(Pseudo) Random Quantum States with Binary Phase*

Zvika Brakerski[1] and Omri Shmueli[2]

[1] Weizmann Institute of Science**
[2] Tel-Aviv University***

Abstract. We prove a quantum information-theoretic conjecture due to Ji, Liu and Song (CRYPTO 2018) which suggested that a uniform superposition with random *binary* phase is statistically indistinguishable from a Haar random state. That is, any polynomial number of copies of the aforementioned state is within exponentially small trace distance from the same number of copies of a Haar random state.
As a consequence, we get a provable elementary construction of *pseudorandom* quantum states from post-quantum pseudorandom functions. Generating pseudorandom quantum states is desirable for physical applications as well as for computational tasks such as quantum money.
We observe that replacing the pseudorandom function with a $(2t)$-wise independent function (either in our construction or in previous work), results in an explicit construction for *quantum state t-designs* for all t. In fact, we show that the circuit complexity (in terms of both circuit size and depth) of constructing t-designs is bounded by that of $(2t)$-wise independent functions. Explicitly, while in prior literature t-designs required linear depth (for $t > 2$), this observation shows that polylogarithmic depth suffices for all t.
We note that our constructions yield pseudorandom states and state designs with only real-valued amplitudes, which was not previously known. Furthermore, generating these states require quantum circuit of restricted form: applying one layer of Hadamard gates, followed by a sequence of Toffoli gates. This structure may be useful for efficiency and simplicity of implementation.

1 Introduction

Randomness is one of the most fundamental resources for computation, and is indispensable for algorithms, complexity theory and cryptography. It is also a

* The full version of this paper is available at https://arxiv.org/abs/1906.10611.

** Email: zvika.brakerski@weizmann.ac.il. Supported by the Israel Science Foundation (Grant No. 468/14), Binational Science Foundation (Grants No. 2016726, 2014276), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).

*** Email: omrishmueli@mail.tau.ac.il. Supported by the Zevulun Hammer Scholarship from the Council for Higher Education in Israel, and by Israel Science Foundation Grant No. 18/484, and by Len Blavatnik and the Blavatnik Family Foundation.

foundational tool for science in general, for purposes of describing and modeling natural phenomena. As our understanding of nature expands to quantum phenomena, the importance of understanding the uniform distribution over quantum states, and being able to sample from it, naturally emerges.

Quantum states can be described as unit vectors in a high-dimensional complex Hilbert space. Thus, a random quantum state is just a random unit vector on this abstract sphere. This distribution is also referred to as the Haar measure over quantum states. We note that this is a continuous distribution, even if the Hilbert space is finite dimensional (i.e. can be described by a finite number of qubits). Since quantum states cannot be duplicated, the ability to generate random quantum states refers to the ability to generate multiple copies of the same random state vector. (In fact, a single copy of a quantum random state is identical to a classical random state.) Haar random quantum states have numerous computational and physical applications. The former includes optimal quantum communication channels [8], efficient quantum POVM measurements [14] which are in turn useful in quantum state tomography, and gate fidelity estimation [4]. The latter includes constructing physical models of quantum thermalization [13].

Since random states have infinitely long descriptions (and super-exponential even if restricting to some finite precision), there is extensive literature studying approximate notions and specifically the notion of ϵ-approximate t-designs. These are distributions whose t-tensor (i.e. taking t copies of a sample from this distribution) are ϵ indistinguishable from (a t-tensor of) Haar (using the standard notion for statistical indistinguishability known as trace distance). We adopt the standard asymptotic convention and require by default that ϵ is negligible in our "security parameter", which we associate with the logarithm of the dimension of the Hilbert space. In this work we focus on quantum states over n qubits (i.e. 2^n dimensional Hilbert space), so we associate our security parameter with n. However, our methods are extendable to any finite-dimensional space (with efficient representation). There is extensive literature studying (approximate) designs with bounded t, which also carry physical significance, see e.g. [2,4,5,10,9,7]. Indeed, it is possible to efficiently generate t-designs using quantum circuits of size $\mathrm{poly}(t, n)$. Up to asymptotics, this matches the information theoretic bound (however, the important aspect of the depth complexity of generating t-designs remained open, to the best of our knowledge), and one cannot hope to efficiently generate t-designs for super-polynomial t.

Asymptotically Random States, Pseudorandom States and the JLS Conjecture. Ji, Liu and Song [6] (henceforth JLS) recently proposed to extend the notion of approximate designs. They proposed the notion of a *pseudorandom quantum state* (PRS) which has a finite description but is *computationally indistinguishable* from Haar given a t-tuple, for *any* $t = \mathrm{poly}(n)$. Thus, for any computationally bounded purpose (experiment, naturally occurring process) a PRS is indistinguishable from a Haar state, regardless of the number of copies. They also showed that PRS are useful for cryptographic applications such as quantum money.

Furthermore, [6] proposed an insightful template for constructing PRS. They start by showing that given quantum RAM access to exponentially many classical random bits, it is possible to construct a $\mathrm{negl}(n)$-approximate $n^{\omega(1)}$-design. Let us call such a distribution ARS, for Asymptotically Random State.[3] An ARS is a statistical notion of PRS which has asymptotic limitations but no computational restrictions. Then, replacing the exponential random string with a quantum-query-resistant classically-computable pseudorandom function (PRF), the PRS construction naturally follows from ARS. The existence of such PRF is implied by the existence of quantum secure one-way functions [15].

The ARS construction of JLS is quite straightforward to describe. Generate a uniform superposition over all strings $x \in \{0,1\}^n$. This is described in the standard Dirac notation as $\sum_x |x\rangle$ (with some normalization factor). Then, assign a random quantum phase to each component x, i.e. generate $\sum_x \alpha_x |x\rangle$ for random independent roots of unity α_x. To cope with finite precision, α_x is taken to a finite but exponential resolution $\alpha_x = \omega_{2^n}^{f(x)}$, where $f : \{0,1\}^n \to [2^n]$ is a random function and ω_{2^n} is the 2^n-th root of unity. Given RAM access to the truth table of f, this state can be efficiently computed using Quantum Fourier Transform (QFT) modulo 2^n.

JLS then conjecture (but were unable to prove) that a much simpler construction, where $\alpha_x = (-1)^{f(x)}$, should also imply ARS. That is, replacing the "high-resolution" random phase, by the simplest binary phase. While this is only one of a few conjectures made in that work, it is the only one relevant to our work and we thus refer to it simply as the JLS conjecture.

Conjecture 1.1 ([6], restated). The distribution over n-qubit quantum states defined by

$$2^{-n/2} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle$$

where $f : \{0,1\}^n \to \{0,1\}$ is a random function, is an ARS.

To highlight the gap between the conjecture and the provable ARS construction of JLS, let us describe a crucial point in the analysis of JLS. The analysis is based on an equivalence relation between t-tuples of n-bit strings, which naturally arises from the expression for statistical distance from Haar. The tuples $(x_1, \ldots, x_t)$, $(y_1, \ldots, y_t)$ are equivalent if their histograms (i.e. the number of times each n-bit string appears) are equal modulo 2^n. Since $t < 2^n$ this condition is equivalent to requiring that the tuples are permutations of each other, which makes it possible to analyze the equivalence classes of this relation and for the analysis to go through.

In the binary setting, the equivalence relates tuples whose histograms are equal modulo 2. Thus the equivalence classes can no longer be described simply as a set and all of its permutations, and they don't even have the same size

[3] Actually, their ARS, as well as the one proven in this work, is even stronger: they show that for all t, their distribution is $O(t^2)/2^n$-approximate t-design.

anymore. This creates many additional terms in the so called density matrix of the state (which is a complex matrix of exponential dimensions $2^{tn} \times 2^{tn}$). In order to prove the conjecture, one will have to show that the effect of these exponentially many new terms on the spectrum of the matrix is negligible and there seems to be no straightforward handle for this analysis. We resolve this problem in this work.

Our Results – Proving the Conjecture. We prove the JLS conjecture, in fact we prove that the binary ARS implied by the conjecture has comparable properties to the prior construction (that used complex phase).

Theorem 1.2 (Main Result). *The distribution over n-qubit quantum states defined by*

$$2^{-n/2} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle$$

where $f : \{0,1\}^n \to \{0,1\}$ is a random function, is a $\frac{4t^2}{2^n}$-approximate t-design for all t, and thus an ARS.

This result has various implications that we describe below. We furthermore hope that our techniques will be useful for analyzing similarly complicated quantum states.

We make two additional observations that refer to the requirement from a function f to be plugged into either our theorem or that of JLS in order to imply PRS and quantum t-designs.

1. If we wish to obtain a PRS, the requirement of using a full-fledged quantum secure PRF can be relaxed. In fact, it is sufficient to have a function f that is indistinguishable from random while allowing only uniform superposition queries (as opposed to arbitrary superposition queries). This leads to a quantum notion which is somewhat analogous to the classical notion of weak pseudorandom functions [11], an object that can be of interest for independent investigation and possibly more efficient constructions than PRFs.
2. If we only wish to obtain a t-design, it is sufficient to replace f with a $(2t)$-wise independent function, using the fact that given t-quantum-query access, a $(2t)$-wise independent function is perfectly indistinguishable from a completely random function [15].

Implications. We find the JLS conjecture compelling from aesthetic, conceptual and perhaps even practical reasons. In terms of aesthetics, it is bothersome that one would need to go into exponentially fine-grained resolution on the phase in order to generate an ARS/PRS, being able to achieve the same parameters with a more coarse resolution (and as we show next without compromising on parameters) seems to be a more desirable state of affairs. Conceptually, the result shows that ARS, which is for all efficiently observable purposes identical to

a Haar random state, can be generated using only real-valued phases. Recalling that the Haar distribution is defined over complex vectors, it appears not obvious that it can be approximated for all observable purposes by real-valued vectors.

In terms of computational complexity, our construction uses circuits with restricted structure known in the literature as **HT** [12]. Concretely, the circuit contains a single parallel layer of Hadamard gates, followed by a circuit of Toffoli gates. This model is considered fairly weak, and in particular **HT** circuits are weakly classically simulatable (i.e. any distribution samplable by an **HT** circuit followed by measurement is also classically samplable). Result shows that even such a restricted model of quantum computation is enough to approximate the Haar measure.

Lastly, from a practical standpoint, replacing the function f by an efficient quantum-resilient PRF yields a very simple construction of a PRS, requiring only an **HT** circuit with the same circuit size and depth (up to asymptotics) as that of the PRF. Prior provable PRS candidates do not enjoy this property and appear to require a more complicated implementation (that in particular seem to need performing the Quantum Fourier Transform modulo 2^n, or a similar procedure) to allow for the high-resolution of complex phase.

In the context of generating t-designs, using our aforementioned observation and replacing f with a $(2t)$-wise independent function (in either our theorem or JLS) implies a t-design construction with circuit size $\mathrm{poly}(t, n)$ and depth $O(\log t \cdot \log n)$. We are not aware of prior constructions of t designs with $o(n)$ depth for $t > 2$ in the literature. Moreover, the t-design construction which is implied by our result can be implemented by an **HT** circuit with the same circuit size and depth (up to asymptotics) as that of the $(2t)$-wise independent function.

Proof High-Level Overview. Formally speaking, the proof follows by bounding the trace norm of the difference between the density matrix of t-copies of the state with binary phase, and the density matrix of t-copies of the state with 2^n roots of unity. However, one needs not know much about density matrices, it suffices to say that we have a complex Hermitian matrix of dimensions $2^{tn} \times 2^{tn}$, where the sum of all eigenvalues is 0, and we want to bound the sum of all absolute values of eigenvalues. It is thus sufficient to consider only positive or only negative eigenvalues.

Each row of the matrix corresponds to a tuple $(x_1, \ldots, x_t) \in (\{0,1\}^n)^t$ and each column corresponds to a tuple $(y_1, \ldots, y_t) \in (\{0,1\}^n)^t$. The entry in location $(x_1, \ldots, x_t), (y_1, \ldots, y_t)$ is nonzero if the aforementioned "histogram condition" holds on the tuples.[4] In a bit more detail, up to a global 2^{-tn} scaling factor, if the modulo-2 histogram condition holds but the modulo-2^n condition (i.e. permutation) does not hold then the entry will be 1, but if both hold then there is a cancellation and the entry will be 0.

[4] Recall that the (modulo-2) histogram condition states that $(x_1, \ldots, x_t), (y_1, \ldots, y_t)$ are equivalent if for all z, the number of times z appears in the first tuple and the number of times it appears in the second tuple have the same parity.

We start by observing that the matrix can be decomposed into "combinatorial blocks", each representing an equivalence class of the histogram relation. We analyze the properties of these blocks. We then provide two structural lemmas that together imply the theorem:

1. We provide a non-trivial upper bound on the rank of the matrix. While it is tempting to disregard the cancellations and just count the number of nonzero blocks and their respective rank, this implies an upper bound that is too coarse. We must therefore carefully take into account the cancellations induced by permutations in order to obtain a usable bound.
2. We provide an upper bound on the absolute value of each negative eigenvalue. We do this by computing the characteristic polynomial of the matrix (the polynomial whose roots are the eigenvalues), which amounts to a product of the characteristic polynomials of the blocks. Within each block we obtain a closed form formula for the characteristic polynomial and show that its root cannot exceed a bound that is determined by the cardinality of the respective equivalence class (properly normalized).

Combining the two lemmas by multiplying the rank bound with the eigenvalue absolute value bound implies the theorem.

Paper Organization. We use standard quantum and cryptographic notations and definitions, essentially following [6], see short summary in Section 2. Our construction is presented in Section 3 and proven in Section 4.

2 Preliminaries

For $m \in \mathbb{N}$, we denote $[m] := \{1, \cdots, m\}$. For a natural number N, denote by $\omega_N := e^{\frac{2\pi i}{N}}$ the complex root of unity of order N. Also for N, denote by $\mathcal{S}(N)$ the set of unit vectors in $\mathbb{C}^N$, by $\mathcal{D}(N)$ the set of $N \times N$ density matrices over $\mathbb{C}$, and by $\mathcal{U}(N)$ the set of $N \times N$ unitary matrices over $\mathbb{C}$. Note that for $n \in \mathbb{N}$, $\mathcal{S}(2^n)$ is the set of n-qubit pure quantum states, $\mathcal{D}(2^n)$ is the set of n-qubit mixed states, and $\mathcal{U}(2^n)$ is the set of n-qubit unitaries. When we consider quantum algorithms, we usually think of them as a uniform family of quantum circuits.

When we consider eigenvalues and singular values of matrices throughout this paper, we implicitly refer to eigenvalues and singular values that possibly repeat, e.g. $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ for matrix with n, possibly identical eigenvalues.

The trace distance, defined below, is a generalization of statistical distance to the quantum setting and represents the maximal distinguishing probability between quantum states.

Definition 2.1 (Trace distance). *Let $\rho_1, \rho_2 \in \mathcal{D}(2^n)$ be two density matrices of n-qubit mixed states. The trace distance between them is*

$$\mathrm{TD}(\rho_1, \rho_2) := \frac{1}{2} \left\| \rho_1 - \rho_2 \right\|_1 \ ,$$

where for a hermitian matrix M, $\|M\|_1 = \sum_i |\lambda_i|$, where λ_i are the eigenvalues of M.

The following is a basic fact that shows that classical circuits are a subset of quantum circuits. Recall that the Toffoli gate implements the 3-qubit unitary defined by $|x, y, z\rangle \to |x, y, z \oplus xy\rangle$.

Proposition 2.2 (Toffoli gate is universal for classical computation). *Let $f : \{0,1\}^n \to \{0,1\}^m$ be a function and let C be a classical circuit that computes f. Define the unitary $U_f : |x\rangle|y\rangle \to |x\rangle|y \oplus f(x)\rangle$. Then there exists a quantum circuit of size $O(|C|)$ consisting only of Toffoli gates that computes U_f (possibly using auxiliary $|0\rangle$ qubits).*

HT circuits are quantum circuits of a restricted structure, defined as follows.

Definition 2.3 (HT Circuit). *A quantum circuit C is an* **HT** *circuit if the first layer of the circuit consists of only Hadamard gates on a subset of the qubits, and the rest of the circuit consists of only Toffoli gates.*

2.1 Pseudorandom Functions and k-Wise Independent Functions

Here we define pseudorandom functions with quantum security (QPRFs).

Definition 2.4 (Quantum-Secure Pseudorandom Function (QPRF)). *Let $\mathcal{K} = \{\mathcal{K}_n\}_{n\in\mathbb{N}}$ be an efficiently samplable key distribution, and let $\mathsf{PRF} = \{\mathsf{PRF}_n\}_{n\in\mathbb{N}}$, $\mathsf{PRF}_n : \mathcal{K}_n \times \{0,1\}^n \to \{0,1\}^n$ be an efficiently computable function. We say that PRF is a quantum-secure pseudorandom function if for every efficient non-uniform quantum algorithm A that can make quantum queries there exists a negligible function $\mathrm{negl}(\cdot)$ s.t. for every $n \in \mathbb{N}$,*

$$\left| \Pr_{k\leftarrow\mathcal{K}_n}[A^{\mathsf{PRF}(k,\cdot)}() = 1] - \Pr_{f\leftarrow(\{0,1\}^n)^{(\{0,1\}^n)}}[A^f() = 1] \right| \leq \mathrm{negl}(n) \ .$$

In [15], QPRFs were proved to exist under the assumption that post-quantum one-way functions exist.

We define k-wise independent functions are keyed functions s.t. when the key is sampled uniformly at random, then any k different inputs to the function generate k-wise independent random variables.

Definition 2.5 ($k(n)$-Wise Independent Function). *Let $k(n) : \mathbb{N} \to \mathbb{N}$ be a function, $\mathcal{K} = \{\mathcal{K}_n\}_{n\in\mathbb{N}}$ be a key distribution, and let $f = \{f_n\}_{n\in\mathbb{N}}$, $f_n : \mathcal{K}_n \times \{0,1\}^n \to \{0,1\}^n$ be a function. Thus, f is a $k(n)$-wise independent function if for all n, for every distinct $k(n)$ input values $x_1, \cdots, x_{k(n)} \in \{0,1\}^n$,*

$$\forall y_1, \cdots, y_{k(n)} \in \{0,1\}^n : \Pr_{s\leftarrow\mathcal{K}_n}[f(s, x_1) = y_1 \wedge \cdots \wedge f(s, x_{k(n)}) = y_{k(n)}] = 2^{-n\cdot k(n)} \ .$$

It is not a part of the standard definition, but it is usually the case that we consider $\mathcal{K}$ to be efficiently samplable and f to be efficiently computable.

2.2 Quantum Randomness and Pseudorandomness

The Haar Measure on Quantum States. Intuitively, the Haar measure on quantum states is the quantum analogue of the classical uniform distribution over bit strings, that is, we can think of it as the uniform (continuous) probability distribution on quantum states. Recall that an n-qubit quantum state can be viewed as a unit vector in $\mathbb{C}^{2^n}$, thus the Haar measure on n qubits is the uniform distribution over all unit vectors in $\mathbb{C}^{2^n}$.

Formally, the density matrix representing the distribution of drawing a random Haar vector and outputting t copies of it is given below.

Definition 2.6 (n-Qubits, t-Copy Random Haar State). *Let $t, n \in \mathbb{N}$, we define the n-qubits t-copy random Haar mixed state to be*

$$\rho_{(t,n,H)} := \mathbb{E}_{|\psi\rangle \leftarrow \mu(2^n)} \left[(|\psi\rangle\langle\psi|)^{\otimes t} \right] \ ,$$

where $\mu(2^n)$ is the continuous distribution over $\mathbb{C}^{2^n}$ that is invariant under unitary transformations (it is known that there is only one such distribution).

Approximate Quantum State t-Designs. Approximate t-designs are quantum distributions that are approximately random when the number of output copies of the sampled state is restricted. The formal definition follows.

Definition 2.7 (n-Qubits, ε-Approximate State t-Design). *Let $\varepsilon \in [0,1], t \in \mathbb{N}$, and let $\mathcal{Q}$ be a quantum distribution over n-qubit states. We say that $\mathcal{Q}$ is an ε-approximate state t-design if*

$$\mathrm{TD}\big(\mathbb{E}_{|\psi\rangle \leftarrow \mathcal{Q}}[(|\psi\rangle\langle\psi|)^{\otimes t}],\ \rho_{(t,n,H)}\big) \leq \varepsilon \ .$$

For the sake of completeness, we give a definition for quantum state t-design generators.

Definition 2.8 ($\varepsilon(n)$-Approximate State $t(n)$-Design Generator). *Let $\varepsilon(n) : \mathbb{N} \to [0,1]$, $t(n) : \mathbb{N} \to \mathbb{N}$ be functions. We say that a pair of quantum algorithms (K, G) is an $\varepsilon(n)$-approximate state $t(n)$-design generator if the following holds:*

- ***Key Generation.*** *For all n, $K(1^n)$ always outputs a classical key k.*
- ***State Generation.*** *For all n and for all k in the image of $K(1^n)$, there exists an n-qubit pure state $|\phi_k\rangle$ s.t. $G(1^n, k) = |\phi_k\rangle$.*
- ***Approximate Quantum Randomness.*** *For all n, the distribution $|\phi_k\rangle_{k \leftarrow K(1^n)}$ is an n-qubit, $\varepsilon(n)$-approximate state $t(n)$-design.*

Note that we define the generator as two algorithms instead of one, to highlight the fact that a state that is sampled can be generated multiple times on demand.

For the purposes of this work it is convenient to define the notion of Asymptotically Random States (ARS) as follows.

Definition 2.9 (Asymptotically Random State (ARS)). *An Asymptotically Random State (*ARS*) is shorthand for an asymptotic sequence of* $\mathrm{negl}(n)$*-approximate $n^{\omega(1)}$-designs.*

Quantum Pseudorandomness. The notion of pseudorandom quantum states was introduced in [6], was shown to be implied by QPRFs, and is defined below.

Definition 2.10 (Pseudorandom Quantum State (PRS)). *A pair of quantum polynomial-time algorithms (K, G) is a Pseudorandom State Generator (PRS Generator) if the following holds:*

- ***Key Generation.*** *For all n, $K(1^n)$ always outputs a classical key k.*
- ***State Generation.*** *For all n and for all k in the image of $K(1^n)$, there exists an n-qubit pure state $|\phi_k\rangle$ s.t. $G(1^n, k) = |\phi_k\rangle$.*
- ***Security.*** *For any polynomial $t(\cdot)$ and a non-uniform efficient quantum algorithm A there exists a negligible function $\mathrm{negl}(\cdot)$ such that for all $n \in \mathbb{N}$,*

$$\left| \Pr_{k \leftarrow K(1^n)}[A(|\phi_k\rangle^{\otimes t(n)}) = 1] - \Pr_{|\psi\rangle \leftarrow \mu}[A(|\psi\rangle^{\otimes t(n)}) = 1] \right| \leq \mathrm{negl}(n) \ ,$$

where μ is the Haar measure on $\mathcal{S}(2^n)$.

If the above holds, we say that the ensemble $\mathsf{PRS} = \{\mathsf{PRS}_n\}_{n\in\mathbb{N}}$, where PRS_n is the distribution $|\phi_k\rangle_{k\leftarrow K(1^n)}$, is a Pseudorandom Quantum State (PRS) which is generated by (K, G).

In the above definition, the number of qubits in the pseudorandom states can also be parameterized (i.e. $G(1^n, k)$ can output $m(n)$-qubit states and not necessarily n-qubit states), but in the current work we will ignore this.

3 Construction

The following construction will be the base of both our pseudorandom state and quantum state t-design constructions.

Definition 3.1 (Binary Phase State Generator for F). *Let $\mathcal{K} = \{\mathcal{K}_n\}_{n\in\mathbb{N}}$ be a key space and let $F = \{F_n\}_{n\in\mathbb{N}}$ be a keyed (boolean) function $F_n : \mathcal{K}_n \times \{0,1\}^n \to \{0,1\}$. $\mathsf{G}^F_{\mathrm{bin}}$ is the procedure that takes as input a $k \in \mathcal{K}_n$ and outputs the superposition*

$$|\phi_k\rangle := 2^{-n/2} \sum_{x\in\{0,1\}^n} (-1)^{F_k(x)}|x\rangle \ .$$

The following claim establishes that $\mathsf{G}^F_{\mathrm{bin}}$ is efficiently implementable when F is.

Claim. If F is computable by a classical circuit of size $s(n)$ and depth $d(n)$, then $\mathsf{G}^F_{\mathrm{bin}}$ is computable by an **HT** circuit of size $O(s(n))$ and depth $d(n) + 1$.

Proof. The algorithm of $\mathsf{G}^F_{\mathrm{bin}}$ will get as input a key k and generate the state $|+\rangle^{\otimes n}|-\rangle$ by performing $(n + 1)$ Hadamard gates (in parallel) $H^{\otimes(n+1)}$ on the

ancillary classical state $|0\rangle^{\otimes n}|1\rangle$, then execute the F_k circuit (which can be realized quantumly by Toffoli gates) on the state $|+\rangle^{\otimes n}|-\rangle$. After the execution of F_k, the state is

$$2^{-n/2}\sum_{x\in\{0,1\}^n}(-1)^{F_k(x)}|x\rangle|-\rangle,$$

thus by tracing out the last qubit we get the output state $|\phi_k\rangle$.

We note that previous candidates required a more involved generation process which required applying quantum Fourier transform modulo 2^n, or a similar procedure.

3.1 Our Pseudorandom Quantum State (PRS) Generator and its Properties

Recall the definition of a PRS (see Definition 2.10) and of a QPRF (Definition 2.4). We present our construction of a PRS candidate with binary phase as follows.

Claim. If F is a QPRF then $\mathsf{G}^F_{\text{bin}}$ (along with the key generation algorithm of F) is a secure PRS generator.

Proof. First, it's clear that the key generation algorithm K of our PRS is the key generation algorithm of F (that for input 1^n, samples $k \leftarrow \mathcal{K}_n$), and that the state generation algorithm G of our PRS is $\mathsf{G}^F_{\text{bin}}$.

Now, we argue that by the quantum-security of F, for any polynomial number of copies $t(n)$, the distribution $|\phi_k\rangle_{k\leftarrow\mathcal{K}_n}$ is computationally indistinguishable (by quantum adversaries) from a random binary phase state, that is, the distribution over n-qubit quantum states defined by

$$2^{-n/2}\sum_{x\in\{0,1\}^n}(-1)^{f(x)}|x\rangle \ ,$$

where $f:\{0,1\}^n\to\{0,1\}$ is a truly random function.

By Theorem 1.2, a random binary phase state is an ARS (Definition 2.9), which in particular means that a random Haar state and a random binary phase state are computationally indistinguishable for any polynomial number of copies. By the the triangle inequality of computational indistinguishability, we deduce that for any polynomial number of copies, the quantum distribution $|\phi_k\rangle_{k\leftarrow\mathcal{K}_n}$ and the Haar distribution are computationally indistinguishable, which completes our proof.

Remark 3.2. We note that in our security proof we did not use the full power of quantumly secure PRFs. Indeed, if we consider the QPRF unitary U_{PRF_k} : $|x\rangle|y\rangle \to |x\rangle|y\oplus\mathsf{PRF}_k(x)\rangle$, then in order for the PRS to be secure, it is only needed that the QPRF will be secure when the input register is in the uniform superposition $|+\rangle^{\otimes n}$ (and moreover, the output register is $|-\rangle$). In particular, we don't even need the QPRF to be secure against chosen classical queries. This

can be thought of as a quantum analog of the classical notion of weak PRFs [11]. In the classical setting, it is conjectured that weak PRFs reside in a lower complexity class than full fledged PRFs [1]. If similar behavior can be shown in the quantum case it could improve the efficiency of PRS constructions.

We leave the investigation of this new notion (which we propose to call quantumly weak PRFs) to future works.

We conclude with observing that by our result, the complexity of PRSs is no greater than that of QPRFs, and is moreover implementable by **HT** circuits.

Corollary 3.3. *Let* $\mathsf{PRF} = \{\mathsf{PRF}_n\}_{n\in\mathbb{N}}$ *be a* QPRF*. Thus there is a* PRS *generator construction* (K, G) *implemented by* **HT** *circuits, where* K *is implemented by circuits of the same size and depth as that of the key sampling algorithm of* PRF*, and* G *is implemented by circuits of the same size and depth (up to asymptotics) as that of* PRF*.*

3.2 Shallow-Circuit Approximate t-Design Generators

We note that by a simple observation, we can replace the truly random function f in Theorem 1.2 with a $2t$-wise independent function to gain an elementary and efficient construction of quantum state approximate t-designs. Formally, we use the following fact.

Fact 3.4 ([15], Fact 2) *The behavior of any quantum algorithm making at most* q *queries to a* $2q$*-wise independent function is identical to its behavior when the queries are made to a random function.*

This implies that when f is a $2t$-wise independent function, then the state from Theorem 1.2 is a $\frac{4t^2}{2^n}$-approximate t-design. We note that this observation can also be applied to the ARS from [6], and it would imply a different (but seemingly less efficient) construction of t-designs.

Corollary 3.5. *The distribution over* n*-qubit quantum states defined by*

$$2^{-n/2}\sum_{x\in\{0,1\}^n}(-1)^{f(x)}|x\rangle$$

where $f:\{0,1\}^n\to\{0,1\}$ *is a* $2t$*-wise independent function, is a* $\frac{4t^2}{2^n}$*-approximate* t*-design.*

More explicitly, combining the above with Claim 3 implies that that when f is a $2t$-wise independent function, $\mathsf{G}^f_{\mathrm{bin}}$ is an approximate t-design generator (along with the key generation algorithm of f). The following corollary relates the complexity of t-design generators with that of the $2t$-wise independent functions.

Corollary 3.6. *Let* $t(n):\mathbb{N}\to\mathbb{N}$ *be a function and let* $f=\{f_n\}_{n\in\mathbb{N}}$*,* $f_n:\mathcal{K}_n\times\{0,1\}^n\to\{0,1\}$ *be a* $(2t(n))$*-wise independent function. Thus there is an* $\frac{O(t(n)^2)}{2^n}$*-approximate quantum state* $t(n)$*-design generator* (K, G) *implemented*

by **HT** *circuits, where K is implemented by circuits of the same size and depth as that of the key sampling algorithm of f, and G is implemented by circuits of the same size and depth (up to asymptotics) as that of f.*

Finally, we can instantiate with known construction of k-wise independent functions to obtain the following.

Corollary 3.7. *For every function $t(n) : \mathbb{N} \to \mathbb{N}$, there exists a $\frac{O(t(n)^2)}{2^n}$-approximate quantum state $t(n)$-design generator, implemented by* **HT** *circuits of* $\mathrm{poly}(t(n), n)$ *size and $O(\log t(n) \cdot \log n)$ depth.*

Proof. We recall the most elementary construction of k-wise independent distributions over 2^n variables. Consider the field $\mathbb{F} = \mathbb{F}_{2^n}$ and recall that $\mathbb{F}$ elements correspond to degree $(n-1)$ formal polynomials with binary coefficients. Thus there is a natural bijection between $\mathbb{F}$ and $\{0,1\}^n$ that allows to represent $\mathbb{F}$ elements as elements in $\{0,1\}^n$. This representation allows to perform field arithmetic operations using circuits of size $\mathrm{poly}(n)$ and depth $O(\log n)$.

A k-wise independent distribution over $\mathbb{F}^{\mathbb{F}}$ is defined by the evaluations of a random degree $(k-1)$ polynomial over $\mathbb{F}$, on all elements in $\mathbb{F}$. The computational complexity of evaluating such a polynomial is $\mathrm{poly}(k, n)$ and its depth is $O(\log k \cdot \log n)$. Plugging in $k = 2t$ completes the proof (note that we only require $2t$-wise independence over $\{0,1\}^{\mathbb{F}}$ so our instantiation is actually a slight overkill).

4 Proof of Theorem 1.2

We introduce the following notation.

Notation 4.1 (Complex phase state by f) *For a function $f : \{0,1\}^n \to [2^n]$ we denote*

$$|f\rangle_{(2^n)} := 2^{-n/2} \sum_{x \in \{0,1\}^n} \omega_{2^n}^{f(x)} |x\rangle \ .$$

when it is clear from the context, the subscript 2^n will be dropped from $|f\rangle_{(2^n)}$.

Notation 4.2 (Binary phase state by f) *For a function $f : \{0,1\}^n \to \{0,1\}$ we denote*

$$|f\rangle_{(2)} := 2^{-n/2} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle \ .$$

when it is clear from the context, the subscript 2 will be dropped from $|f\rangle_{(2)}$.

Notation 4.3 (t-copy random complex phase mixed state) *For $t, n \in \mathbb{N}$, denote*

$$\rho_{(t,n,2^n)} := \mathbb{E}_f[(|f\rangle_{(2^n)}\langle f|_{(2^n)})^{\otimes t}] \ ,$$

where the expectation is taken over a uniformly random function $f : \{0,1\}^n \to [2^n]$.

Notation 4.4 (t-copy random binary phase mixed state) *For $t, n \in \mathbb{N}$, denote*

$$\rho_{(t,n,2)} := \mathbb{E}_f[(|f\rangle_{(2)}\langle f|_{(2)})^{\otimes t}] \ ,$$

where the expectation is taken over a uniformly random function $f : \{0,1\}^n \to \{0,1\}$.

In [6] It is shown that the random complex phase state is an ARS.

Lemma 4.5 ([6], Lemma 2). *Let $n, t \in \mathbb{N}$, then*

$$\mathrm{TD}\left(\rho_{(t,n,2^n)}, \rho_{(t,n,H)}\right) = \prod_{i\in[t-1]}\left(1-\frac{i}{2^n}\right) - \prod_{i\in[t-1]}\left(1-\frac{2\cdot i}{2^n+i}\right) \ .$$

We will show that a random binary phase state is asymptotically statistically close to a random complex phase state. More precisely, we will prove the following lemma.

Lemma 4.6. *Let $n, t \in \mathbb{N}$, then*

$$\mathrm{TD}\left(\rho_{(t,n,2)}, \rho_{(t,n,2^n)}\right) \leq \prod_{i\in[t-1]}\left(1+\frac{i}{2^n}\right) - \prod_{i\in[t-1]}\left(1-\frac{i}{2^n}\right) \ .$$

Using the triangle inequality of trace distance and Lemmas 4.5 and 4.6 (below, in the first inequality), we show that a random binary phase state is an ARS. In the following, assume that $t < \frac{\sqrt{2^n}}{2}$, otherwise the upper bound on the trace distance trivially holds:

$$\begin{aligned}
&\mathrm{TD}\left(\rho_{(t,n,2)}, \rho_{(t,n,H)}\right) \leq \prod_{i\in[t-1]}\left(1+\frac{i}{2^n}\right) - \prod_{i\in[t-1]}\left(1-\frac{2\cdot i}{2^n+i}\right) \leq \\
&\left(1+\frac{t}{2^n}\right)^t - \left(1-\frac{2\cdot t}{2^n+t}\right)^t \leq \left(1+\frac{t}{2^n}\right)^t - \left(1-\frac{2\cdot t}{2^n}\right)^t \underset{(*)}{\leq} \\
&1+\frac{2\cdot t^2}{2^n} - \left(1-\frac{2\cdot t}{2^n}\right)^t \underset{(**)}{\leq} 1+\frac{2\cdot t^2}{2^n} - \left(1-\frac{2\cdot t^2}{2^n}\right) = \frac{4\cdot t^2}{2^n} \ ,
\end{aligned}$$

where $(*)$ is due to one variant of Bernoulli's inequality $(\forall r > 1, x \in [0, \frac{1}{2(r-1)}) : (1+x)^r \leq 1+2rx)$, and $(**)$ follows from the more popular variant of Bernoulli's inequality $(\forall r \notin (0,1), x \geq -1 : (1+x)^r \geq 1+rx)$.

Therefore, all that remains is to prove Lemma 4.6, which will require most technical effort.

4.1 Proof of Lemma 4.6

Denote the difference matrix $\rho_n := \rho_{(t,n,2)} - \rho_{(t,n,2^n)}$. The proof of the lemma contains two main components. First, an upper bound on the number of non-zero eigenvalues of ρ_n.

Lemma 4.7. *Let $n \in \mathbb{N}$ and let $t \in \{1, 2, \cdots, 2^n - 1\}$, thus the number of non-zero eigenvalues of ρ_n is upper bounded by*

$$\binom{2^n + t - 1}{t} - \binom{2^n}{t} .$$

Second, a lower bound on the minimal (as in most negative) eigenvalue of ρ_n.

Lemma 4.8. *Let $n \in \mathbb{N}$ and let $t \in \{1, 2, \cdots, 2^n - 1\}$, thus for all eigenvalues λ of ρ_n we have $-\frac{t!}{2^{tn}} \leq \lambda$.*

Note that this will give an upper bound on the absolute values of all negative eigenvalues of ρ_n.

Given the last two lemmas, we can prove Lemma 4.6.

Proof. Let $n \in \mathbb{N}$ and let $t \in \{1, 2, \cdots, 2^n - 1\}$[5]. ρ_n is a difference between two density matrices, and because that trace is linear and density matrices have a trace of 1, the trace of ρ_n is 0. Also recall that the sum of eigenvalues of a matrix is equal to its trace, so, the positive and negative eigenvalues of ρ_n balance each other to 0, and thus, a bound on the sum of absolute values of all eigenvalues of ρ_n can be obtained by bounding the sum of the absolute values of its negative eigenvalues. Formally:

$$\mathrm{TD}\left(\rho_{(t,n,2)}, \rho_{(t,n,2^n)}\right) = \frac{1}{2} \left\|\rho_n\right\|_1 =$$

$$\frac{1}{2} \cdot \sum_{\lambda \text{ eigenvalue of } \rho_n} |\lambda| = \sum_{\lambda \text{ negative eigenvalue of } \rho_n} |\lambda| .$$

Using Lemma 4.7 and Lemma 4.8, we obtain an upper bound on the last sum, which yields the wanted inequality.

$$\sum_{\lambda \text{ negative eigenvalue of } \rho_n} |\lambda| \leq \left(\binom{2^n + t - 1}{t} - \binom{2^n}{t} \right) \frac{t!}{2^{tn}} =$$

$$\frac{(2^n + t - 1)!}{2^{tn}(2^n - 1)!} - \frac{(2^n)!}{2^{tn}(2^n - t)!} = \frac{(2^n)!\left(\prod_{i\in[t-1]}(2^n + i)\right)}{2^{tn}(2^n - 1)!} - \frac{(2^n)!\left(\prod_{i\in[t-1]}(2^n - i)\right)}{2^{tn}(2^n - 1)!} =$$

$$\frac{\prod_{i\in[t-1]}(2^n + i)}{2^{(t-1)n}} - \frac{\prod_{i\in[t-1]}(2^n - i)}{2^{(t-1)n}} = \prod_{i\in[t-1]} \frac{2^n + i}{2^n} - \prod_{i\in[t-1]} \frac{2^n - i}{2^n} =$$

$$\prod_{i\in[t-1]} \left(1 + \frac{i}{2^n}\right) - \prod_{i\in[t-1]} \left(1 - \frac{i}{2^n}\right) .$$

[5] For $t \geq 2^n$ the bound trivially holds: Note that for $t \geq 2^n$ the bound's expression is minimized for $t = 2^n$ and $n = 1$, which yields 1 as a trivial bound on any trace distance.

4.2 The Structure of the Matrix ρ_n

We identify the structure of ρ_n in order to prove Lemma 4.13, which will be used in both proofs of Lemmas 4.7, 4.8. We do this by first describing $\rho_{(t,n,2^n)}$ and $\rho_{(t,n,2)}$. More precisely, we will derive combinatorial expressions for $\rho_{(t,n,2^n)}$ and $\rho_{(t,n,2)}$, and as a consequence we'll have an expression for their difference ρ_n.

The Structure of $\rho_{(t,n,2^n)}$. We will start with a formula for the entries of $\rho_{(t,n,2^n)}$ (a similar analysis was done for this matrix in [6]); for convenience, the definition is restated:

$$\rho_{(t,n,2^n)} = \mathbb{E}_{f\leftarrow[2^n]^{\{0,1\}^n}}\left[(|f\rangle\langle f|)^{\otimes t}\right] = \mathbb{E}_{f\leftarrow[2^n]^{\{0,1\}^n}}\left[|f\rangle^{\otimes t}\langle f|^{\otimes t}\right] \ .$$

Observe that for a function $f : \{0,1\}^n \to [2^n]$,

$$|f\rangle^{\otimes t} = \left(2^{-n/2}\sum_{x\in\{0,1\}^n}\omega_{2^n}^{f(x)}|x\rangle\right)^{\otimes t} = 2^{-tn/2}\sum_{\mathbf{x}=(x_1,\cdots,x_t)\in\{0,1\}^{n\times t}}\omega_{2^n}^{(\sum_{i\in[t]}f(x_i))}|\mathbf{x}\rangle \ .$$

Now we can compute $\rho_{(t,n,2^n)}$:

$$\rho_{(t,n,2^n)} = \mathbb{E}_f\left[|f\rangle^{\otimes t}\langle f|^{\otimes t}\right] =$$

$$\mathbb{E}_f\left[\left(2^{-tn/2}\sum_{\mathbf{x}=(x_1,\cdots,x_t)\in\{0,1\}^{n\times t}}\omega_{2^n}^{(\sum_{i\in[t]}f(x_i))}|\mathbf{x}\rangle\right)\cdot\right.$$

$$\left.\left(2^{-tn/2}\sum_{\mathbf{y}=(y_1,\cdots,y_t)\in\{0,1\}^{n\times t}}\omega_{2^n}^{(-\sum_{i\in[t]}f(y_i))}\langle\mathbf{y}|\right)\right] =$$

$$2^{-tn}\sum_{\mathbf{x},\mathbf{y}\in\{0,1\}^{n\times t}}|\mathbf{x}\rangle\langle\mathbf{y}|\cdot\mathbb{E}_f\left[\omega_{2^n}^{(\sum_{i\in[t]}f(x_i)-\sum_{i\in[t]}f(y_i))}\right] \ ,$$

So, for $\mathbf{x},\mathbf{y}\in\{0,1\}^{n\times t}$, the $(\mathbf{x},\mathbf{y})$-th entry of $\rho_{(t,n,2^n)}$ is

$$2^{-tn}\cdot\mathbb{E}_f\left[\omega_{2^n}^{(\sum_{i\in[t]}f(x_i)-\sum_{i\in[t]}f(y_i))}\right] \ .$$

Now, define:

Definition 4.9 ((t,n) permutations). *Let $\mathbf{x},\mathbf{y}\in\{0,1\}^{t\times n}$, and denote $\mathbf{x} = (x_1,\cdots,x_t), \mathbf{y} = (y_1,\cdots,y_t)$, where $\forall i\in[t] : x_i, y_i \in \{0,1\}^n$. We say that $\mathbf{x},\mathbf{y}$, are (t,n) permutations of each other (or just permutations of each other) if there exists a permutation $\pi\in S_t$ s.t.*

$$(x_1,\cdots,x_t) = (y_{\pi(1)},\cdots,y_{\pi(t)}) \ .$$

Note that an equivalent convenient characterization of the two strings $\mathbf{x}, \mathbf{y}$ being permutations of each other is that the multisets $\{x_1, \cdots, x_t\}, \{y_1, \cdots, y_t\}$ are equal.

Observe that when $\mathbf{x}$ and $\mathbf{y}$ are permutations of each other, then for every f we have $\sum_{i\in[t]} f(x_i) = \sum_{i\in[t]} f(y_i)$ and thus the expected value is 1 and the entry's value is 2^{-tn}. We would like to also claim that if $\mathbf{x}, \mathbf{y}$ are not permutations of each other then the entry is 0, and it turns out we indeed can. Observe that if $\mathbf{x}, \mathbf{y}$ are not permutations of each other then there exists a string $s \in \{0,1\}^n$ that appears a different number of times in $\mathbf{x}$ and $\mathbf{y}$, and we can say that the $(\mathbf{x}, \mathbf{y})$-th entry is

$$2^{-tn} \cdot \mathbb{E}_f\left[\omega_{2^n}^{(\sum_{i\in[t]} f(x_i) - \sum_{i\in[t]} f(y_i))}\right] = 2^{-tn} \cdot \beta \cdot \mathbb{E}_f\left[\omega_{2^n}^{a\cdot f(s)}\right] \ ,$$

where $\beta \in \mathbb{R}$ is some real number (which we won't care about) and $a \in \{-t, \cdots, -1, 1, \cdots, t\}$ is the (non-zero) difference between the number of appearances of s in $\mathbf{x}$ and $\mathbf{y}$ (last equality follows from the fact that the expectation of a product of independent random variables is the product of expectations). Now we will use our restriction on t, which is that t is strictly smaller then 2^n. Combined with the fact that $a \neq 0$, it is necessarily the case that $\omega_{2^n}^a \neq 1$ (if t could be as big as 2^n then a will be able to be 2^n or some integer multiple of it, which will yield $\omega_{2^n}^a = 1$). After this restriction we obtain:

$$\mathbb{E}_f\left[\omega_{2^n}^{a\cdot f(s)}\right] = \sum_{i\in\{0,1,\cdots,2^n-1\}} 2^{-n} \cdot \omega_{2^n}^{a\cdot i} = 2^{-n} \cdot \left(\frac{\omega_{2^n}^{a\cdot 2^n} - 1}{\omega_{2^n}^a - 1}\right) = 0 \ ,$$

Finally, the above yields a combinatorial description of $\rho_{(t,n,2^n)}$:

$$\forall \mathbf{x}, \mathbf{y} \in \{0,1\}^{n\times t} : \rho_{(t,n,2^n)}[\mathbf{x}, \mathbf{y}] = \begin{cases} 2^{-tn} & \mathbf{x}, \mathbf{y} \text{ are permutations} \\ 0 & \mathbf{x}, \mathbf{y} \text{ are not permutations} \end{cases} .$$

The Structure of $\boldsymbol{\rho_{(t,n,2)}}$. By the same reasoning as in the case of $\rho_{(t,n,2^n)}$, we obtain that the $(\mathbf{x}, \mathbf{y})$-th entry of $\rho_{(t,n,2)}$ is

$$2^{-tn} \cdot \mathbb{E}_f\left[(-1)^{(\sum_{i\in[t]} f(x_i) - \sum_{i\in[t]} f(y_i))}\right] \ ,$$

where this time f is a random function from $\{0,1\}^n$ to $\{0,1\}$ (rather than from $\{0,1\}^n$ to $[2^n]$). Because $(-1) = (-1)^{-1}$, the entry is simplified to

$$2^{-tn} \cdot \mathbb{E}_f\left[(-1)^{(\sum_{i\in[t]} f(x_i) + \sum_{i\in[t]} f(y_i))}\right] \ .$$

Like in the case of $\rho_{(t,n,2^n)}$, we would like a nice and clean combinatorial predicate to describe the entries of the matrix, and as we'll see in a bit, the matrix $\rho_{(t,n,2)}$ indeed have the same general structure as $\rho_{(t,n,2^n)}$ but with different predicate on $\mathbf{x}, \mathbf{y}$.

First, define the following:

Definition 4.10 ((t,n) stabilizations). *Let $\mathbf{x},\mathbf{y} \in \{0,1\}^{t\times n}$, and denote $\mathbf{x} = (x_1,\cdots,x_t), \mathbf{y} = (y_1,\cdots,y_t)$, where $\forall i \in [t] : x_i, y_i \in \{0,1\}^n$. We say that $\mathbf{x},\mathbf{y}$, are (t,n) stabilizations of each other (or just stabilizations of each other) if in the concatenated string $(\mathbf{x}\ \mathbf{y}) = (x_1,\cdots,x_t,y_1,\cdots,y_t)$, for every $s \in \{0,1\}^n$, s appears an even number of times (this, of course, includes appearing 0 times).*

We note that the stabilization relation (which is all pairs that stabilize each other) is an equivalence relation over the set $\{0,1\}^{n\times t}$ (just like the permutation relation, which we didn't mention it being an equivalence relation, but it can easily be seen as one). It is clear that the stabilization relation is reflexive ($\mathbf{x}$ is always stabilizing $\mathbf{x}$), and it is also easy to verify that it is symmetric. To see why it is also transitive, we will use an additional characterization:

Definition 4.11. *For a string $\mathbf{z} = (z_1,\cdots,z_t) \in \{0,1\}^{n\times t}$ with $\forall i \in [t] : z_i \in \{0,1\}^n$, $\mathrm{Odd}(\mathbf{z})$ is the set of strings from $\{0,1\}^n$ that appear an odd number of times in the sequence $(z_1,\cdots,z_t)$.*

For example, if $n = 3$ and $t = 7$ then $\mathrm{Odd}(101,111,101,000,011,111,111) = \{111,000,011\}$.

We claim that two strings $\mathbf{x},\mathbf{y}$ are stabilizations of each other if and only if $\mathrm{Odd}(\mathbf{x}) = \mathrm{Odd}(\mathbf{y})$. It is easy to verify the correctness of this claim, and also the fact that this claim implies the transitivity of the stabilization relation.

To identify the elements of $\rho_{(t,n,2)}$ it remains to observe that when $\mathbf{x},\mathbf{y}$ are stabilizations of each other then the entry is 2^{-tn}, and when they are not, then we have $\mathrm{Odd}(\mathbf{x}) \neq \mathrm{Odd}(\mathbf{y})$ and it can be verified that the entry is 0, which yields the following description of $\rho_{(t,n,2)}$:

$$\forall \mathbf{x},\mathbf{y} \in \{0,1\}^{n\times t} : \rho_{(t,n,2)}[\mathbf{x},\mathbf{y}] = \begin{cases} 2^{-tn} & \mathbf{x},\mathbf{y} \text{ are stabilizations} \\ 0 & \mathbf{x},\mathbf{y} \text{ are not stabilizations} \end{cases}.$$

Conclusion. Note that if $\mathbf{x},\mathbf{y}$ are permutations then they necessarily stabilize each other, but the opposite is not true generally, furthermore, it is fairly easy to find stabilizing pairs that are not permutations, for instance $(111,000,101,101,000)$ and $(110,111,111,111,110)$. We'll call a pair of strings that suffice this demand (i.e. stabilize each other but are not permutations) remotely stabilized, that is:

Definition 4.12 ((t,n) remote stabilizations). *Let $\mathbf{x},\mathbf{y} \in \{0,1\}^{t\times n}$, we say that $\mathbf{x},\mathbf{y}$ are (t,n) remote stabilizations of each other (or just remote stabilizations of each other) if they are stabilizations of each other but are not permutations of each other.*

In contrast to the cases of permutation and stabilization, remote stabilization is not an equivalence relation, and thus (generally speaking) it is harder to work with it. The stabilization relation is symmetric, but it is not reflexive, and in fact it is anti-reflexive, because a string is always a permutation of itself (and thus not a remote stabilization of itself), and it is also not transitive, because a (non-empty) relation which is symmetric and anti-reflexive can't be transitive.

As said above, two strings that are permutations of each other are necessarily stabilizations of each other (in other words, the permutation relation is a refinement of the stabilization relation), and we deduce that ρ_n has no negative terms and is also binary (scaled by 2^{-tn}). Finally, this proves the characterization lemma of ρ_n.

Lemma 4.13. *Let $n \in \mathbb{N}$ and let $t \in \{1, 2, \cdots, 2^n - 1\}$, then the entries of ρ_n can be given by the following formula:*

$$\forall \mathbf{x}, \mathbf{y} \in \{0,1\}^{n \times t} : \rho_n[\mathbf{x}, \mathbf{y}] = \begin{cases} 2^{-tn} & \mathbf{x}, \mathbf{y} \text{ are remote stabilizations} \\ 0 & \mathbf{x}, \mathbf{y} \text{ are not remote stabilizations} \end{cases} .$$

4.3 Proof of Lemma 4.7

Proof. We will give an upper bound on the number of non-zero eigenvalues of ρ_n. ρ_n is hermitian (and in particular diagonalizable) and thus the sum of dimensions of its eigenspaces sums up to the order of the matrix, which is 2^{tn}. Also recall that the 0-eigenspace of ρ_n is its kernel, thus by the rank-nullity theorem, the dimension of the 0-eigenspace plus the rank of ρ_n equals the order of the matrix, 2^{tn}. This means that the rank of ρ_n equals the sum of dimensions of non-zero eigenspaces of ρ_n, which is exactly the number of non-zero (possibly identical) eigenvalues of ρ_n, thus, by giving an upper bound of $\text{rank}(\rho_n)$, we get an upper bound on the number of non-zero eigenvalues of ρ_n.

It is a well known fact in linear algebra that elementary row operations does not change the rank of a matrix, it is also known that the rank of a matrix is bounded from above by the number of non-zero rows (the rank is the dimension of the row space, which in turn cannot be more than the number of non-zero rows), thus our bound on the rank of ρ_n will come from looking at ρ'_n, a row-equivalent matrix to ρ_n, and bounding its number of non-zero rows.

ρ'_n is obtained by the following procedure: Recall that the permutation relation and the stabilization relation are both equivalence relations on $\{0,1\}^{t \times n}$ and thus induce equivalence classes. It will be useful (also for the proof of the next lemma) to define the following:

Definition 4.14 (Sentinel of an Equivalence Class). *Let C be an equivalence class of one of the two equivalence relations above (either the permutation relation or the stabilization relation). We define $\mathbf{x}_C \in \{0,1\}^{tn}$ the sentinel of C to be the element in C with the largest lexicographic order (where the lexicographic order of strings is as usual, with the most significant bit on the left, and least significant bit on the right).*

Observation 1 *Let P be a permutation class of $\{0,1\}^{tn}$. Then, every pair in it $\mathbf{x}, \mathbf{y} \in P$ have the same set of remote stabilizers, and thus have identical rows in ρ_n.*

This means we can erase a bunch of redundant rows from ρ_n; for each permutation class P, take the sentinal row $\mathbf{x}_P$ of P and subtract it from all other

rows of strings from P. In the obtained matrix ρ'_n, the only non-zero rows are of sentinels.

The number of sentinels is exactly the number of equivalence classes of the permutation relation, which in turn is the number of different multisets of t elements from $\{0,1\}^n$ (note that a permutaion class can be defined by a multiset from $\{0,1\}^n$ of size t), and that number is known as common knowledge in combinatorics, usually referred to as "n multichoose k", in our case, $\binom{2^n+t-1}{t}$.

Observation 2 *Let P be a permutation class of a multiset of t distinct elements (essentially, a permutation class of a set of size t with elements from $\{0,1\}^n$), then each of its elements have no remote stabilizers.*

The above observation basically says that strings of t distinct elements are a special case where every stabilizer of them is also a permutation of them. This observation is useful to us because it means that for every permutation class P of t distinct elements, all rows of P are zero-rows in the original ρ_n (and thus so in ρ'_n).

Furthermore, the reason that observation 2 is important to our proof comes from the fact that there are $\binom{2^n}{t}$ such permutation classes, which is an overwhelming precentage from the total number of permutation classes $\binom{2^n+t-1}{t}$. To conclude, we said that in ρ'_n, only the sentinels can possibly have non-zero rows, and that there are $\binom{2^n+t-1}{t}$ sentinels in total, but now we add the information that out of these $\binom{2^n+t-1}{t}$ sentinels, $\binom{2^n}{t}$ have zero rows, and thus, there are at most

$$\binom{2^n+t-1}{t} - \binom{2^n}{t}$$

non-zero rows in ρ'_n (and as a side note, there are in fact more zero-rows, for instance, for permutation classes of a multisets of the same element appearing t times, but we won't care about these as their precentage is negligible). This concludes our proof of Lemma 4.7.

4.4 Proof of Lemma 4.8

We will give a lower bound on the most negative eigenvalue of ρ_n. Recall that ρ_n is hermitian and thus has only real eigenvalues. Let $\lambda \in \mathbb{R}$, we know that λ is an eigenvalue ρ_n if and only if $\det(\rho_n - \lambda I) = 0$. Denote by $\mathcal{A}$ the set of negative relative sizes of the permutation classes (along with 0),

$$\mathcal{A} := \left\{ -\frac{|P|}{2^{tn}} \;\middle|\; P \text{ is a permutation class} \right\} \cup \{0\} \ ,$$

where for a permutation class, its size is the number of different possible permutations of it, e.g. if P is a permutation class of a multiset of the same element t times, then $|P| = 1$, if P is a permutation class of a multiset of t distinct elements (in this specific case it is also a set) then $|P| = t!$, and if P is a permutation

class of a multiset of $(t-2)$ distinct elements plus an additional distinct element that appears twice, then $|P| = \binom{t}{2} \cdot (t-2)!$. We will show that there are no eigenvalues of ρ_n smaller then all elements of $\mathcal{A}$, which will give us a lower bound of $(-\frac{t!}{2^{tn}})$ on the minimal eigenvalue of ρ_n (as this is the minimal element in $\mathcal{A}$).

We prove the eigenvalue lower bound by calculating the determinant $\det(\rho_n - \lambda I)$ for $\lambda \in \mathbb{R} \setminus \mathcal{A}$, and showing that it cannot be 0. Recall that the permutation relation is a refinement of the stabilization relation, which means that every stabilization class can be divided into a bunch of permutation classes, also recall that in the proof of Lemma 4.7 we saw that for some stabilization classes, they are exactly a single permutation class and not a few (for example, according to the second observation in the proof, this is the case for permutation classes of sets of size t) - we'll call such stabilization classes *trivial stabilization classes*.

By the observations from Lemma 4.7 and by the restriction $\lambda \notin \mathcal{A}$, the calculation of the determinant is enabled and we obtain that for every $\lambda \in (\mathbb{R} \setminus \mathcal{A})$, the value of the determinant $\det(\rho_n - \lambda \cdot I)$ is,

$$\prod_{S \text{ trivial stabilization class}} (-\lambda)^{|S|} \cdot \prod_{S \text{ non-trivial stabilization class}} \left(\alpha_S \cdot \beta_S \cdot \gamma_S\right) ,$$

where,

$$\alpha_S = \prod_{\mathbf{x} \text{ non-sentinel in } S} (-\lambda) ,$$

$$\beta_S = \prod_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(-\lambda - \frac{|P|}{2^{tn}}\right) ,$$

$$\gamma_S = -\lambda + \left(\lambda + \frac{|P_S|}{2^{tn}}\right) \cdot \sum_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(\frac{\frac{|P|}{2^{tn}}}{(\lambda + \frac{|P|}{2^{tn}})}\right) .$$

The full version of this calculation is in [3].

Using The Determinant to Show The Lower Bound. Given the above determinant result, we can now finally prove Lemma 4.8.

Proof. Assume towards contradiction that there is a real number $\lambda' < -\frac{t!}{2^{tn}}$ (note that this implies $\lambda' \notin \mathcal{A}$) such that it is an eigenvalue of ρ_n, thus $\det(\rho_n - \lambda' \cdot I) = 0$ and thus it is necessarily the case that one of the terms in the above product (of the determinant) has to be 0. Due to $\lambda' \notin \mathcal{A}$, it can be seen that the only terms that can possibly be 0 in the above product are the terms γ_S for non-trivial S, so let's check what happens in these terms.

Let S be a non-trivial stabilization class, and consider the term γ_S in the product above:

$$-\lambda' + \left(\lambda' + \frac{|P_S|}{2^{tn}}\right) \cdot \sum_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(\frac{\frac{|P|}{2^{tn}}}{(\lambda' + \frac{|P|}{2^{tn}})}\right) .$$

We have,

$$\forall P \text{ permutation class in } S, \text{ including } P_S : \lambda' < -\frac{|P|}{2^{tn}} \ ,$$

and thus

$$\left(\lambda' + \frac{|P_S|}{2^{tn}}\right) < 0, \sum_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(\frac{\frac{|P|}{2^{tn}}}{(\lambda' + \frac{|P|}{2^{tn}})}\right) < 0 \ ,$$

which implies

$$\left(\lambda' + \frac{|P_S|}{2^{tn}}\right) \cdot \sum_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(\frac{\frac{|P|}{2^{tn}}}{(\lambda' + \frac{|P|}{2^{tn}})}\right) > 0 \ .$$

Finally, due to $-\lambda'$ being in particular positive, the term has to be positive as well:

$$-\lambda' + \left(\lambda' + \frac{|P_S|}{2^{tn}}\right) \cdot \sum_{P \text{ permutation class in } S \text{ with } P \neq P_S} \left(\frac{\frac{|P|}{2^{tn}}}{(\lambda' + \frac{|P|}{2^{tn}})}\right) > 0 \ ,$$

in contradiction to $\det(\rho_n - \lambda' I) = 0$.

Acknowledgments We thank Henry Yuen and Vinod Vaikuntanathan for insightful discussions. In particular thanks to Henry for pointing us to the [6] result. We thank the anonymous reviewers for their useful comments. We also thank Aram Harrow for providing advice regarding the state of the art.

References

1. A. Akavia, A. Bogdanov, S. Guo, A. Kamath, and A. Rosen. Candidate weak pseudorandom functions in ac0 mod2. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 21, page 33, 2014.
2. A. Ambainis and J. 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.
3. Z. Brakerski and O. Shmueli. (pseudo) random quantum states with binary phase. *CoRR*, abs/1906.10611, 2019.
4. C. Dankert, R. Cleve, J. Emerson, and E. Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.
5. A. W. Harrow and R. A. Low. Random quantum circuits are approximate 2-designs. *Communications in Mathematical Physics*, 291(1):257–302, 2009.
6. Z. Ji, Y. Liu, and F. Song. Pseudorandom quantum states. In H. Shacham and A. Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018.

7. R. Kueng and D. Gross. Qubit stabilizer states are complex projective 3-designs. *arXiv preprint arXiv:1510.02767*, 2015.
8. S. Lloyd. Capacity of the noisy quantum channel. *Physical Review A*, 55(3):1613, 1997.
9. Y. Nakata, M. Koashi, and M. Murao. Generating a state t-design by diagonal quantum circuits. *New Journal of Physics*, 16(5):053043, 2014.
10. Y. Nakata and M. Murao. Diagonal-unitary 2-design and their implementations by quantum circuits. *International Journal of Quantum Information*, 11(07):1350062, 2013.
11. M. Naor and O. Reingold. Synthesizers and their application to the parallel construction of pseudo-random functions. *Journal of Computer and System Sciences*, 58(2):336–375, 1999.
12. M. Nest. Classical simulation of quantum computation, the gottesman-knill theorem, and slightly beyond. *arXiv preprint arXiv:0811.0898*, 2008.
13. S. Popescu, A. J. Short, and A. Winter. Entanglement and the foundations of statistical mechanics. *Nature Physics*, 2(11):754, 2006.
14. J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves. Symmetric informationally complete quantum measurements. *Journal of Mathematical Physics*, 45(6):2171–2180, 2004.
15. M. Zhandry. How to construct quantum random functions. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 679–687. IEEE, 2012.