Efficient Unitary Designs from Random Sums and Permutations

Chi-Fang Chen
Institute for Quantum Information and Matter
California Institute of Technology
Pasadena, CA, USA
achifchen@gmail.com

Jordan Docter
Department of Computer Science
Stanford University
Stanford, CA, USA
jdocter@stanford.edu

Michelle Xu
Stanford Institute for Theoretical Physics
Stanford University
Stanford, CA, USA
mdx@stanford.edu

Adam Bouland
Department of Computer Science
Stanford University
Stanford, CA, USA
abouland@stanford.edu

Fernando G.S.L. Brandão
Institute for Quantum Information and Matter
California Institute of Technology
Pasadena, CA, USA
fbrandao@caltech.edu

Patrick Hayden
Stanford Institute for Theoretical Physics
Stanford University
Stanford, CA, USA
phayden@stanford.edu

2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS) | 979-8-3315-1674-1/24/$31.00 ©2024 IEEE | DOI: 10.1109/FOCS61266.2024.00037

***Abstract*—A unitary k-design is an ensemble of unitaries that matches the first k moments of the Haar measure. In this work, we provide two efficient constructions of k-designs on n-qubits using new random matrix theory techniques.**

Our first construction is based on exponentiating sums of random i.i.d. Hermitian matrices and uses $\tilde{O}(k^2n^2)$-many gates. In the spirit of central limit theorems, we show that this random sum approximates the Gaussian Unitary Ensemble (GUE). We then show that the product of just two exponentiated GUE matrices is already approximately Haar random. Our second construction is based on products of exponentiated sums of random permutations and uses $\tilde{O}(k\,\mathrm{poly}(n))$ many gates. The k dependence is optimal (up to polylogarithmic factors) and is inherited from the efficiency of existing k-wise independent permutations. Furthermore, replacing random permutations with quantum-secure pseudorandom permutations (PRPs), we also obtain a pseudorandom unitary (PRU) ensemble that is secure under nonadaptive queries.

A central feature of both proofs is a new connection between the polynomial method in quantum query complexity and the large-dimension (N) expansion in random matrix theory. In particular, the first construction uses the polynomial method to control high moments of certain random matrix ensembles without requiring delicate Weingarten calculations. In doing so, we define and solve a moment problem on the unit circle, asking whether a finite number of equally weighted points can reproduce a given set of moments. In our second construction, the key step is to exhibit an orthonormal basis for irreducible representations of the partition algebra that has a low-degree large-N expansion. This allows us to show that the distinguishing probability is a low-degree rational polynomial of the dimension N.

***Index Terms*—unitary designs, polynomial method, pseudorandomness, random matrices**

I. Introduction

Pseudorandom states and unitaries are fundamental tools in quantum information. By efficiently creating ensembles of states and unitaries that mimic the Haar measure, we can access Haar random properties without paying a cost exponential in the number of qubits. Broadly speaking, two types of quantum pseudorandomness have been previously considered. The first is *information-theoretic* pseudorandomness, where indistinguishability is determined by the statistical closeness of an ensemble. Such ensembles that match the first k moments of the Haar measure are called unitary or state k-designs, and are the natural quantum analogs of k-wise independent functions or permutations. These yield information-theoretic k-copy security and suffice for many applications, such as randomized benchmarking [EAŻ05], [KLR$^+$08], [DCEL09], cryptography [DLT02], [ABW09], shadow tomography [HKP20], communication [HHWY08], [SDTR13], and phase retrieval [KL17]. The second is *computational* pseudorandomness, where closeness is determined by the computational hardness of the distinguishing task. Pseudorandom states and unitaries (PRSs and PRUs) are ensembles that are computationally indistinguishable from Haar [JLS18]. These are the quantum analogs of pseudorandom functions (PRFs) or permutations (PRPs), and have found significant applications both in quantum cryptography (e.g., [Kre21], [KQST23]) and in the complexity of physical systems, e.g., [BFV19], [KTP20], [ABF$^+$24], [YE23]. This relaxed security notion allows one to obtain properties impossible in information-theoretic settings, such as small ensemble (key) sizes, many copy security, and low entanglement [JLS18], [ABF$^+$24].

Many questions remain open along both branches of the quantum pseudorandomness family tree. On the k-design side, a significant open problem has been to efficiently generate unitary k-designs on n qubits. It is known that this task requires at least $\Omega(nk)$ quantum gates [BHH16] but so far existing constructions have not achieved the bound, despite much work in the area [Web15], [ZKGG16], [Zhu17], [MY23], [HL09b], [HL09a], [BHH16], [HM23], [Haf22], [HJ19], [DCEL09], [CLLW16],

A.B. was supported in part by the DOE QuantISED grant DE-SC0020360, the AFOSR under grant FA9550-21-1-0392, and by the U.S. DOE Office of Science under Award Number DE-SC0020377. P.H. acknowledges support from AFOSR (award FA9550-19-1-0369), DOE (Q-NEXT), CIFAR, and the Simons Foundation. M.X. acknowledges support from NSF GRFP.

2575-8454/24/$31.00 ©2024 IEEE
DOI 10.1109/FOCS61266.2024.00037

[OBK+17], [NHKW17], [BNZZ19], [NZO+21], [BNOZ22], [OSP23], [KKS+23]. The first efficient and systematic constructions of approximate unitary k-designs which work for large values of k were given by [BHH16], who showed that random local circuits achieved the goal in $\mathcal{O}(k^{10.5}n^2)$ time. Then, [Haf22] improved the analysis of this construction, yielding $\tilde{\mathcal{O}}(k^5n^2)$ time. This stands in contrast to constructions of related k-wise independent objects, which achieve the optimal linear scaling in k, such as quantum state designs (e.g., [BS19]) as well as for analogous classical objects, namely k-wise independent functions [Jof74], [WC81], [ABI86] or approximate k-wise independent permutations [Kas07], [CK23]. We note that linear scaling[1] in k has been achieved[2] only in restricted cases, such as the limit of large local dimension [HHJ21], or if the number of moments matched is small ($k \ll \mathcal{O}(\sqrt{n})$) [NHKW17]. This problem is not only a fundamental and natural question in theoretical computer science, but has also gathered the attention of the quantum gravity community, because the linear growth in circuit complexity associated with an efficient k-design ensemble (see e.g., [RY17]) is believed to play a role in resolving certain paradoxes in the structure of quantum spacetime [BS18].

On the *computational* pseudorandomness side, a significant open problem is to construct pseudorandom unitaries (PRUs) from standard cryptographic assumptions, such as the existence of quantum-secure one-way functions (OWFs). While several efficient constructions of pseudorandom state ensembles are known [JLS18], [BS19], [ABF+24], [GTB23], [JMW23], [Ma23], progress has been much more difficult in the unitary case. [JLS18] proposed a few candidates, but security has yet to be proven. However, a number of objects intermediate between a PRS and a PRU have been rigorously constructed. For example, Ananth et al. constructed a pseudorandom function-like state, a generalization of a PRS that allows one to create polynomially many independent PRSs [AQY22], [AGQY22]. Subsequently, Lu et al. defined and constructed pseudorandom state scramblers, which are ensembles of unitaries that produce a PRS from any *fixed* input state [LQS+23]. There has also been a recent construction of parallel-secure pseudorandom isometries between spaces of differing dimensions [AGKL23], as well as another variant of a pseudorandom state scrambler for real (orthogonal) states [BM24]. However, the existence of efficient pseudorandom unitaries with general query security remains open.

In this work, we make progress on both quantum pseudorandomness open problems simultaneously. First, we construct a unitary k-design by summing "simple" random matrices. Let $\tilde{\mathcal{O}}(\cdot)$ absorb polylogarithmic dependence on n, k, ϵ^{-1}, and consider approximate k-designs in the parallel query model

$$\mathcal{N}_{2k,\boldsymbol{U}}[\boldsymbol{\rho}] := \mathbb{E}_{\boldsymbol{U}}[\boldsymbol{U}^{\otimes k}\boldsymbol{\rho}\boldsymbol{U}^{\dagger\otimes k}]$$

with the diamond norm

$$\|\mathcal{N}_1 - \mathcal{N}_2\|_\diamond := \|(\mathcal{N}_1 - \mathcal{N}_2) \otimes \mathcal{I}\|_{1-1}.$$

Theorem I.1 (Random sums)**.** *For any $k \leq 2^{O(n/\log n)}$, there exists an efficient quantum algorithm that generates an ϵ-approximate unitary k-design (in diamond norm) using $\tilde{O}\left(k^2n^2\log(\epsilon^{-1})\right)$ quantum gates, and $\tilde{O}(kn^2)$ bits of classical randomness.*

Second, we give a construction of a unitary k-design using quantum gates scaling near-linearly with k, which matches the lower bound on k-dependence up to logarithmic factors. Here, the improved efficiency is inherited from existing classical k-wise independent permutations [Kas07], [CK23][3].

Theorem I.2 (Random permutations)**.** *For any $k \leq 2^{O(n)}$, using existing efficient classical k-wise independent permutations, there exists an efficient quantum algorithm that generates an ϵ-approximate unitary k-design (in diamond norm) using $\tilde{\mathcal{O}}(k\,\mathrm{poly}(n))$ one and two-qubit quantum gates, and $\tilde{O}(kn)$ bits of classical randomness.*

The same construction using computationally pseudorandom permutations (instead of k-wise independent permutations) gives a pseudorandom unitary ensemble with nonadaptive security.

Corollary I.1 (Parallel PRU from quantum-secure OWF)**.** *The existence of a one-way function secure against quantum attack implies an efficient quantum algorithm to generate parallel-secure pseudorandom unitaries.*

Our algorithm is fundamentally different than existing approaches for generating unitary k-designs [BHH16], [Haf22], [HHJ21]. In particular, instead of using random quantum circuits, we exponentiate Hamiltonians, which are sums of independent random terms. Our proof uses tools from random matrix theory to create efficient unitary designs.

A. The random sum construction

Rather than employ the mixing properties of *products* of random matrices to achieve convergence, we aggregate randomness from *sums* of random matrices [CDX+24]. In particular, we focus on the Gaussian Unitary Ensemble (GUE, Hermitian matrices with i.i.d. Gaussian entries) a the stepping stone to Haar random unitaries. In a nutshell, our main construction is the following:

(i) Consider a *finite* sum over i.i.d. Hermitian matrices H_j which match the first q-moments of the GUE

$$\boldsymbol{H} \sim \frac{1}{\sqrt{m}}\sum_{j=1}^{m}\boldsymbol{H}_j \quad \text{where} \quad \mathbb{E}[\boldsymbol{H}_j^{\otimes k}] \approx \mathbb{E}[\boldsymbol{G}^{\otimes k}]$$

for each $k = 1, \ldots, q$.

[1]We emphasize that this refers to the quantum gate count — which is typically the bottleneck in applications — rather than the scaling in the number of bits of classical randomness, which was recently achieved by [OSP23].

[2]We note there have also been plausible arguments that certain continuous time Brownian motions should attain linear scaling [JBS23] but the cost of simulating them on a quantum computer remains to be analyzed.

[3]The latter construction [CK23] is much simpler and uses finite-field arithmetic. The classical efficiency of the earliest construction of Kassabov [Kas07] has been folklore and has been fully verified in [CHH+24].

(ii) Then, the product of two exponentiations of such matrices is approximately Haar for appropriate $\theta = \mathcal{O}(1)$:

$$\mathrm{e}^{\mathrm{i}\theta \boldsymbol{H}} \mathrm{e}^{\mathrm{i}\theta \boldsymbol{H}'} \overset{(1)}{\approx} \underbrace{\mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}} \mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}'}}_{=:\boldsymbol{W}} \overset{(2)}{\approx} \boldsymbol{U}_{Haar}. \tag{1}$$

The first approximation (1) results from a matrix central limit theorem: sums of random matrices matching low moments converge to the GUE matrix $\boldsymbol{G}$. Remarkably, random matrices as simple as signed random Pauli strings $\boldsymbol{\sigma}$, which match the second moments of GUE ($\mathbb{E}[\boldsymbol{\sigma} \otimes \boldsymbol{\sigma}] = \mathbb{E}[\boldsymbol{G} \otimes \boldsymbol{G}]$), can reproduce the very complex GUE. This elementary approach circumvents the spectral gap calculations key to prior works [BHH16], [HHJ21], [Haf22]. More precisely, we significantly generalize a matrix Lindeberg principle from [CDB+23] to control the convergence rate on m with a suitable choice of q concerning a very different distance.

The second approximation (2) is a highly nontrivial conversion from the GUE to Haar random unitaries. Although Gaussian matrices have a Haar-random basis (i.e. they are unitarily invariant under conjugation), they have spectra very different from Haar—roughly, the eigenvalues of GUE matrices are semicircle-distributed while the eigenvalues of Haar-random unitaries are evenly distributed around the unit circle. This difference remains even after we exponentiate $\boldsymbol{G}$ to get unitarity $e^{i\boldsymbol{G}\theta}$. This difference between an $N \times N$ Haar unitary $\boldsymbol{U}$ and an exponentiated GUE $\boldsymbol{G}$ is exhibited by their trace moments:

$$\frac{1}{N} \mathrm{Tr}[\mathrm{e}^{\mathrm{i}\boldsymbol{G}\theta_p}] \not\approx \frac{1}{N} \mathrm{Tr}[\boldsymbol{U}^p] \approx 0. \tag{2}$$

Surprisingly, we show that the product of merely two Gaussian exponentials $\boldsymbol{W} = \mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}} \mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}'}$ *is close* to a Haar random unitary at particularly chosen values of the numerical *constant* θ (independent of the system size). That is, the product $\boldsymbol{W}$ is an ϵ-approximate unitary k-design for a very high value of k ($k = 2^{\Omega(n)}$) and a small value of ϵ. Intuitively, the two independent Haar-random bases wash away the structure in the spectra. Proving this requires both developing a novel query complexity lower bound specifically for distinguishing ensembles with symmetries, as well as new random matrix theory results for bounding moments of the ensemble $\boldsymbol{W}$ using a novel large-N polynomial method. We describe how this works in section II-A.

B. The random permutation construction

The basic building blocks of our second construction [CBB+24] are *random phased permutations*, that is random permutation with random *complex* signs:

$$\boldsymbol{Z} := \boldsymbol{D}_z \cdot \boldsymbol{S}, \qquad \text{(random phased permutations)}$$

$$\text{where} \quad \boldsymbol{S} \overset{unif}{\sim} \boldsymbol{S}(N)$$

(uniformly random permutations on N elements)

$$(\boldsymbol{D}_z)_{ij} = \delta_{ij} z_i \quad \text{where} \quad z_i \overset{i.i.d.}{\sim} \boldsymbol{U}(1).$$

(random diagonal complex phases)

Though our ultimate goal is pseudorandomness, in most of the discussion it suffices to treat the matrix $\boldsymbol{Z}$ as truly random since it will eventually be replaced by classical pseudorandom counterparts.

The central object built from random phased permutations $\boldsymbol{Z}$ is the matrix exponential

$$\mathrm{e}^{\mathrm{i}\theta_m \boldsymbol{A}_m} \quad \text{for} \quad \theta_m = \mathcal{O}(1). \tag{3}$$

The $2m$-sparse Hermitian matrix $\boldsymbol{A}_m$ is a sum of i.i.d. random matrices $\boldsymbol{Z}$ and their adjoints

$$\boldsymbol{A}_m := \frac{1}{\sqrt{2m}} \sum_{a=1}^{m} (\boldsymbol{Z}_a + \boldsymbol{Z}_a^\dagger) \quad \text{where} \quad \boldsymbol{Z}_a \overset{i.i.d.}{\sim} \boldsymbol{Z}, \tag{4}$$

and can be thought of as the adjacency matrix for a random graph weighted by random phases. The full construction is a product of their i.i.d. exponentials (sandwiched by random phased permutation),

$$\boldsymbol{Z} \cdot \mathrm{e}^{\mathrm{i}\theta_m \boldsymbol{A}_m^{(1)}} \cdots \mathrm{e}^{\mathrm{i}\theta_m \boldsymbol{A}_m^{(\ell)}} \cdot \boldsymbol{Z}' \approx \boldsymbol{U}_{Haar} \tag{5}$$

The key insight of this ensemble is that it forms an approximate k-design for a *superpolynomial* value of k, using only very few, specifically $\tilde{\mathcal{O}}(m\ell)$, random permutations from $\boldsymbol{S}(N)$. Since sampling uniformly random permutations requires exponentially many bits of randomness, we are not evading known lower bounds on the cardinality of a unitary design [BHH16], [RY17]. Crucially, however, the number of iterations ℓ can be small. For example, setting $m = 2$ and $\ell = n$ yields an approximate k-design for very large value of $k = 2^{\Omega(n)}$, but even $m = 2$ and $\ell = \log^2(n)$ suffices for a superpolynomial design; on the contrary, spectral gap approaches (e.g., [BHH16]) to unitary designs often cost $\Omega(nk)$ rounds of products.

In a nutshell, multiplying sparse matrices is an efficient way to get a dense matrix, but controlling the diamond distance from Haar requires careful analysis. Prior approaches to unitary k-designs have often been based on the spectral gaps of random walks (e.g., [BHH16]), which bootstrap the statistical distance from a comparatively more tractable spectral gap. However, this approach necessarily requires a log-dimensional factor $\mathcal{O}(\log(N^k)) = \mathcal{O}(nk)$ multiplicative blowup in the gate complexity due to the conversion from 2-norm to 1-norm. For cryptographic applications, the attacker may perform an arbitrary polynomial number of queries, so security requires a *fixed* poly-time construction of a superpolynomial k-design, posing a fundamental barrier for spectral gap approaches.

The essence of our alternative argument can be captured in the following observation:

$$\begin{aligned} \boldsymbol{G}_{Ginibre} &\overset{(1)}{\approx} \frac{1}{\sqrt{m^\ell}} \sum_{i=1}^{m^\ell} \boldsymbol{Z}_i \\ &\overset{(2)}{\approx} \left(\frac{1}{\sqrt{m}} \sum_{i=1}^{m} \boldsymbol{Z}_i^{(\ell)} \right) \cdots \left(\frac{1}{\sqrt{m}} \sum_{i=1}^{m} \boldsymbol{Z}_i^{(1)} \right), \end{aligned} \tag{6}$$

where the $\boldsymbol{Z}_i$ are i.i.d. random phased permutations. The first approximation (1) is again a central limit theorem as in (1). It allows us to obtain a very nice random matrix, $\boldsymbol{G}$, drawn from the Ginibre ensemble [Gin65], which is a complex Gaussian matrix that is both left and right unitarily invariant. However,

the CLT-type convergence rate is too slow (polynomial in the number of summands) and will incur a large $\Omega(\mathrm{poly}(k))$-cost scaling with the number of queries k.

The crux of our argument for circumventing the large number of i.i.d. copies (which is a root cause of the k^2 scaling in theorem I.1) is the second approximation (2): a product of sums reproduces the statistics of an *independent* sum but using many fewer ($\tilde{\mathcal{O}}(m\ell)$) copies of $\boldsymbol{Z}_i$. Roughly, in the large-N limit, distinct words of permutations are effectively independent of each other.

As an instructive example, when the dimension N is large, the following correlated words of $\boldsymbol{Z}_1, \boldsymbol{Z}_2$ acting on $|i\rangle$ are almost independent of each other:

$$\begin{aligned}&(\boldsymbol{Z}_1|i\rangle, \boldsymbol{Z}_2|i\rangle, \boldsymbol{Z}_1\boldsymbol{Z}_2|i\rangle, \boldsymbol{Z}_2\boldsymbol{Z}_1|i\rangle)\\ &\overset{dist}{\approx} (\boldsymbol{Z}_1|i\rangle, \boldsymbol{Z}_2|i\rangle, \boldsymbol{Z}_3|i\rangle, \boldsymbol{Z}_4|i\rangle).\end{aligned} \tag{7}$$

Indeed, knowing $\boldsymbol{Z}_1|i\rangle$ and $\boldsymbol{Z}_2|i\rangle$ tells us nothing about $\boldsymbol{Z}_2\boldsymbol{Z}_1|i\rangle$, unless the very unlikely collision $\boldsymbol{Z}_2|i\rangle \propto |i\rangle$ occurs. Applying this intuition to the product of sums, we can get m^ℓ independent $\boldsymbol{Z}$s from merely $\tilde{\mathcal{O}}(m\ell)$ many truly independent $\boldsymbol{Z}$s!

At first glance, the above analysis may appear strange, as if we get "more randomness for free" from a much smaller ($m\ell$) number of independent elements. Careful thought reveals the approximation in Eq. (6) is possible because we only consider *low-moment* properties of the permutations, i.e., we are fixing $k \ll N$ and then taking a large-N limit. In the $N \to \infty$ limit, the large amount of randomness in the permutations themselves effectively "decouples" the different terms, as collisions become vanishingly improbable. Strictly speaking, this precise statement only holds in a *non-adaptive* setting, i.e., when the inputs are fixed in advance. Of course, sequential/adaptive queries would reveal correlations between these words – for example, if one were able to query $\boldsymbol{Z}_2$ *after* knowing the result of $\boldsymbol{Z}_1|i\rangle$, it would coincide with $\boldsymbol{Z}_2\boldsymbol{Z}_1|i\rangle$. However, the key point is that under non-adaptive queries[4], such attacks are not possible, and the words effectively decouple.

Unfortunately, proving the validity of the argument above in finite dimensions N is nontrivial. There are nonasymptotic correlations, however tiny, between the distinct words because those words are made of only $\tilde{\mathcal{O}}(m\ell)$-many independent random phased permutations. To complete the proof, the remaining and most substantial technical argument is a framework to control those finite-N corrections effectively (see section II-B).

C. From permutations, designs, to PRUs

So far, the above "random permutation" construction assumes truly random permutations, which are costly to sample from. To achieve algorithmic efficiency, the key point is that by substituting *pseudorandom* permutations in their place, we establish the link between classical pseudorandom permutations and quantum pseudorandom unitaries, showing the aforementioned theorem I.2 and corollary I.1. First, we consider substituting information-theoretically secure pseudorandom permutations (i.e., k-wise independent permutations) in place of the truly random permutations to obtain an efficient unitary k-design. This requires choosing suitable parameters m and ℓ.

Corollary I.2 (Unitary k-designs from k'-wise independence (informal))**.** *There is a constant c such that for $k \leq 2^{cn}$, an ϵ-approximate quantum unitary k-design can be efficiently implemented by applying our construction with k'-wise independent (discrete) phases and k'-wise independent permutations in place of the truly random phases/permutations, where*

$$k' = \mathcal{O}(k \log(k/\epsilon)). \tag{8}$$

Here, we are making key use of the fact that our construction only relies on $k' = \tilde{\mathcal{O}}(k)$-th moments of the permutations via efficient Hamiltonian simulation algorithms (e.g. [BCC+15], [GSLW19]). Thus, our result "lifts" the efficient construction of k'-wise independent permutations and functions to the construction of unitary k-designs. This is in a similar spirit to recent results of Brakerski and Shmueli that lift $2k$-wise independent functions to approximate quantum state k-designs [BS19].

To achieve an efficient unitary design, we now need to leverage the fact that k'-wise independent permutations and functions can be implemented in only $\mathcal{O}(\mathrm{poly}(n)k')$ time. However, one caveat is that the best known classical k'-wise independent permutations are only approximately k'-wise independent. Fortunately, explicit constructions of high-accuracy classical k-wise independent permutations do exist, with sufficiently low error. This yields unitary k-designs that are almost algorithmically and information-theoretically optimal (in terms of k dependence and the diamond distance).

To achieve computational pseudorandomness, we substitute cryptographically secure pseudorandom permutations (i.e., PRPs) in place of the truly random permutations to obtain a parallel-secure PRU:

Corollary I.3 (Quantum secure-PRP implies parallel-PRU)**.** *Suppose quantum-secure pseudorandom permutations exist. Then our construction $\boldsymbol{V}$ gives a pseudorandom unitary under nonadaptive queries.*

Note that the existence of quantum-secure PRPs (with inverses) only requires assuming the existence of quantum-secure pseudorandom functions [Zha16], which is a standard cryptographic assumption. Furthermore, if there exists a low-depth implementation of quantum-secure PRPs, then our PRU can also be low-depth with suitable parameters m, ℓ and standard quantum algorithmic implementations.

II. A Polynomial Method for Random Matrices

Both of our constructions feature classical random matrix ensembles: random Gaussian matrices, random permutations, and Haar random unitaries. In principle, since these families are reasonably nice objects, we could brute force through the combinatorics for whichever quantity f we care about, whether

[4]We conjecture our ensemble also gives adaptive security, but we note this would require further proof ideas, such as defining a more refined notion of independence of different words.

that is their moments or certain distinguishing probability. With more work, we have all the data necessary to perform a *large-N expansion*: the Weingarten calculus and its variants yield sums over diagrams with coefficients being (rational) functions of $\frac{1}{N}$

$$f(\frac{1}{N}) = f_0 + \frac{f_1}{N} + \cdots . \tag{9}$$

To prove our quantitative results, however, it is necessary to control the total contribution coming from all orders in the $1/N$ expansion. It appears very challenging to systematically capture the fine-grained combinatorics and *cancellations* required to deliver strong enough nonasymptotic bounds required in the k-design context.

Our main strategy can be summarized by the following guiding principle:

"Suitable" random matrix properties at finite dimension N are controlled by the large-N limit.

Conceptually, we are looking for an *interpolation* argument: instead of directly calculating a complicated random matrix quantity at finite N, we start with the much simpler large-N limit, and control the finite-N corrections by arguing that the function "changes slowly" as a function of $\frac{1}{N}$.

The crucial mathematical tool is a basic fact in polynomial approximation: Markov's "other" inequality states that a real-valued, *low-degree* polynomial $f(x)$ whose values are *bounded* on an interval will have bounded derivatives.

Lemma II.1 (Markov's inequality for polynomials [Mar89], [Mar16], [RC66], [EZ64])**.** *Let $f(x) : \mathbb{R} \to \mathbb{R}$ be a real polynomial of degree d. Then,*

$$\sup_{x\in[0,1]} |f'(x)| \leq 2d^2 \max_{x\in[0,1]} |f(x)| .$$

This seemingly innocent property of polynomials has found profound implications in establishing lower bounds in classical circuit [Bei93] and quantum query complexity, a strategy referred to as the *polynomial method*, e.g. [BBC+01], [AS04], [BSS01], [NW99], [Kut05], [Raz03]. The standard recipe for applying the above inequality to lower-bounding the number of queries required for distinguishing tasks goes roughly as follows: cook up a quantity f that depends on an (often nonobvious) interpolation parameter x such that

- (**End-points**.) The values $f(x_1)$ and $f(x_2)$ correspond to the acceptance probabilities for the two cases to be distinguished.
- (**Low-degree**.) The function $f(x)$ is a low-degree polynomial with a degree d, roughly the number of queries.
- (**Apriori bound**.) The function $f(x)$ can be *extended* to a larger interval $[x_1, x_3]$ with $x_3 - x_1 \gg x_2 - x_1$ such that its (unknown) value remains bounded.

Then, the distinguishing probability must not change too quickly between the two cases provided the degree d is small:

$$|f(x_1) - f(x_2)| \leq 2d^2 \frac{|x_1 - x_2|}{|x_1 - x_3|} \max_{x\in[x_1,x_3]} |f(x)| . \tag{10}$$

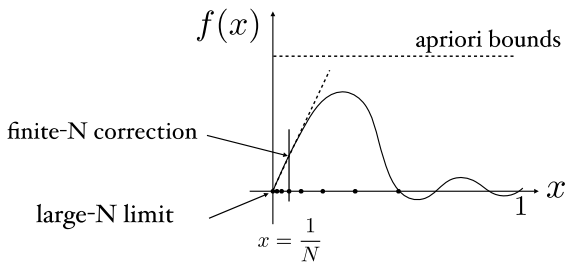

Fig. 1. Interpolating the $\frac{1}{N}$ expansion. For suitable quantities dependent on the dimension N, we control the finite-N behavior by interpolating from the large-N limit. This argument requires that (1) the quantity of interest be a low-degree rational function of N, (2) the function obeys a good apriori bound, and (3) the large-N limit can be effectively calculated.

While explicitly understanding *all possible* adversarial quantum algorithms is essentially impossible, we often do have sufficient structural restrictions to guarantee that $f(x_1) - f(x_2)$ is a bounded degree polynomial which, remarkably, proves to be sufficient.

For our purposes, the key insight is to draw a direct connection between the large-N expansion in random matrix theory and polynomial approximation by setting the interpolation parameter to be (fig. 1)

$$x = \frac{1}{N},$$

$$\text{such that} \quad \left| f\left(\frac{1}{\infty}\right) - f\left(\frac{1}{N}\right) \right| \leq \frac{2d^2}{N} \max_{0\leq x\leq 1} |f(x)| , \tag{11}$$

which requires constructing a suitable function that fulfills the requirements of the polynomial methods, which is the most substantial part of our proofs.

A. The random sum construction

A key step in the "random sum" construction (ii) is that the moments for the product $\boldsymbol{W} = \mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}} \mathrm{e}^{\mathrm{i}\theta \boldsymbol{G}'}$ are very close to the Haar value

$$|\mathbb{E} \operatorname{Tr}[\boldsymbol{W}^p]| \leq O\left(\frac{\operatorname{poly}(p)}{\operatorname{poly}(N)}\right) \approx 0.$$

We are not aware of existing non-asymptotic bounds of this type; what is well-known is the infinite dimension limit $N \to \infty$ (from e.g., free probability),

$$\lim_{N\to\infty} |\mathbb{E} \operatorname{Tr}[\boldsymbol{W}^p]| = 0,$$

which is insufficient if we care about how these quantities scale as a function of both N and p.

In principle, such quantities can be computed exactly by Weingarten calculus and diagrammatic expansion (the GUE matrices are diagonal in Haar random bases), but it often requires a difficult, tedious, and nonsystematic calculation. Generally, the number of diagrams in the Weingarten calculation increases *factorially* in the number of moments p one considers—even at leading order in $1/N$—and nontrivial cancellations in the

Weingarten series are difficult to control. Simply applying the triangle inequality to the Weingarten series yields very loose upper bounds, like $\mathcal{O}(\frac{p!}{N})$ as they cannot capture these intricate cancellations. Fortunately, for our usage, we do not need the exact value of these quantities but rather just an upper bound to the magnitude of the correction: for example, does it scale like $\mathcal{O}(p/N)$, or $\mathcal{O}(2^p/N^2)$?

Our solution is to apply the polynomial method for $f(\frac{1}{N}) = \mathbb{E}\,\mathrm{Tr}[\boldsymbol{W}^p]$, going through the checklist:

- (**Low-degree in** $\frac{1}{N}$) The function $f(\frac{1}{N})$ is low-degree (rational) in $\frac{1}{N}$, rooted in Weingarten calculus.
- (**Large-**N **limit**) The value at $N \to \infty$ is known to be $f(0) = 0$.
- (**Apriori bounds**) For each $1 \le N' \le \infty$, $\left|f(\frac{1}{N})\right| \le 1$.

This method can be used to prove much tighter upper bounds on moments without a hard Weingarten calculus: it yields bounds of the form $O\left(\frac{\mathrm{poly}(p)}{\mathrm{poly}(N)}\right)$, which is qualitatively stronger than the naive bound in p dependence and captures the *cancellations* occurring in the Weingarten sum. This observation drastically simplifies many of our moment calculations and circumvents hard Weingarten calculations, instead replacing them with (more tractable) challenges in applying the polynomial method. The alternative problem that we had to resolve to obtain an apriori bound was to manually extend the function across different dimensions $N' \neq N$, by solving a moment problem for unitary matrices (see Section III).

B. The random permutation construction

Applying the polynomial method to the "random permutation" construction requires a rather different implementation. In particular, we choose the function of interest f to be the distinguishing probability between the unitaries $\boldsymbol{V} = \boldsymbol{Z}_L \left(\prod_{j=1}^{\ell} \mathrm{e}^{\mathrm{i}\theta_m \boldsymbol{A}_m^{(j)}}\right) \boldsymbol{Z}_R$ and $\boldsymbol{U}_{Haar}$ by defining the key quantity

$$f_{\boldsymbol{\rho},\boldsymbol{O}}\left(\frac{1}{N}\right) := \mathrm{Tr}\left[\boldsymbol{O}\mathcal{N}_1\boldsymbol{\rho}\right] - \mathrm{Tr}\left[\boldsymbol{O}\mathcal{N}_2\boldsymbol{\rho}\right] \quad \text{for each} \quad \boldsymbol{\rho} \quad \text{and} \quad \boldsymbol{O} \tag{12}$$

where $\mathcal{N}_1 = \mathcal{N}_{2k,\boldsymbol{V}}$ and $\mathcal{N}_2 = \mathcal{N}_{2k,\boldsymbol{U}_{Haar}}$ are the corresponding channels.

We aim to control the finite-N distinguishing probability $f_{\boldsymbol{\rho},\boldsymbol{O}}(\frac{1}{N})$ from the large-N limit $f_{\boldsymbol{\rho},\boldsymbol{O}}(\frac{1}{\infty})$ *for arbitrary fixed* $\boldsymbol{\rho}$ *and* $\boldsymbol{O}$ that the quantum attacks may use. We spell out the structural requirements for $f_{\boldsymbol{\rho},\boldsymbol{O}}(\frac{1}{N})$ for interpolation:

Low-degree rational functions of N. The random permutations are very nice objects that are defined consistently across dimensions $N' \neq N$, inducing channels $\mathcal{N}_1^{(N')}$ and $\mathcal{N}_2^{(N')}$. In particular, the average over permutations can be calculated using a well-defined diagrammatic expansion, with coefficients that are low-degree rational functions of N.[5] However, the problem is that the test state $\boldsymbol{\rho}$ and the test operator $\boldsymbol{O}$, which are selected adversarially, do not obviously apply in other dimensions. Finding the suitable extension requires a few arguments: first, the permutation symmetry substantially reduces the number of parameters. Specifically, the Schur-Weyl duality for the commutant of permutations restricts the effective input state $\boldsymbol{\rho}$ to a much smaller object, the *partition algebra* $\boldsymbol{P}_k(N)$, whose *algebraic structure* only depends on the number of copies k, and is in particular *independent* of N as long as the dimension is large $N \ge 2k$ [HR05], [HJ20].[6] While the partition algebra has been studied in the algebraic combinatorics literature, our argument necessitates explicitly describing the embedding with respect to the computational basis in order to verify that the embedding "does not change too quickly" as N increases. This requires a detailed foray into the structure of the partition algebra and writing down an explicit orthogonal basis as a diagrammatic sum to establish that the basis coefficients are rational polynomials in $\frac{1}{N}$.

The above understanding of the algebraic structure allows us to handcraft the suitable extension $f_{\boldsymbol{\rho},\boldsymbol{O}}$ for other dimensions $N' \neq N$

$$f_{\boldsymbol{\rho},\boldsymbol{O}}\left(\frac{1}{N'}\right) = \mathrm{Tr}\left[\boldsymbol{O}^{(N')}\mathcal{N}_1^{(N')}\boldsymbol{\rho}^{(N')}\right] - \mathrm{Tr}\left[\boldsymbol{O}^{(N')}\mathcal{N}_2^{(N')}\boldsymbol{\rho}^{(N')}\right] \tag{13}$$

by defining a family of test operators

$$\left(\boldsymbol{O}^{(N')}, \boldsymbol{\rho}^{(N')}\right) \quad \text{for each} \quad 2k \le N' \le \infty \tag{14}$$

$$\text{such that} \quad \left(\boldsymbol{O}^{(N)}, \boldsymbol{\rho}^{(N)}\right) = (\boldsymbol{O}, \boldsymbol{\rho})\,. \tag{15}$$

The operators $\boldsymbol{\rho}^{(N)}, \boldsymbol{O}^{(N)}$ for other dimensions N' relate to the original $\boldsymbol{\rho}, \boldsymbol{O}$ in dimension N by substituting the basis we found for the partition algebra. In the end, the function $f_{\boldsymbol{\rho},\boldsymbol{O}}(\frac{1}{N})$ is (approximately[7]) a rational polynomial of N with degree $\mathrm{poly}(k)$ and poles at small integers $1, \ldots, \tilde{\mathcal{O}}(2k\ell)$. [8]

A priori bounds. The expression is a difference between probabilities such that

$$\left|f_{\boldsymbol{\rho},\boldsymbol{O}}\left(\frac{1}{N'}\right)\right| \le 2 \quad \text{for each integer} \quad N'. \tag{16}$$

Large-N **limits**. The large-N limits of each channel coincide a pair of nicer ones: $\mathcal{N}_1$ to the sum over independent permutations $\sum w_{\boldsymbol{i}} \boldsymbol{W}_{\boldsymbol{i}} \to \sum w_{\boldsymbol{i}} \boldsymbol{Z}_{\boldsymbol{i}}$, and $\mathcal{N}_2$ to the "Gaussian" model: $\boldsymbol{U}_{Haar} \to \boldsymbol{G}_{Ginibre}$.

$$\begin{aligned} \mathrm{Tr}\left[\boldsymbol{O}_{N'}\mathcal{N}_1\boldsymbol{\rho}_{N'}\right] &\overset{N'\to\infty}{=} \mathrm{Tr}\left[\boldsymbol{O}_{N'}\mathcal{N}_1^{(free)}\boldsymbol{\rho}_{N'}\right] \\ \mathrm{Tr}\left[\boldsymbol{O}_{N'}\mathcal{N}_2\boldsymbol{\rho}_{N'}\right] &\overset{N'\to\infty}{=} \mathrm{Tr}\left[\boldsymbol{O}_{N'}\mathcal{N}_2^{(Ginibre)}\boldsymbol{\rho}_{N'}\right]. \end{aligned} \tag{17}$$

The independent sums $\mathcal{N}_1^{(free)}$ and Gaussian $\mathcal{N}_1^{(Ginibre)}$ can be compared by the Lindeberg principle [CDB+23], [CBB+24],

[5] Qualitatively, this can be regarded as the analog of Weingarten expansion in the unitary case.

[6] In fact, the partition algebras stabilize and are all *isomorphic* for large enough dimension $N \ge 2k$.

[7] We need to truncate the rapidly converging Taylor expansion for the exponential function.

[8] The resulting rational function in $1/N$ will have poles that can be handled by extended versions of Markov-type inequalities.

with an error suppressed by $\mathrm{poly}(1/m^\ell)$ typical in central limit theorems. While the most general Lindeberg argument works in any finite dimension N, the large-N limits simplify the calculations.

III. The Moment Problem

A important ingredient in our proof of Theorem I.1 is finding solutions to the moment problem mentioned above in Section II. In general, a moment problem is defined as follows: given a list of values $\alpha_1, \alpha_2, \ldots \alpha_p$ and a probability space, find a measure whose ith moment is α_i. Variants of this problem arise in several places in our proof. First, as just discussed previously, it arises in our application of the polynomial method to bounding the value of random matrix theory quantities. In that context, we have some diagonal matrix $\boldsymbol{D}$ of dimension N, and we wish to find diagonal matrices of different dimensions which yield the same moments. We need to solve the moment problem in many different dimensions for the proof to work, as these solutions yield bounds on the polynomial at different points.

Separately, the moment problem arises in our polynomial method query lower bound to show that the matrices $e^{i\boldsymbol{G}\theta}e^{i\boldsymbol{G}'\theta}$ are T-query indistinguishable from Haar-random for a high value of T. In particular, after showing the trace moments of our ensemble are close to Haar, we define a linear interpolating path between our moments and Haar in some parameter x. The acceptance probability of our query algorithm is then a low degree univariate polynomial in x, and must jump in some small range if the algorithm distinguishes our ensemble from Haar. To show distinguishability, and hence this jump, is not possible, we must show the polynomial is bounded at many points along the interpolation path. This requires proving that unitary matrices exist that have a wide range of moment vectors.

The moment problem is well-studied in the mathematics literature—see Chapter 11 of [Sch91] for a related variant. However, to the best of our knowledge, off-the-shelf statements from the current literature are too weak to use for our proofs. The basic issue is that in our context, we need to show there exist *finite-dimensional* unitary matrices that have certain trace moments. This is a more "discrete" version of the moment problem than has been previously considered.[9] It is analogous to asking, given some list of empirical moments tabulated from finitely many samples, could those empirical moments have been exactly reproduced by a different number of samples?

To overcome this, we produce a solution to the moment problem for finite-dimensional unitary matrices. Given a list of moments, we start by finding a "nearby" moment vector that admits a discrete solution with a matrix. We then show that this solution can be perturbed to cover the nearby moment vector, while maintaining the discreteness of the measure. This follows from the Jacobian of the moment vector—i.e., the derivative of the moment vector under perturbing individual points—being full rank. Therefore, perturbations of the input point completely cover a small ball around the starting point, a region that is large enough for our applications.

[9]To the best of our knowledge, the closest result that exists is for "atomic measures" which places mass at discrete points but with real-valued probabilities. However, this is not enough, as for N-dimensional matrices, the probability masses must be integer multiples of $1/N$.

IV. Discussions

A. Related work

This article is merged from two pre-prints [CBB+24], [CDX+24]. In the early stages, the conceptual connection between large-N limit and Markov inequality was passed between the concurrent works [CGVTvH24], [CBB+24], [CDX+24] through the first author. Nevertheless, in each scenario, the technical arguments needed to realize this idea appear rather different, and we include self-contained expositions. When [CDX+24] was nearly completion, we became aware of the independent related work of Haah, Liu, and Tan [HLT24] achieving similar dependence on k via a different construction; when [CBB+24] was near completion, we became aware of the independent work of Metger, Poremba, Sinha, and Yuen [MPSY24], which obtain similar results for parallel-secure PRUs and k-designs via a completely different construction and analysis.

B. Acknowledgments

We thank Chris Bowman, Dmitry Grinko, Jeongwan Haah, Aram Harrow, Jonas Haferkamp, Tom Halverson, William He, Hsin-Yuan Huang, Martin Kassabov, Yunchao Liu, Fermi Ma, Tony Metger, Quynh Nguyen, Ryan O'Donnell, Alexander Poremba, Arun Ram, Makrand Sinha, Norah Tan, Joel Tropp, Jorge Garza Vargas, Shunyu Yao, and Henry Yuen for stimulating discussions. We thank the Simons Institute for the Theory of Computing.

References

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

[ABI86] Noga Alon, László Babai, and Alon Itai, *A fast and simple randomized parallel algorithm for the maximal independent set problem*, Journal of Algorithms **7** (1986), no. 4, 567–583.

[ABW09] Andris Ambainis, Jan Bouda, and Andreas Winter, *Non-malleable encryption of quantum information*, Journal of Mathematical Physics **50** (2009), no. 4, 042106.

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

[AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen, *Pseudorandom (function-like) quantum state generators: New definitions and applications*, Theory of Cryptography Conference, Springer, 2022, pp. 237–265.

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

[AS04] Scott Aaronson and Yaoyun Shi, *Quantum lower bounds for the collision and the element distinctness problems*, Journal of the ACM (JACM) **51** (2004), no. 4, 595–605.

[BBC+01] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald De Wolf, *Quantum lower bounds by polynomials*, Journal of the ACM (JACM) **48** (2001), no. 4, 778–797.

[BCC+15] Dominic W Berry, Andrew M Childs, Richard Cleve, Robin Kothari, and Rolando D Somma, *Simulating hamiltonian dynamics with a truncated taylor series*, Physical review letters **114** (2015), no. 9, 090502.

[Bei93] Richard Beigel, *The polynomial method in circuit complexity*, [1993] Proceedings of the Eigth Annual Structure in Complexity Theory Conference, IEEE, 1993, pp. 82–95.

[BFV19] Adam Bouland, Bill Fefferman, and Umesh Vazirani, *Computational pseudorandomness, the wormhole growth paradox, and constraints on the ads/cft duality*, arXiv preprint arXiv:1910.14646, 2019.

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

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

[BNOZ22] Eiichi Bannai, Yoshifumi Nakata, Takayuki Okuda, and Da Zhao, *Explicit construction of exact unitary designs*, Advances in Mathematics **405** (2022), 108457.

[BNZZ19] Eiichi Bannai, Mikio Nakahara, Da Zhao, and Yan Zhu, *On the explicit constructions of certain unitary t-designs*, Journal of Physics A: Mathematical and Theoretical **52** (2019), no. 49, 495301.

[BS18] Adam R Brown and Leonard Susskind, *Second law of quantum complexity*, Physical Review D **97** (2018), no. 8, 086015.

[BS19] Zvika Brakerski and Omri Shmueli, *(pseudo) random quantum states with binary phase*, Theory of Cryptography Conference, Springer, 2019, pp. 229–250.

[BSS01] Howard Barnum, Michael Saks, and Mario Szegedy, *Quantum decision trees and semidefinite programming.*, Tech. report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States), 2001.

[CBB+24] Chi-Fang Chen, Adam Bouland, Fernando GSL Brandão, Jordan Docter, Patrick Hayden, and Michelle Xu, *Efficient unitary designs and pseudorandom unitaries from permutations*, arXiv preprint arXiv:2404.16751 (2024).

[CDB+23] Chi-Fang (Anthony) Chen, Alexander M Dalzell, Mario Berta, Fernando GSL Brandão, and Joel A Tropp, *Sparse random hamiltonians are quantumly easy*, 2023.

[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.

[CGVTvH24] Chi-Fang Chen, Jorge Garza-Vargas, Joel A Tropp, and Ramon van Handel, *A new approach to strong convergence*, arXiv preprint arXiv:2405.16026 (2024).

[CHH+24] Chi-Fang Chen, Jeongwan Haah, Jonas Haferkamp, Yunchao Liu, Tony Metger, and Xinyu Tan, *Incompressibility and spectral gaps of random circuits*, arXiv preprint arXiv:2406.07478 (2024).

[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** (2023), no. 11, 7983–8021.

[CLLW16] Richard Cleve, Debbie W. Leung, Li Liu, and Chunhao Wang, *Near-linear constructions of exact unitary 2-designs*, Quantum Inf. Comput. **16** (2016), no. 9, 721–756.

[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** (2009), no. 1, 012304.

[DLT02] David P DiVincenzo, Debbie W Leung, and Barbara M Terhal, *Quantum data hiding*, IEEE Transactions on Information Theory **48** (2002), no. 3, 580–598.

[EAŻ05] Joseph Emerson, Robert Alicki, and Karol Życzkowski, *Scalable noise estimation with random unitary operators*, Journal of Optics B: Quantum and Semiclassical Optics **7** (2005), no. 10, S347.

[EZ64] H. Ehlich and K. Zeller, *Schwankung von polynomen zwischen gitterpunkten*, Mathematische Zeitschrift **86** (1964), 41–44.

[Gin65] Jean Ginibre, *Statistical Ensembles of Complex, Quaternion, and Real Matrices*, Journal of Mathematical Physics **6** (1965), no. 3, 440–449.

[GSLW19] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe, *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 193–204.

[GTB23] Tudor Giurgica-Tiron and Adam Bouland, *Pseudorandomness from subset states*, arXiv preprint arXiv:2312.09206, 2023.

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

[HHJ21] Jonas Haferkamp and Nicholas Hunter-Jones, *Improved spectral gaps for random quantum circuits: large local dimensions and all-to-all interactions*, Physical Review A **104** (2021), no. 2, 022417.

[HHWY08] Patrick Hayden, Michał Horodecki, Andreas Winter, and Jon Yard, *A decoupling approach to the quantum capacity*, Open Systems & Information Dynamics **15** (2008), no. 01, 7–19.

[HJ19] Nicholas Hunter-Jones, *Unitary designs from statistical mechanics in random quantum circuits*, 2019.

[HJ20] Tom Halverson and Theodore N. Jacobson, *Set-partition tableaux and representations of diagram algebras*, Algebraic Combinatorics **3** (2020), no. 2, 509–538 (en).

[HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill, *Predicting many properties of a quantum system from very few measurements*, Nature Physics **16** (2020), no. 10, 1050–1057.

[HL09a] Aram W Harrow and Richard A Low, *Efficient quantum tensor product expanders and k-designs*, International Workshop on Approximation Algorithms for Combinatorial Optimization, Springer, 2009, pp. 548–561.

[HL09b] ———, *Random quantum circuits are approximate 2-designs*, Communications in Mathematical Physics **291** (2009), 257–302.

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

[HM23] Aram W Harrow and Saeed Mehraban, *Approximate unitary t-designs by short random quantum circuits using nearest-neighbor and long-range gates*, Communications in Mathematical Physics (2023), 1–96.

[HR05] Tom Halverson and Arun Ram, *Partition algebras*, European Journal of Combinatorics **26** (2005), no. 6, 869–921.

[JBS23] Shao-Kai Jian, Gregory Bentsen, and Brian Swingle, *Linear growth of circuit complexity from brownian dynamics*, Journal of High Energy Physics **2023** (2023), no. 8, 1–42.

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

[JMW23] Fernando Granha Jeronimo, Nir Magrafta, and Pei Wu, *Subset states and pseudorandom states*, arXiv preprint arXiv:2312.15285, 2023.

[Jof74] A. Joffe, *On a Set of Almost Deterministic k-Independent Random Variables*, The Annals of Probability **2** (1974), no. 1, 161 – 162.

[Kas07] Martin Kassabov, *Symmetric groups and expander graphs*, Inventiones mathematicae **170** (2007), no. 2, 327–354.

[KKS+23] Ágoston Kaposi, Zoltán Kolarovszki, Adrian Solymos, Tamás Kozsik, and Zoltán Zimborás, *Constructing generalized unitary group designs*, International Conference on Computational Science, Springer, 2023, pp. 233–245.

[KL17] Shelby Kimmel and Yi-Kai Liu, *Phase retrieval using unitary 2-designs*, 2017 International Conference on Sampling Theory and Applications (SampTA), 2017, pp. 345–349.

[KLR+08] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, *Randomized benchmarking of quantum gates*, Phys. Rev. A **77** (2008), 012307.

[KQST23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal, *Quantum cryptography in algorithmica*, Proceedings of the 55th Annual ACM Symposium on Theory of Computing, 2023, pp. 1589–1602.

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

[KTP20] Isaac Kim, Eugene Tang, and John Preskill, *The ghost in the radiation: Robust encodings of the black hole interior*, Journal of High Energy Physics **2020** (2020), no. 6, 1–65.

[Kut05] Samuel Kutin, *Quantum lower bound for the collision problem with small range*, Theory of Computing **1** (2005), no. 1, 29–36.

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

[Ma23] Fermi Ma, December 2023, Personal communication.

[Mar89] A.A. Markov, *On a problem of d.i. mendeleev (russian)*, Zapishi Imp. Akad. Nauk **I12** (1889), 1–24.

[Mar16] V.A. Markov, *Uber polynome die in einem gegebenen intervalle mölichst wenig von null abweichen*, Math. Annalen **77** (1916), 213–258.

[MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen, *Simple constructions of linear-depth t-designs and pseudorandom unitaries*, arXiv preprint arXiv:2404.12647 (2024).

[MY23] Yosuke Mitsuhashi and Nobuyuki Yoshioka, *Clifford group and unitary designs under symmetry*, PRX Quantum **4** (2023), no. 4, 040331.

[NHKW17] Yoshifumi Nakata, Christoph Hirche, Masato Koashi, and Andreas Winter, *Efficient quantum pseudorandomness with nearly time-independent hamiltonian dynamics*, Physical Review X **7** (2017), no. 2, 021006.

[NW99] Ashwin Nayak and Felix Wu, *The quantum query complexity of approximating the median and related statistics*, Proceedings of the thirty-first annual ACM symposium on Theory of computing, 1999, pp. 384–393.

[NZO+21] Yoshifumi Nakata, Da Zhao, Takayuki Okuda, Eiichi Bannai, Yasunari Suzuki, Shiro Tamiya, Kentaro Heya, Zhiguang Yan, Kun Zuo, Shuhei Tamate, et al., *Quantum circuits for exact unitary t-designs and applications to higher-order randomized benchmarking*, PRX Quantum **2** (2021), no. 3, 030339.

[OBK+17] Emilio Onorati, Oliver Buerschaper, Martin Kliesch, Winton Brown, Albert H Werner, and Jens Eisert, *Mixing properties of stochastic quantum hamiltonians*, Communications in Mathematical Physics **355** (2017), 905–947.

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

[Raz03] Alexander A Razborov, *Quantum communication complexity of symmetric predicates*, Izvestiya: Mathematics **67** (2003), no. 1, 145.

[RC66] T. J. Rivlin and E. W. Cheney, *A comparison of uniform approximations on an interval and a finite subset thereof*, SIAM Journal of Numerical Analysis **3** (1966), no. 2, 311–320.

[RY17] Daniel A. Roberts and Beni Yoshida, *Chaos and complexity by design*, Journal of High Energy Physics **2017** (2017), no. 4, 121.

[Sch91] Konrad Schmüdgen, *The moment problem on the unit circle*, ch. 11, pp. 257–279, Springer-Verlag, New York, NY, 1991.

[SDTR13] Oleg Szehr, Frédéric Dupuis, Marco Tomamichel, and Renato Renner, *Decoupling with unitary approximate two-designs*, New Journal of Physics **15** (2013), no. 5, 053022.

[WC81] Mark N Wegman and J Lawrence Carter, *New hash functions and their use in authentication and set equality*, Journal of computer and system sciences **22** (1981), no. 3, 265–279.

[Web15] Zak Webb, *The clifford group forms a unitary 3-design*, 2015.

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

[Zha16] Mark Zhandry, *A note on quantum-secure prps*, 2016.

[Zhu17] Huangjun Zhu, *Multiqubit clifford groups are unitary 3-designs*, Physical Review A **96** (2017), no. 6, 062336.

[ZKGG16] Huangjun Zhu, Richard Kueng, Markus Grassl, and David Gross, *The clifford group fails gracefully to be a unitary 4-design*, arXiv preprint arXiv:1609.08172, 2016.