf{AR}} \otimes \Lambda_{\mathsf{B_1 \dots B_t}} |+\rangle_{\mathsf{B_1 \dots B_t}}) .$$

Consequently, it suffices to show that

$$\left\| \mathop{\mathbb{E}}_{V \sim (PF)_{s(n)}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| - \mathop{\mathbb{E}}_{V \sim \mathrm{Haar}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| \right\|_1 = \mathrm{negl}(n) , \tag{6.3}$$

which then implies Eq. (6.2).

Distinct state as output of gate teleportation. The next step is to write the state $|A(V)_\Lambda\rangle$ as the output of the gate teleportation map from Lemma 6.4. For this, we first observe that we can rewrite

$$|A(V)_\Lambda\rangle = \sum_{(y_1,\dots,y_t) \in \mathrm{distinct}(s(n),t)} A_t(U_{\mathsf{AB_t}} |y_t\rangle\langle y_t|_{\mathsf{B_t}} \otimes I)A_{t-1}(U_{\mathsf{AB_{t-1}}} |y_{t-1}\rangle\langle y_{t-1}|_{\mathsf{B_{t-1}}} \otimes I)$$
$$\cdots A_1(U_{\mathsf{AB_1}} |y_1\rangle\langle y_1|_{\mathsf{B_1}} \otimes I)(|A_0\rangle_{\mathsf{AR}} \otimes |+\rangle_{\mathsf{B_1 \dots B_t}}) . \tag{6.4}$$

To make this look like an output state from Lemma 6.4, in each step we can insert identities on all systems except the one that $|y_i\rangle\langle y_i|$ is acting on. More formally, we have that

$$|A(V)_\Lambda\rangle = \sum_{(x_1,\dots,x_t) \in S} A_t(U_{\mathsf{AB_t}} \otimes I) |x_t\rangle\langle x_t| A_{t-1}(U_{\mathsf{AB_{t-1}}} \otimes I) |x_{t-1}\rangle\langle x_{t-1}|$$
$$\cdots A_1(U_{\mathsf{AB_1}} \otimes I) |x_1\rangle\langle x_1| (|A_0\rangle_{\mathsf{AR}} \otimes |+\rangle_{\mathsf{B_1 \dots B_t}})$$

for the following subset of strings (where $D = |R| + |A| + |B_1| + \cdots + |B_t|$ is the total dimension of the space used by algorithm $\mathcal{A}$):

$$S = \left\{ (x_1, \dots, x_t) \in [D]^t \mid ((x_1)_{[\mathsf{B_1}]}, \dots, (x_t)_{[\mathsf{B_t}]}) \in \mathrm{distinct}(s(n), t) \right\}, \tag{6.5}$$

where $(x_i)_{[\mathsf{B_i}]}$ denotes the substring of x_i corresponding to register B_i. In other words, S contains tuples of strings; each string x_i can be associated to a basis ket $|x_i\rangle_{\mathsf{RAB_1 \dots B_t}}$ on $\mathcal{A}$'s total workspace; and S includes all tuples such that when we only look at the part of x_i that corresponds to register $\mathsf{B_i}$ (which is where $|y_i\rangle\langle y_i|$ acts in Eq. (6.4)), all of these substrings are distinct (since all y_i in Eq. (6.4) are distinct). Using the shorthand $W_i = U_{\mathsf{AB_i}} \otimes \mathbb{1}_{\mathsf{B_{\setminus i} \otimes \mathbb{1}_R}}$, we can write this more compactly as

$$|A(V)_\Lambda\rangle = \sum_{(x_1,\dots,x_t) \in S} A_t W_t |x_t\rangle\langle x_t| A_{t-1} W_{t-1} |x_{t-1}\rangle\langle x_{t-1}| \cdots A_1 W_1 |x_1\rangle\langle x_1| \left(|A_0\rangle_{\mathsf{AR}} \otimes |+\rangle_{\mathsf{B_1 \dots B_t}}\right)$$

It then follows from Lemma 6.4 that

$$|A(V)_\Lambda\rangle\langle A(V)_\Lambda| = \mathcal{E}\left(\left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right) |\Omega_S\rangle\langle\Omega_S| \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right)^\dagger \right), \tag{6.6}$$

where $\mathcal{E}$ is the gate teleportation channel defined in Lemma 6.4 and $\mathsf{C_i'} \equiv \mathsf{C_i} \equiv \mathsf{ARB_1 \dots B_t}$ are copies of the total workspace.

Using the relative-error property on the distinct subspace. Observe that in the state

$$\left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right) |\Omega_S\rangle\langle\Omega_S| \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right)^\dagger$$

from Eq. (6.6), the different $U_{\mathsf{AB_i}}$ (which appear, tensored with identity, in W_i) act on distinct strings. This is ensured by the fact that in the definition of the set S (Eq. (6.5)), we have the condition that the substrings $(x_i)_{[\mathsf{B_i}]}$ on the systems $\mathsf{B_i}$ (which is part of the input system to $U_{\mathsf{AB_i}}$) are distinct. It therefore follows from Lemma 6.3 that

$$\begin{aligned} &\mathop{\mathbb{E}}_{U \sim PF} \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right) |\Omega_S\rangle\langle\Omega_S| \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right)^\dagger \\ \leq (1 + O(t^2/d)) &\mathop{\mathbb{E}}_{U \sim \mathrm{Haar}} \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right) |\Omega_S\rangle\langle\Omega_S| \left(\mathbb{1}_{\mathsf{C_1}} \otimes (W_1)_{\mathsf{C_1'}} \otimes \dots \mathbb{1}_{\mathsf{C_t}} \otimes (W_t)_{\mathsf{C_t'}} \right)^\dagger, \end{aligned}$$

where $W_i = U_{\mathsf{AB_i}} \otimes I$ as before.

We can insert this into Eq. (6.6) and use the fact that the gate teleportation channel $\mathcal{E}$ is manifestly completely positive to find that

$$\mathop{\mathbb{E}}_{V \sim (PF)_{s(n)}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| \leq (1 + O(t^2/d)) \mathop{\mathbb{E}}_{V \sim \mathrm{Haar}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| \,. \tag{6.7}$$

To see how this implies Eq. (6.3), note that using the variational definition of trace distance, there exists an M such that $0 \leq M \leq \mathbb{1}$ and[10]

$$\begin{aligned} \text{l.h.s. of Eq. (6.3)} &= \mathrm{Tr}\left[M \left(\mathop{\mathbb{E}}_{V \sim (PF)_{s(n)}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| - \mathop{\mathbb{E}}_{V \sim \mathrm{Haar}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| \right) \right] \\ &\leq O(t^2/d)\, \mathrm{Tr}\left[M \mathop{\mathbb{E}}_{V \sim \mathrm{Haar}} |A(V)_\Lambda\rangle\langle A(V)_\Lambda| \right] \\ &\leq O(t^2/d)\,. \end{aligned}$$

[10] We do not need absolute value signs because the expression in parentheses is the difference between two quantum states and therefore traceless. Consequently we can always replace $M \mapsto \mathbb{1} - M$ to switch the sign of the trace expression.

The first inequality follows by inserting Eq. (6.7) and remembering that M is positive semi-definite, so the map $X \mapsto \mathrm{Tr}[MX]$ is operator-monotone. The second inequality uses the fact that the trace expression is at most 1, which holds because $M \leq \mathbb{1}$ and $\mathbb{E}_{V\sim\mathrm{Haar}} |A(V)_\Lambda\rangle\!\langle A(V)_\Lambda|$ is a quantum state. This proves Eq. (6.3) and completes the proof. □

6.1 Towards PRUs with adaptive security

We briefly comment on possible ways to extend the proof of Theorem 6.2 to adaptively secure PRUs. The reason why the current proof of Theorem 6.2 only works for isometries is the following. We can analyse the state $|A(V)_\Lambda\rangle$ using the relative error property of the PF ensemble (Lemma 6.3). This part of the analysis is general and does not require fixing part of the input state to the unitary to $|+\rangle$. The part of the analysis that limits us to isometries is relating the final state $|A(V)\rangle$ of an adaptive algorithm to the state $|A(V)_\Lambda\rangle$, where all the queries to V (or rather its unitary extension $U = PF$) are restricted to distinct strings. If we fix part of the input to the $|+\rangle$-state as we do in Theorem 6.2, this fixed part of the input already ensures that the unitaries $U = PF$ are only queried on distinct strings (except with negligible weight). This is the case because for $s(n) = \omega(\log n)$, the "collision probability" (i.e. the probability of getting the same computational basis string) when measuring $|+\rangle^{\otimes s(n)}$ is negligible in n.

To extend this to PRUs, we need to be able to ensure that no adaptive query algorithm queries the PF-ensemble twice with non-negligible weight on the same computational basis string. The natural approach to this is to prepend the PF ensemble with another ensemble of unitaries that sufficiently scrambles the input state to make sure that the PF ensemble is not queried twice on the same computational basis string. In the non-adaptive case, this role was played by a random Clifford unitary. Unfortunately, it is unclear whether random Clifford unitaries still work for this purpose in the adaptive case.

More formally, we make the following conjecture.

Conjecture 6.5. *There exists an ensemble of efficient*[11] *n-qubits unitaries $W \sim \mathcal{W}$ such that for all $t = \mathrm{poly}(n)$ and all initial states $|B_0\rangle$ and sequences of unitaries $B_1, \ldots, B_t$ on $n + \mathrm{poly}(n)$ qubits,*

$$\left\| \mathop{\mathbb{E}}_{W\in\mathcal{W}} |B(W)\rangle\!\langle B(W)| - \mathop{\mathbb{E}}_{W\in\mathcal{W}} |B(W)_\Lambda\rangle\!\langle B(W)_\Lambda| \right\|_1 = \mathrm{negl}(n)\,,$$

where

$$\begin{aligned} |B(W)\rangle &= B_t(W \otimes \mathbb{1}_m)B_{t-1}(W \otimes \mathbb{1}_m)\cdots B_1(W \otimes \mathbb{1}_m)\,|B_0\rangle\,, \\ |B(W)_\Lambda\rangle &= \sum_{(x_1,\ldots,x_t)\in\mathrm{distinct}(n,t)} B_t(|x_t\rangle\!\langle x_t|\, W \otimes \mathbb{1}_m)B_{t-1}(|x_{t-1}\rangle\!\langle x_{t-1}|\, W \otimes \mathbb{1}_m)\cdots B_1(|x_1\rangle\!\langle x_1|\, W \otimes \mathbb{1}_m)\,|B_0\rangle\,. \end{aligned}$$

Assuming Conjecture 6.5, it follows from the proof of Theorem 6.2 that the random unitary PFW, with P and F as before and $W \sim \mathcal{W}$, is a PRU with adaptive security: we can use the conjecture to perform the "reduction to distinct string inputs" in the proof of Theorem 6.2 (with $B_i = A_i(PF \otimes I)$) and the rest of the proof goes through unchanged. We leave it as an interesting open problem to prove Conjecture 6.5.

References

[ABF+24] Scott Aaronson, Adam Bouland, Bill Fefferman, Soumik Ghosh, Umesh Vazirani, Chenyi Zhang, and Zixin Zhou. Quantum pseudoentanglement. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2024.

[ABI86] Noga Alon, Laslo Babai, and Alon Itai. A fast and simple randomized parallel algorithm for the maximal independent set problem. *Journal of Algorithms*, 7(4):567–583, 1986.

[11] In fact, it suffices if the ensemble is computationally indistinguishable from an ensemble of computationally efficient unitaries. Then we can replace the inefficient ensemble by the efficient one as in the proof of Theorem 5.2.

[AE07] Andris Ambainis and Joseph Emerson. Quantum t-designs: t-wise independence in the quantum world. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 129–140. IEEE, 2007.

[AGKL23] Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries. *arXiv preprint arXiv:2311.02901*, 2023.

[AL13] Noga Alon and Shachar Lovett. Almost k-wise vs. k-wise independent permutations, and uniformity for general group actions. *Theory of Computing*, 9(15):559–577, 2013.

[AQY22] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In *Annual International Cryptology Conference*, pages 208–236. Springer, 2022.

[BCHJ$^+$21] Fernando GSL Brandão, Wissam Chemissany, Nicholas Hunter-Jones, Richard Kueng, and John Preskill. Models of quantum complexity growth. *PRX Quantum*, 2(3):030316, 2021.

[BF12] Winton Brown and Omar Fawzi. Scrambling speed of random quantum circuits. *arXiv preprint arXiv:1210.6644*, 2012.

[BHH16] Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.

[BL20] Anne Broadbent and Sébastien Lord. Uncloneable quantum encryption via oracles. In *15th Conference on the Theory of Quantum Computation, Communication and Cryptography*, 2020.

[BM21] Sergey Bravyi and Dmitri Maslov. Hadamard-free circuits expose the structure of the clifford group. *IEEE Transactions on Information Theory*, 67(7):4546–4563, July 2021.

[BM24] Zvika Brakerski and Nir Magrafta. Real-valued somewhat-pseudorandom unitaries. *arXiv preprint arXiv:2403.16704*, 2024.

[Bum13] D. Bump. *Lie Groups*. Graduate Texts in Mathematics. Springer New York, 2013.

[CBB$^+$24] Chi-Fang Chen, Adam Bouland, Fernando G. S. L. Brandao, Jordan Docter, Patrick Hayden, and Michelle Xu. Efficient unitary designs and pseudorandom unitaries from permutations. In preparation, 2024.

[CDX$^+$24] Chi-Fang Chen, Jordan Docter, Michelle Xu, Adam Bouland, and Patrick Hayden. Efficient unitary t-designs from random sums. *arXiv preprint arXiv:2402.09335*, 2024.

[CHH$^+$24] Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan. In preparation, 2024.

[Chr06] Matthias Christandl. The Structure of Bipartite Quantum States-Insights from Group Theory and Cryptography. *Ph. D. Thesis*, 2006.

[CK23] Pierre-Emmanuel Caprace and Martin Kassabov. Tame automorphism groups of polynomial rings with property (t) and infinitely many alternating group quotients. *Transactions of the American Mathematical Society*, 376(11):7983–8021, 2023.

[CLLW16] Richard Cleve, Debbie Leung, Li Liu, and Chunhao Wang. Near-linear constructions of exact unitary 2-designs. *Quantum Info. Comput.*, 16(9–10):721–756, jul 2016.

[CLLZ21] Andrea Coladangelo, Jiahui Liu, Qipeng Liu, and Mark Zhandry. Hidden cosets and applications to unclonable cryptography. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part I 41*, pages 556–584. Springer, 2021.

[CMP20] Andrea Coladangelo, Christian Majenz, and Alexander Poremba. Quantum copy-protection of compute-and-compare programs in the quantum random oracle model. *arXiv preprint arXiv:2009.13865*, 2020.

[DCEL09] 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):012304, 2009.

[EFL$^+$24] Netta Engelhardt, Åsmund Folkestad, Adam Levine, Evita Verheijden, and Lisa Yang. Cryptographic censorship. *arXiv preprint arXiv:2402.03425*, 2024.

[FH13] William Fulton and Joe Harris. *Representation theory: a first course*, volume 129. Springer Science & Business Media, 2013.

[GAE07] David Gross, Koenraad Audenaert, and Jens Eisert. Evenly distributed unitaries: On the structure of unitary designs. *Journal of mathematical physics*, 48(5), 2007.

[Haf22] Jonas Haferkamp. Random quantum circuits are approximate unitary t-designs in depth $O(nt^{5+o(1)})$. *Quantum*, 6:795, 2022.

[Har13] Aram W Harrow. The church of the symmetric subspace. *arXiv preprint arXiv:1308.6595*, 2013.

[HBC$^+$22] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, et al. Quantum advantage in learning from experiments. *Science*, 376(6598):1182–1186, 2022.

[HBK23] Tobias Haug, Kishor Bharti, and Dax Enshan Koh. Pseudorandom unitaries are neither real nor sparse nor noise-robust. *arXiv preprint arXiv:2306.11677*, 2023.

[HLT24] Jeongwan Haah, Yunchao Liu, and Xinyu Tan. Efficient approximate unitary designs from random pauli rotations. *arXiv preprint arXiv:2402.05239*, 2024.

[HP07] Patrick Hayden and John Preskill. Black holes as mirrors: quantum information in random subsystems. *Journal of High Energy Physics*, 2007(09):120–120, September 2007.

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *Advances in Cryptology–CRYPTO 2018: 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19–23, 2018, Proceedings, Part III 38*, pages 126–152. Springer, 2018.

[Jof74] A. Joffe. On a Set of Almost Deterministic k-Independent Random Variables. *The Annals of Probability*, 2(1):161 – 162, 1974.

[Kas07] Martin Kassabov. Symmetric groups and expander graphs. *Inventiones mathematicae*, 170(2):327–354, 2007.

[KLR$^+$08] Emanuel Knill, Dietrich Leibfried, Rolf Reichle, Joe Britton, R Brad Blakestad, John D Jost, Chris Langer, Roee Ozeri, Signe Seidelin, and David J Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, 2008.

[KP23] Isaac H. Kim and John Preskill. Complementarity and the unitarity of the black hole S-matrix. *Journal of High Energy Physics*, 2023(2), February 2023.

[Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. *arXiv preprint arXiv:2103.09320*, 2021.

[LQS$^+$23] Chuhan Lu, Minglong Qin, Fang Song, Penghui Yao, and Mingnan Zhao. Quantum pseudorandom scramblers. *arXiv preprint arXiv:2309.08941*, 2023.

[Mel23] Antonio Anna Mele. Introduction to Haar measure tools in quantum information: A beginner's tutorial. *arXiv preprint arXiv:2307.08956*, 2023.

[MGE12] Easwar Magesan, Jay M Gambetta, and Joseph Emerson. Characterizing quantum gates via randomized benchmarking. *Physical Review A*, 85(4):042311, 2012.

[MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Pseudorandom unitaries with non-adaptive security. *arXiv preprint arXiv:2402.14803*, 2024.

[MY22] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Annual International Cryptology Conference*, pages 269–295. Springer, 2022.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition.* Cambridge University Press, 2010.

[OSP23] Ryan O'Donnell, Rocco A Servedio, and Pedro Paredes. Explicit orthogonal and unitary designs. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1240–1260. IEEE, 2023.

[Sus16] Leonard Susskind. Computational complexity and black hole horizons. *Fortschritte der Physik*, 64(1):24–43, 2016.

[vdB21] Ewout van den Berg. A simple method for sampling random Clifford operators. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 54–59. IEEE, 2021.

[Win99] A. Winter. Coding theorem and strong converse for quantum channels. *IEEE Transactions on Information Theory*, 45(7):2481–2485, 1999.

[Zha12] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology – CRYPTO 2012*, pages 758–775, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

[Zha16] Mark Zhandry. A note on quantum-secure PRPs. *arXiv preprint arXiv:1611.05564*, 2016.

[Zha21] Mark Zhandry. How to construct quantum random functions. *Journal of the ACM (JACM)*, 68(5):1–43, 2021.