Derivatives of Quantum Randomness: Separating Pseudorandom Unitaries from Pseudorandom (Function-like) States

Minki Hhan

KAIST, Daejeon, Korea
minkihhan@kaist.ac.kr

Abstract

Quantum computation gives rise to new pseudorandom primitives for states and unitaries, including pseudorandom state generators (PRSGs), pseudorandom function-like state generators (PRFSGs), and pseudorandom unitaries (PRUs). In this paper, we show a full unitary oracle separation between PRFSGs and PRUs.

The separation holds between the strongest state notion and the weakest unitary notion: even adaptively secure, quantum-accessible PRFSGs do not imply non-adaptively secure, forward-only PRUs, even when their implementations are allowed to be non-unitary and use an arbitrary number of ancillary qubits. This reveals a fundamental distinction between pseudorandomness for quantum states and for quantum unitaries.

Our main technical idea is to view a candidate PRU construction with access to state generation oracles as a map from the underlying oracle states to implemented unitaries, and to study the derivatives of this map. These derivatives are inherently low rank, and we exploit this low-rank structure to distinguish the resulting unitaries from truly random ones. We believe this differential perspective may be useful for studying other structural questions about quantum states and unitaries.

AI use disclosure. The author first learned about the Poincaŕe inequality through ChatGPT 5.4 Pro. The mathematical ideas and proofs in this work were developed and completed by the author. ChatGPT 6 Astra was then used to prepare an initial written draft of the full proofs. The author subsequently rewrote and edited the manuscript. The author takes full responsibility for all mathematical claims and the content of this work.

Contents

1 Introduction **3**
- 1.1 Our result . . . 3
- 1.2 Technical overview . . . 4

2 Preliminaries **7**
- 2.1 Distances and tomography . . . 7
- 2.2 Haar randomness and concentration . . . 7
- 2.3 Choi states and purity tests . . . 8
- 2.4 PRFSGs and PRUs . . . 9
- 2.5 The quantum OR test . . . 10

3 Hilbert-valued Functions and Poincaré Inequalities **10**
- 3.1 Functions on spheres and their derivatives . . . 10
- 3.2 Hilbert-valued functions on complex spheres . . . 11
- 3.3 Hilbert-valued functions on products of complex spheres . . . 12

4 CHFS Oracle and Derivatives **14**
- 4.1 The CHFS oracle . . . 14
- 4.2 CHFS oracle derivatives . . . 15
- 4.3 Oracle algorithm derivatives . . . 17
- 4.4 CHFS oracle concentration . . . 20
- 4.5 CHFS oracle simulation . . . 21

5 Oracle Separation of PRUs from PRFSGs **22**
- 5.1 Dichotomy of the averaged Choi state . . . 22
- 5.2 Simulating $\mathcal{G}$ . . . 24
- 5.3 The distinguisher . . . 26
- 5.4 Analysis . . . 27

1 Introduction

Quantum pseudorandomness extends computational randomness from classical strings and functions to intrinsically quantum objects. Ji, Liu, and Song introduced pseudorandom state generators (PRSGs) and pseudorandom unitaries (PRUs) [21]. A PRSG uses a short classical key to prepare a quantum state whose polynomially many copies are computationally indistinguishable from copies of a Haar-random state. A PRU, on the other hand, implements a keyed unitary, or possibly near-unitary, transformation that is computationally indistinguishable from a Haar-random unitary under quantum oracle queries.

PRSGs and PRUs do not have immediate counterparts in the classical theory of pseudorandomness, as they concern two different types of quantum objects: states and transformations. Nevertheless, there is a useful analogy with the classical hierarchy. A PRSG resembles a pseudorandom generator (PRG) in that a short classical seed selects a single pseudorandom object, whereas a PRU resembles a pseudorandom permutation (PRP) or a pseudorandom function (PRF) in that a short key specifies a pseudorandom transformation over a large space. In this analogy, pseudorandom function-like state generators (PRFSGs) [5] may be closer to pseudorandom functions (PRFs): for a fixed key, they associate a pseudorandom object with each classical input. This analogy is imperfect, but it has generated rich cryptographic applications. Starting from PRSGs, quantum pseudorandomness has been used to construct private-key quantum money [21], commitments, encryption, and secure multiparty computation [5, 31, 4, 8]. PRFSGs further support applications such as secret-key encryption and message authentication [5]. PRUs have also found applications in unclonable encryption [32].

Classically, PRGs, PRFs, and PRPs are tightly connected to each other. The GGM construction turns PRGs into PRFs [17], while the Luby-Rackoff construction turns PRFs into PRPs [27]. Thus, at the level of existence, these notions are equivalent. The quantum landscape is considerably more fragmented. At a basic level, quantum pseudorandom primitives can exist relative to oracles where classical pseudorandomness does not exist [22], even with classical oracles [23, 24] or even allowing negligible errors [10, 7].

Even within quantum pseudorandomness, the picture remains fragmented. Constructions of PRSGs and PRFSGs from classical pseudorandomness are well-established [21, 11, 5, 4]. However, constructions of PRUs from classical pseudorandomness were obtained much later and required substantially new techniques [30, 28, 15], while their equivalence with PRS(F)Gs remains unclear. Several intermediate notions were introduced and constructed to capture pseudorandom operations on quantum inputs in the meantime, such as pseudorandom state scramblers [26] and pseudorandom isometries [2].

This raises a natural question: are these notions of quantum pseudorandomness equivalent at the level of existence, as their classical counterparts are? In particular, can PRFSGs be used to construct PRUs? There is an apparent obstacle: a PR(F)SG, or state generation more generally, only needs to operate on a low-dimensional subspace to generate individual pseudorandom states, whereas a PRU must specify a coherent transformation on an entire Hilbert space. This dimensional mismatch has led to partial oracle separations [10, 18] for PRU constructions with severely bounded ancilla space. However, these results leave open whether the ancilla restriction is essential. Since general PRU constructions may use polynomially many ancillary qubits, these results leave open a full oracle separation with unrestricted ancilla space.

1.1 Our result

We give a unitary oracle separation between PRFSGs and PRUs.

Theorem 1.1. *There is a unitary oracle relative to which PRFSGs exist, whereas PRUs do not.*

This separation holds for a strong notion of PRFSGs—adaptively secure, quantum-accessible PRFSGs—and a weak notion of PRUs—non-adaptively secure, forward-only PRUs. Consequently, the same oracle separates any weaker notion of PRFSGs from any stronger notion of PRUs; for example, it separates classically accessible PRFSGs from adaptively secure PRUs with inverse and conjugate access [34]. Furthermore, the impossibility holds even for candidate implementations that are not themselves unitary: they may be implemented by arbitrary efficient quantum channels using polynomially many ancillary qubits and intermediate measurements. That is, our separation even rules out PRU implementations with negligible errors.

Our separation may also be related to the distinction between state synthesis and unitary synthesis. Arbitrary quantum states can be prepared efficiently using one query to a suitable classical oracle [33], whereas an analogous one-query cannot implement arbitrary unitaries [25]. While the connection is unclear, these results similarly highlight a distinction between preparing quantum states and implementing unitary transformations.

Our proof analyzes the derivatives of the oracle PRU construction by interpreting it as a map from oracles (or the states defining oracles) to the space of quantum operations. In this interpretation, the observation that "the state generation can be implemented on a low-dimensional subspace" is captured by the low rank of the gradient, which in turn allows us to show that the norm of the gradient is relatively small on average. Intuitively, the smallness of the derivative makes the output of the map—the resulting quantum operator—insensitive to perturbations, suggesting a concentration phenomenon. A Poincaré inequality formally converts this gradient bound into concentration of the resulting quantum operation. Such concentration is incompatible with the behavior of a random unitary, yielding the separation. To handle non-unitary implementations, we establish a dichotomy: for a suitable input state, either the candidate produces a noticeably mixed output, which already distinguishes it from a unitary, or its output falls into the same concentration regime as above.

1.2 Technical overview

We use the common-Haar function-like state (CHFS) oracle together with the QPSPACE oracle [3, 10, 18]. In the overview, we use the CHFS oracle with a fixed length, say, $d = \lfloor \sqrt{\lambda} \rfloor$ for simplicity. The CHFS oracle is defined by a family of d-qubit pure states $\Phi = (|\phi_x\rangle)_{x\in\{0,1\}^d}$. Given this, the CHFS oracle query is processed by

$$S_d^\Phi = \sum_x |x\rangle\langle x| \otimes S_{\phi_x}$$

in superposition, where S_ϕ swaps between $\left|0^{d+1}\right\rangle$ and $|\phi, 1\rangle$ and is given by

$$S_\phi := I - \left|0^{d+1}\right\rangle\left\langle 0^{d+1}\right| - |\phi,1\rangle\langle\phi,1| + \left|0^{d+1}\right\rangle\left\langle\phi,1\right| + \left|\phi,1\right\rangle\left\langle 0^{d+1}\right|$$

which is inspired by the reflection/swap Haar state oracle [16, 9].

Let $\mathcal{G}^{S_d^\Phi}$ be a candidate PRU. We focus on simulating the application of $\mathcal{G}^{S_d^\Phi}$ on a specific state *without using* S_d^Φ. In the full proof, we use process tomography for logarithmically small d in the simulation. Given the simulation, the quantum OR test [20] over the key together with a strong complexity oracle, say QPSPACE, shows the insecurity of $\mathcal{G}^{S_d^\Phi}$ following, e.g., [13]. We omit the key in the algorithm as it is less relevant to the simulation.

Identity replacement and failures. We briefly explain the attempts to separate PRFSGs and PRUs in [10, 18]. The main idea was to replace the oracle query S_d^Φ by the identity I, i.e., to consider $\mathcal{G}^I$ and then apply this PRU candidate on a Haar random state or the first half of the maximally entangled state $|\Omega\rangle$.

We consider the random-state input for intuition for now. Without an ancillary register, an oracle-independent unitary maps random states into random states. Also note that the oracle query S_d^Φ only perturbs a low-dimensional subspace spanned by $\left|x, 0^{d+1}\right\rangle$ and $|x, \phi_x, 1\rangle$, thus for random state input each query does not change the state much. Therefore, replacing the oracle by the identity does not change the final output state much.

It turns out that ancillary registers break precisely this argument. With a sufficiently large ancilla, an oracle-independent unitary can move the entire input state into the ancilla. For example, for the query register Q and the ancillary register A, it may map

$$|\psi\rangle|0\rangle_A \longmapsto \left|0^d, 0^{d+1}\right\rangle_Q |\psi\rangle_A \begin{cases} \xmapsto{S_d^\Phi} \left|0^d, \phi_{0^d}, 1\right\rangle_Q |\psi\rangle_A \\ \xmapsto{I} \left|0^d, 0^{d+1}\right\rangle_Q |\psi\rangle_A \end{cases}$$

where we omit the unrelated registers. Thus, replacing S_d^Φ by I can drastically change an intermediate state. This alone does not imply a large difference in the induced operation, since the discrepancy may eventually

disappear into the ancillary workspace. However, it breaks the original argument, as the two intermediate states remain far even after averaging over a random input.

Worse, there is a more severe counterexample: replacing the CHFS oracle by the identity may result in a completely different procedure. Fix $x \neq 0^d$. We use the CHFS oracle first on input 0^d and then on input x, uncomputing the input register after the queries. We then apply an arbitrarily chosen unitary V, controlled on whether the d-qubit state register in the ancilla is $\left|0^d\right\rangle$:

$$|\psi\rangle \left|0^{d+1}\right\rangle_A \begin{cases} \xrightarrow{S_d^\Phi} |\psi\rangle \left|\phi_{0^d}, 1\right\rangle_A \xrightarrow[\approx]{S_d^\Phi} |\psi\rangle \left|\phi_{0^d}, 1\right\rangle_A \xrightarrow[\approx]{C_{0^d}V} |\psi\rangle \left|\phi_{0^d}, 1\right\rangle_A \\ \xrightarrow{I} |\psi\rangle \left|0^d, 0\right\rangle_A \xrightarrow{I} |\psi\rangle \left|0^d, 0\right\rangle_A \xrightarrow{C_{0^d}V} V|\psi\rangle \left|0^d, 0\right\rangle_A \end{cases}$$

where in the above equation the first approximation is due to the almost orthogonality of $|\phi_{0^d}\rangle$ and $|\phi_x\rangle$ and the second is due to the almost orthogonality between $|\phi_{0^d}\rangle$ and $\left|0^d\right\rangle$. This gives a more serious obstruction: the purity is almost unaffected, while replacing the CHFS oracle by I can change the induced operation by an arbitrary unitary V. Note that the difference in the final qubit is also a problem.

Random CHFS replacement and derivatives. Our starting point is simple: we replace the original CHFS oracle S_d^Φ by another CHFS oracle $S_d^{\Phi'}$ defined by an independent family of random states $\Phi' = (|\phi'_x\rangle)_{x\in\{0,1\}^d}$. This choice circumvents the second counterexample, as the same near-orthogonality properties still hold for the new CHFS oracle. Still, the first problem persists: we need a new proof strategy to compare the two cases.

We use the following perspective. Consider the output state

$$|\rho_\Phi\rangle = (\mathcal{G}^{S_d^\Phi} \otimes I)\,|\Omega\rangle$$

where, for simplicity, we suppose that $\mathcal{G}^{S_d^\Phi}$ is implemented by a unitary circuit with no ancilla, and identify it with the resulting unitary. The general case will be discussed later. We then interpret this state as the output of the map

$$\text{(product space of states)} \longrightarrow \text{(space of output states)}, \qquad \Phi \longmapsto |\rho_\Phi\rangle\,.$$

This is obviously continuous, and moreover is a differentiable function! Here a derivative comes in, which exhibits the low-dimensionality of the CHFS oracle in a different form.

We use dots to denote differentials; e.g., $\dot\phi$ and $\dot S_\phi$ denote derivatives of ϕ and S_ϕ, respectively. For ease of exposition, we suppress tangent coordinates and other mathematical details, and also use dot notation for gradients. To see how the differential viewpoint captures the low-rank intuition, recall

$$S_\phi := I - \left|0^{d+1}\right\rangle\!\left\langle 0^{d+1}\right| - |\phi,1\rangle\langle\phi,1| + \left|0^{d+1}\right\rangle\!\left\langle \phi,1\right| + \left|\phi,1\right\rangle\!\left\langle 0^{d+1}\right|.$$

A direct calculation shows that its derivative is

$$\dot S_\phi = -\left|\dot\phi,1\right\rangle\!\left\langle\phi,1\right| - \left|\phi,1\right\rangle\!\left\langle\dot\phi,1\right| + \left|0^{d+1}\right\rangle\!\left\langle\dot\phi,1\right| + \left|\dot\phi,1\right\rangle\!\left\langle 0^{d+1}\right|$$

where the identity I vanishes. In particular, the derivative has a constant rank, $\operatorname{rank}(\dot S_\phi) = O(1)$. Accordingly, the derivative of the CHFS oracle becomes

$$\dot S_d^\Phi = \sum_{x\in\{0,1\}^d} |x\rangle\langle x| \otimes \dot S_{\phi_x},$$

which has low rank, $O(2^d)$, compared with the query register dimension $O(2^{2d})$.

Write a q-query oracle algorithm $\mathcal{G}^{S_d^\Phi}$ as

$$W^\Phi = U_q S_d^\Phi U_{q-1} S_d^\Phi \cdots U_1 S_d^\Phi U_0$$

by deferring all intermediate measurements. Differentiating it gives

$$\begin{aligned}\dot{W}^{\Phi} = U_q \dot{S}_d^{\Phi} U_{q-1} S_d^{\Phi} \cdots U_1 S_d^{\Phi} U_0 \\ + U_q S_d^{\Phi} U_{q-1} \dot{S}_d^{\Phi} \cdots U_1 S_d^{\Phi} U_0 \\ + \cdots + U_q S_d^{\Phi} U_{q-1} S_d^{\Phi} \cdots U_1 \dot{S}_d^{\Phi} U_0.\end{aligned}$$

Thus, the derivative of the oracle algorithm has rank at most an $O(q/2^d)$ fraction of the ambient dimension. We exploit this low relative rank to bound the differential of the output state obtained by applying the algorithm to half of the maximally entangled state $|\Omega\rangle$. Intuitively, this controls how much the output state $|\rho_\Phi\rangle$ changes as we perturb the oracle states Φ.

What we actually show is the following squared norm bound

$$\left\| |\dot{\rho}_\Phi\rangle\langle\rho_\Phi| + |\rho_\Phi\rangle\langle\dot{\rho}_\Phi| \right\|^2 \lesssim q^2.$$

Note that this is a derivative of the density matrix $|\rho_\Phi\rangle\langle\rho_\Phi| =: \rho_\Phi$ where $|\dot{\rho}_\Phi\rangle = (\dot{W}^{\Phi} \otimes I)\,|\Omega\rangle$. This turns out to suffice for our purpose. The Poincaré inequality then roughly says that

$$\mathop{\mathbb{E}}_{\Phi_d} \left\| \rho_\Phi - \mathop{\mathbb{E}}_{\Phi}[\rho_\Phi] \right\|_2^2 \leq \frac{\mathbb{E}_\Phi \left[\| \, |\dot{\rho}_\Phi\rangle\langle\rho_\Phi| + |\rho_\Phi\rangle\langle\dot{\rho}_\Phi| \, \|^2 \right]}{2^{d+1}-1} = O(q^2 2^{-d}).$$

Thus, the output state concentrates around the averaged state $\mathbb{E}_\Phi[\rho_\Phi]$, which can be effectively simulated.

Dealing with ancilla and non-unitarity. We now explain why this approach can handle ancillary qubits; it is convenient to discuss the non-unitary case at the same time. Recall that $W^{\Phi} = U_q S_d^{\Phi} U_{q-1} S_d^{\Phi} \cdots U_1 S_d^{\Phi} U_0$ represents the unitary part of the algorithm $\mathcal{G}^{S_d^{\Phi}}$. Let A, R, and B denote the algorithm register, the reference register of the maximally entangled state, and the ancillary register, respectively. We write

$$|\Xi_\Phi\rangle_{ABR} := (W^{\Phi} \otimes I_R)\big(|\Omega\rangle_{AR}\,|0\rangle_B\big), \qquad \rho_\Phi := \mathrm{Tr}_B\,|\Xi_\Phi\rangle\langle\Xi_\Phi|\,.$$

The squared norm that we need to bound is therefore

$$\left\| \mathrm{Tr}_B \left(\left|\dot{\Xi}_\Phi\right\rangle\!\left\langle\Xi_\Phi\right| + \left|\Xi_\Phi\right\rangle\!\left\langle\dot{\Xi}_\Phi\right| \right) \right\|_2^2.$$

Here $\left|\dot{\Xi}_\Phi\right\rangle$ is generated by the low-rank derivative $\dot{W}^{\Phi}$.

Suppose first that ρ_Φ is nearly pure, otherwise the purity test can distinguish $\mathcal{G}$ from random unitaries. We observe that this purity condition gives the decomposition

$$|\Xi_\Phi\rangle_{ABR} = \left|\Xi_\Phi^{\mathrm{id}}\right\rangle_{AR} \otimes |\chi_\Phi\rangle_B + |e_\Phi\rangle\,,$$

for a small error $|e_\Phi\rangle$. Define

$$|\eta_\Phi\rangle_{AR} := (I_{AR} \otimes \langle\chi_\Phi|_B)\left|\dot{\Xi}_\Phi\right\rangle_{ABR}.$$

Then the derivative of the reduced state decomposes as

$$|\eta_\Phi\rangle\!\langle\Xi_\Phi^{\mathrm{id}}| + |\Xi_\Phi^{\mathrm{id}}\rangle\!\langle\eta_\Phi| + E_\Phi, \qquad E_\Phi := \mathrm{Tr}_B\Big(\left|\dot{\Xi}_\Phi\right\rangle\!\left\langle e_\Phi\right| + \left|e_\Phi\right\rangle\!\left\langle\dot{\Xi}_\Phi\right| \Big)\,.$$

The first two terms have the same form as in the pure case, while the last term is small due to the smallness of $|e_\Phi\rangle$ together with our derivative bound.

2 Preliminaries

Notation. All functions and algorithms are parameterized by the security parameter $\lambda \in \mathbb{N}$. We write $[m] = \{1, \ldots, m\}$. For a distribution $\mathcal{D}$ and a finite set S, we write $x \leftarrow \mathcal{D}$ and $x \leftarrow S$ for sampling from $\mathcal{D}$ and uniform sampling from S, respectively.

We identify each quantum register with a finite-dimensional Hilbert space. We write $\mathsf{L}(\mathcal{H})$ for the space of linear operators on $\mathcal{H}$ and $\mathrm{Herm}(\mathcal{H})$ for its real vector space of Hermitian operators. On $\mathrm{Herm}(\mathcal{H})$, the Hilbert–Schmidt inner product is $(X, Y)_{\mathrm{HS}} = \mathrm{Tr}(XY)$.

A quantum state on a register A is represented by a density matrix $\rho \geq 0$ with $\mathrm{Tr}(\rho) = 1$. A pure state is represented by a unit vector $|\psi\rangle$, and we write $|\psi\rangle\langle\psi|$ for its density matrix.

A quantum channel $\mathcal{E}$ from A to A' is a completely positive, trace-preserving linear map; $\mathcal{E}(\rho)$ denotes its output on input ρ. We write $\mathsf{U}(A)$ for the unitaries on A. A unitary U induces the channel $\mathcal{U}(\rho) = U\rho U^\dagger$. The identity unitary and channel on A are denoted by I_A and $\mathcal{I}_A$, respectively.

2.1 Distances and tomography

For a vector v, we write $\|v\|$ for its Euclidean norm. For an operator X, the trace, Hilbert–Schmidt, and operator norms are

$$\|X\|_1 = \mathrm{Tr}\sqrt{X^\dagger X}, \qquad \|X\|_2 = \sqrt{\mathrm{Tr}(X^\dagger X)}, \qquad \|X\|_\infty = \sup_{\|v\|=1} \|Xv\|.$$

We use $\|X\|_2 \leq \|X\|_1$ and $\|YXZ\|_1 \leq \|Y\|_\infty \|X\|_1 \|Z\|_\infty$ for operators of compatible dimensions.

The trace distance between ρ and σ is $\|\rho - \sigma\|_1 / 2$. For any two-outcome measurement, the absolute difference in acceptance probabilities on ρ and σ is at most their trace distance, and some measurement achieves this bound. Applying the same channel to two states cannot increase their trace distance. For pure states $|\psi\rangle, |\phi\rangle$,

$$\frac{\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_1}{2} = \sqrt{1 - |\langle\psi|\phi\rangle|^2} = \frac{\||\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_2}{\sqrt{2}}.$$

For two (possibly unnormalized) states $|u\rangle, |e\rangle \in A \otimes B$, we also have the inequality $\|\mathrm{Tr}_B(|u\rangle\langle e|)\|_2 \leq \||u\rangle\| \, \||e\rangle\|$.

For channels $\mathcal{E}, \mathcal{F}$ from A to A', their diamond-norm distance is

$$\|\mathcal{E} - \mathcal{F}\|_\diamond = \sup_{C, \rho_{AC}} \|(\mathcal{E} \otimes \mathcal{I}_C)(\rho_{AC}) - (\mathcal{F} \otimes \mathcal{I}_C)(\rho_{AC})\|_1,$$

where the supremum is over all finite-dimensional C and all states ρ_{AC} on $A \otimes C$. Replacing gates in a quantum computation changes its output state, in trace norm, by at most the sum of the diamond-norm distances of the replaced gates.

We use the following process tomography result from [19].

Theorem 2.1. *Given oracle access to a unitary Z in dimension D and $\zeta, \delta \in (0, 1)$, an algorithm outputs a classical description of a unitary $\widetilde{Z}$ such that $\left\|\mathcal{Z} - \widetilde{\mathcal{Z}}\right\|_\diamond \leq \zeta$ with probability at least $1 - \delta$. It uses $O(D^2 \zeta^{-1} \log(2/\delta))$ queries and time polynomial in $D, 1/\zeta, \log(2/\delta)$.*

2.2 Haar randomness and concentration

Haar-random pure states and unitaries are sampled with respect to the normalized Haar measure in the relevant dimension. We sometimes write $\psi \leftarrow \mathcal{H}$ for sampling from the Haar distribution.

The following concentration bound follows from [29, 22].

Lemma 2.2. *Let $U_1, \ldots, U_m$ be independent Haar-random unitaries of dimensions $d_1, \ldots, d_m \geq d$. Let $g(U)$ be the acceptance probability of a quantum algorithm making at most $q \geq 1$ queries to $\bigoplus_{j=1}^m U_j$. Then, for every $t > 0$,*

$$\mathbb{P}[g(U) \geq \mathbb{E}\, g(U) + t] \leq \exp\left(-\frac{(d-2)t^2}{24q^2}\right).$$

The same bound holds for $\mathbb{P}[g(U) \leq \mathbb{E}\, g(U) - t]$.

We also use the Chernoff-Hoeffding bound.

Lemma 2.3. *Let $X_1, \ldots, X_r$ be independent Bernoulli variables with $\mathbb{E}\, X_i = p$. For $0 < p < a < 1$,*

$$\mathbb{P}\left[\frac{1}{r}\sum_{i=1}^r X_i \geq a\right] \leq \exp\left(-r\left[a\log\frac{a}{p} + (1-a)\log\frac{1-a}{1-p}\right]\right).$$

2.3 Choi states and purity tests

Let A have dimension d_A, and let R be an ancillary register of the same dimension. The maximally entangled state and the Choi state of a channel $\mathcal{E}$ on A are

$$|\Omega\rangle_{AR} := \frac{1}{\sqrt{d_A}}\sum_{a=1}^{d_A} |a\rangle_A |a\rangle_R, \qquad J(\mathcal{E}) := (\mathcal{E}\otimes\mathcal{I}_R)(|\Omega\rangle\langle\Omega|).$$

Thus $J(\mathcal{E})$ is prepared by applying $\mathcal{E}$ to the A register of $|\Omega\rangle_{AR}$. If $\mathcal{V}$ is induced by a unitary V, then $J(\mathcal{V})$ is the pure state with state vector $(V\otimes I_R)|\Omega\rangle$.

We use the swap test and the purity test.

Lemma 2.4 (Swap and purity tests)**.** *The swap test on $\rho\otimes\sigma$ passes with probability $(1+\mathrm{Tr}(\rho\sigma))/2$. On two independent copies of ρ, it is the purity test and fails with probability $(1-\mathrm{Tr}(\rho^2))/2$.*

We will use the following lemmas for almost pure states.

Lemma 2.5 (Purification near a pure state)**.** *Let $|\Xi\rangle_{AB}$ be a purification of ρ_A. If $\|\rho - |\psi\rangle\langle\psi|\|_1 \leq \varepsilon \leq 2$ for a pure state $|\psi\rangle_A$, then some unit vector $|\eta\rangle_B$ satisfies $\||\Xi\rangle - |\psi\rangle|\eta\rangle\| \leq \sqrt{\varepsilon}$.*

Proof. Let $w = (\langle\psi|\otimes I_B)|\Xi\rangle$ and $p = \|w\|^2 = \mathrm{Tr}(|\psi\rangle\langle\psi|\,\rho)$. The measurement $\{|\psi\rangle\langle\psi|, I - |\psi\rangle\langle\psi|\}$ accepts $|\psi\rangle\langle\psi|$ with probability one and ρ with probability p. Hence $1 - p \leq \|\rho - |\psi\rangle\langle\psi|\|_1/2 \leq \varepsilon/2$. If $p > 0$, let $|\eta\rangle = w/\sqrt{p}$; then $\||\Xi\rangle - |\psi\rangle|\eta\rangle\|^2 = 2 - 2\sqrt{p} \leq \varepsilon$. If $p = 0$, then $\varepsilon = 2$ and any unit vector $|\eta\rangle$ suffices. □

Lemma 2.6 (Trace norm near pure states)**.** *Let ρ, σ be density operators and P, Q rank-one projectors on the same space. If $\varepsilon_\rho = \|\rho - P\|_1$ and $\varepsilon_\sigma = \|\sigma - Q\|_1$, then $\|\rho - \sigma\|_1 \leq (1+\sqrt{2})(\varepsilon_\rho + \varepsilon_\sigma) + \sqrt{2}\,\|\rho - \sigma\|_2$.*

Proof. The triangle inequality and $\|P - Q\|_1 = \sqrt{2}\,\|P - Q\|_2$ give

$$\|\rho - \sigma\|_1 \leq \varepsilon_\rho + \varepsilon_\sigma + \sqrt{2}\,\|P - Q\|_2.$$

Also, $\|P - Q\|_2 \leq \varepsilon_\rho + \|\rho - \sigma\|_2 + \varepsilon_\sigma$. Combining the two inequalities proves the claim. □

For a channel $\mathcal{E}$ on A, define its Choi error and purity defect by

$$e(\mathcal{E}) := \inf_{V\in\mathsf{U}(A)} \|J(\mathcal{E}) - J(\mathcal{V})\|_1, \qquad \Delta(\mathcal{E}) := 1 - \mathrm{Tr}\big(J(\mathcal{E})^2\big).$$

Given two forward calls to $\mathcal{E}$, we can test the purity of its Choi state: prepare two copies of $J(\mathcal{E})$ and apply the purity test. It fails with probability $\Delta(\mathcal{E})/2$; for a unitary channel, it never fails. The following lemma relates this probability to the Choi error.

Lemma 2.7 (Purity and Choi error)**.** *Every channel $\mathcal{E}$ from A to itself satisfies*

$$\Delta(\mathcal{E}) \le e(\mathcal{E}) \le 5\sqrt{\Delta(\mathcal{E})}.$$

Proof. Let $\rho = J(\mathcal{E})$. Choose a unitary channel $\mathcal{V}$ attaining the infimum in the definition of $e(\mathcal{E})$. Since $J(\mathcal{V})$ is pure, the measurement $\{J(\mathcal{V}), I - J(\mathcal{V})\}$ accepts $J(\mathcal{V})$ with probability one. Hence $\mathrm{Tr}(J(\mathcal{V})\rho) \ge 1 - \|\rho - J(\mathcal{V})\|_1/2 = 1 - e(\mathcal{E})/2$. Since $\mathrm{Tr}(\rho^2) \ge \mathrm{Tr}(J(\mathcal{V})\rho)^2$, we have $\Delta(\mathcal{E}) \le e(\mathcal{E})$.

Conversely, let p be the largest eigenvalue of ρ and $|\gamma\rangle$ a corresponding eigenvector. Since $p \ge \mathrm{Tr}(\rho^2) = 1 - \Delta(\mathcal{E})$, $\|\rho - |\gamma\rangle\langle\gamma|\|_1 = 2(1-p) \le 2\Delta(\mathcal{E})$. As $\mathrm{Tr}_A \rho = I_R/d_A$ and partial trace cannot increase trace distance,

$$\|\mathrm{Tr}_A |\gamma\rangle\langle\gamma| - I_R/d_A\|_1 \le \||\gamma\rangle\langle\gamma| - \rho\|_1 \le 2\Delta(\mathcal{E}).$$

Write $|\gamma\rangle$ in Schmidt form and choose a unitary V such that

$$|\gamma\rangle = \sum_{j=1}^{d_A} \sqrt{s_j}\,|a_j\rangle_A\,|r_j\rangle_R\,, \qquad (V \otimes I_R)\,|\Omega\rangle = \frac{1}{\sqrt{d_A}} \sum_{j=1}^{d_A} |a_j\rangle_A\,|r_j\rangle_R\,.$$

Then

$$\||\gamma\rangle - (V \otimes I_R)\,|\Omega\rangle\|^2 = \sum_j \left(\sqrt{s_j} - \frac{1}{\sqrt{d_A}}\right)^2 \le \sum_j |s_j - 1/d_A| \le 2\Delta(\mathcal{E}).$$

The pure-state distance bound and the triangle inequality give $e(\mathcal{E}) \le 2\Delta(\mathcal{E}) + 2\sqrt{2\Delta(\mathcal{E})} \le 5\sqrt{\Delta(\mathcal{E})}$, where the last inequality uses $0 \le \Delta(\mathcal{E}) \le 1$. □

2.4 PRFSGs and PRUs

The main targets of this paper are the pseudorandom function-like state generators (PRFSGs) and pseudorandom unitaries (PRUs) [4, 21]. By default we use the definitions in the oracle world.

Definition 2.8 (PRFSG)**.** *Let κ, m, n be such that $\kappa, m = \omega(\log\lambda)$. We say an oracle QPT algorithm $\mathsf{Gen}^{\mathcal{O}}$ is a quantum-accessible, adaptively secure* pseudorandom function-like state generator *(PRFSG) if the following hold:*

- *$\mathsf{Gen}^{\mathcal{O}}$ maps $(k, x) \in \{0,1\}^{\kappa(\lambda)} \times \{0,1\}^{m(\lambda)}$ to a fresh $n(\lambda)$-qubit state in a new register;*
- *For every oracle QPT adversary $\mathcal{A}$,*

$$\left|\mathbb{P}_{k\leftarrow\{0,1\}^{\kappa(\lambda)}}[\mathcal{A}^{\mathcal{O},\mathsf{Gen}_k^{\mathcal{O}}(\cdot)} = 1] - \mathbb{P}_{G_{\mathrm{Haar}}}[\mathcal{A}^{\mathcal{O},G_{\mathrm{Haar}}(\cdot)} = 1]\right| \le \mathrm{negl}(\lambda)$$

where G_{Haar} samples an independent Haar state $|\psi_x\rangle$ for each x at the beginning of the game. On each query, it applies the isometry $|x\rangle_X \mapsto |x\rangle_X\,|\psi_x\rangle_Y$. The algorithm $\mathcal{A}$ is allowed to query the oracle coherently and adaptively.

For the pseudorandom unitaries, we allow the construction to be a channel, i.e., allow some errors as long as they are indistinguishable from random unitaries. In the main theorem, our adversary will query PRU candidates non-adaptively only in the forward direction.

Definition 2.9 (PRU)**.** *Let $n = \omega(\log\lambda)$. We say an oracle QPT algorithm $\mathcal{G}^{\mathcal{O}}$ is a* pseudorandom unitary *(PRU) if the following hold:*

- *For each $k \in \{0,1\}^\lambda$, $\mathcal{G}_k^{\mathcal{O}}$ is a channel from $n(\lambda)$ qubits to $n(\lambda)$ qubits.*
- *For every oracle QPT adversary $\mathcal{A}$,*

$$\left|\mathbb{P}_{k\leftarrow\{0,1\}^\lambda}[\mathcal{A}^{\mathcal{O},\mathcal{G}_k^{\mathcal{O}}} = 1] - \mathbb{P}_{U\leftarrow\mathsf{U}(2^{n(\lambda)})}[\mathcal{A}^{\mathcal{O},\mathcal{U}} = 1]\right| \le \mathrm{negl}(\lambda).$$

2.5 The quantum OR test

We use the unitary QPSPACE oracle defined as follows.

Definition 2.10. *The unitary oracle* QPSPACE *takes a quantum register A, a classical Turing machine M, a time bound t encoded in binary, and an abort qubit. If M outputs a valid unitary circuit on A within t steps, the oracle applies that circuit to A. Otherwise, it flips the abort qubit. The oracle preserves the description registers and acts coherently on superpositions of (M, t).*

Recall the quantum OR test [20, 9].

Lemma 2.11 (Quantum OR test)**.** *Let $\{M_i\}_{i=1}^N$ be two-outcome measurements on the same state ξ, and let $0 < a < 1/2$ and $b > 0$. There is a quantum OR test acting on one copy of ξ with the following guarantees. If some M_i accepts with probability at least $1 - a$, the test accepts with probability at least $(1 - a)^2/7$. If every M_i accepts with probability at most b, the test accepts with probability at most $4Nb$.*

We allow the measurements to make QPSPACE oracle queries. The following corollary is implicit in the previous works, see e.g., [13, Appendix A].

Corollary 2.12 (QPSPACE implementation)**.** *In Lemma 2.11, suppose that M_i is described by unitary QPT circuits with* QPSPACE *access that are generated uniformly from i and a common polynomial-length classical input. Suppose that ξ consists of polynomially many qubits and $\log N = \mathrm{poly}(\lambda)$. Then the quantum OR test has a uniform QPT implementation with access to* QPSPACE*.*

Proof. The unitary OR construction of [13, Appendix A] can be executed using QPSPACE. The proof of [9, Lemma 5.2] extends this implementation to binary measurements given uniformly by polynomial-size circuits with QPSPACE access. Since the family $\{M_i\}_{i=1}^N$ has such implementations, ξ has polynomially many qubits, and $\log N = \mathrm{poly}(\lambda)$, this result gives the claimed QPT implementation with access to QPSPACE. □

3 Hilbert-valued Functions and Poincaré Inequalities

This section introduces the mathematical background on the basic differential analysis of functions from product spheres to Hilbert spaces, and the Poincaré inequality for those functions.

Throughout this paper, every function differentiated is C^1[1]. All derivatives are taken with respect to real parameters and, for vector- or matrix-valued maps, are computed entrywise in fixed bases.

3.1 Functions on spheres and their derivatives

We first review the notions and results about the functions on spheres. Let $\mathbb{S}(\mathbb{C}^D) := \{\psi \in \mathbb{C}^D : \|\psi\| = 1\}$ be the complex unit sphere. We regard it as the $(2D-1)$-dimensional real unit sphere in $\mathbb{R}^{2D}$, with real inner product $\mathrm{Re}\langle u, v\rangle$. Its tangent space at $\psi \in \mathbb{S}(\mathbb{C}^D)$ is

$$T_\psi \mathbb{S}(\mathbb{C}^D) = \{v \in \mathbb{C}^D : \mathrm{Re}\langle \psi, v\rangle = 0\}.$$

If $e_1, \ldots, e_{D-1}$ is a complex orthonormal basis of the complex orthogonal complement $\psi^\perp := \{v \in \mathbb{C}^D : \langle \psi, v\rangle = 0\}$ of ψ, then

$$\{i\psi, e_1, ie_1, \ldots, e_{D-1}, ie_{D-1}\} \tag{1}$$

forms a real orthonormal basis of $T_\psi \mathbb{S}(\mathbb{C}^D)$.

[1]The number 1 here simultaneously denotes C^1 and the footnote index. A C^1 function is a differentiable function with a continuous derivative.

We consider a function $f : \mathbb{S}(\mathbb{C}^D) \to \mathbb{R}$. At $\psi \in \mathbb{S}(\mathbb{C}^D)$, the derivative of f along a direction $\phi \in T_\psi \mathbb{S}(\mathbb{C}^D)$ is defined by

$$\frac{\partial f}{\partial \phi}(\psi) := \frac{\mathrm{d}}{\mathrm{d}t} f(\psi(t))\Big|_{t=0},$$

where $\psi(t)$ is any differentiable curve on $\mathbb{S}(\mathbb{C}^D)$ with $\psi(0) = \psi$ and $\frac{\mathrm{d}}{\mathrm{d}t}\psi(t)\big|_{t=0} = \phi$. This quantity depends only on ϕ, not on the choice of the curve.

For the basis chosen in Eq. (1), the gradient of f at ψ is represented by

$$\nabla f(\psi) = \left(\frac{\partial f}{\partial i\psi}(\psi), \frac{\partial f}{\partial e_1}(\psi), \frac{\partial f}{\partial i e_1}(\psi), \ldots, \frac{\partial f}{\partial e_{D-1}}(\psi), \frac{\partial f}{\partial i e_{D-1}}(\psi) \right).$$

We are mainly interested in the squared gradient norm

$$\|\nabla f(\psi)\|^2 = \left| \frac{\partial f}{\partial i\psi}(\psi) \right|^2 + \sum_{j=1}^{D-1} \left(\left| \frac{\partial f}{\partial e_j}(\psi) \right|^2 + \left| \frac{\partial f}{\partial i e_j}(\psi) \right|^2 \right),$$

which is independent of the choice of the real orthonormal tangent basis.

We recall the Poincaré inequality on the complex sphere.

Theorem 3.1 (Poincaré inequality on the complex sphere)**.** *Let $f : \mathbb{S}(\mathbb{C}^D) \to \mathbb{R}$ be a C^1 function, and let ψ be uniformly distributed on $\mathbb{S}(\mathbb{C}^D)$. Then*

$$\mathbb{E}_\psi \left| f(\psi) - \mathbb{E}_\psi f(\psi) \right|^2 = \mathrm{Var}(f(\psi)) \le \frac{1}{2D-1} \mathbb{E}_\psi \|\nabla f(\psi)\|^2.$$

3.2 Hilbert-valued functions on complex spheres

We next extend the preceding definitions to functions taking values in a finite-dimensional real Hilbert space $\mathcal{K}$. Recall that a finite-dimensional real Hilbert space is a real inner-product space; we write $\langle \cdot, \cdot \rangle_{\mathcal{K}}$ for the inner product of $\mathcal{K}$ and $\| \cdot \|_{\mathcal{K}}$ for the induced norm.

Let $b_1, \ldots, b_m$ be an orthonormal basis of $\mathcal{K}$. Then, for every $x, y \in \mathcal{K}$,

$$x = \sum_{r=1}^{m} \langle b_r, x \rangle_{\mathcal{K}} b_r, \quad \langle x, y \rangle_{\mathcal{K}} = \sum_{r=1}^{m} \langle b_r, x \rangle_{\mathcal{K}} \langle b_r, y \rangle_{\mathcal{K}}, \quad \|x\|_{\mathcal{K}}^2 = \sum_{r=1}^{m} \left| \langle b_r, x \rangle_{\mathcal{K}} \right|^2.$$

A map $F : \mathbb{S}(\mathbb{C}^D) \to \mathcal{K}$ is said to be C^1 if its coordinates in an orthonormal basis are C^1. Expectations of $\mathcal{K}$-valued random variables are taken coordinatewise.

For $\phi \in T_\psi \mathbb{S}(\mathbb{C}^D)$, define its directional derivative by

$$\frac{\partial F}{\partial \phi}(\psi) := \frac{\mathrm{d}}{\mathrm{d}t} F(\psi(t))\Big|_{t=0},$$

where $\psi(t)$ is any differentiable curve on $\mathbb{S}(\mathbb{C}^D)$ with $\psi(0) = \psi$ and $\frac{\mathrm{d}}{\mathrm{d}t}\psi(t)\big|_{t=0} = \phi$. The coordinatewise differentiation is

$$\left\langle b_r, \frac{\partial F}{\partial \phi}(\psi) \right\rangle_{\mathcal{K}} = \frac{\partial F_r}{\partial \phi}(\psi) \qquad \text{where} \quad F_r(\psi) := \langle b_r, F(\psi) \rangle_{\mathcal{K}}.$$

For the real orthonormal tangent basis in Eq. (1), write

$$\nabla F(\psi) = \left(\frac{\partial F}{\partial i\psi}(\psi), \frac{\partial F}{\partial e_1}(\psi), \frac{\partial F}{\partial i e_1}(\psi), \ldots, \frac{\partial F}{\partial e_{D-1}}(\psi), \frac{\partial F}{\partial i e_{D-1}}(\psi) \right)$$

and define the squared gradient norm[2] by

$$\|\nabla F(\psi)\|_{\mathcal{K}}^2 := \left\|\frac{\partial F}{\partial i\psi}(\psi)\right\|_{\mathcal{K}}^2 + \sum_{j=1}^{D-1}\left(\left\|\frac{\partial F}{\partial e_j}(\psi)\right\|_{\mathcal{K}}^2 + \left\|\frac{\partial F}{\partial ie_j}(\psi)\right\|_{\mathcal{K}}^2\right).$$

Again, the gradient norm is independent of the choice of basis.

For a $\mathcal{K}$-valued random variable X, we define its variance by

$$\mathrm{Var}(X) := \mathbb{E}\|X - \mathbb{E}X\|_{\mathcal{K}}^2.$$

The Poincaré inequality can be naturally extended to this setting.

Corollary 3.2 (Hilbert-valued Poincaré inequality)**.** *Let $F : \mathbb{S}(\mathbb{C}^D) \to \mathcal{K}$ be a C^1 function, and let ψ be uniformly distributed on $\mathbb{S}(\mathbb{C}^D)$. Then*

$$\mathbb{E}_\psi\|F(\psi) - \mathbb{E}_\psi F(\psi)\|_{\mathcal{K}}^2 = \mathrm{Var}(F(\psi)) \le \frac{1}{2D-1}\mathbb{E}_\psi\|\nabla F(\psi)\|_{\mathcal{K}}^2.$$

Proof. Let $b_1, \dots, b_m$ be an orthonormal basis of $\mathcal{K}$. Write $F_r(\psi) := \langle b_r, F(\psi)\rangle_{\mathcal{K}}$, which is a real-valued function on $\mathbb{S}(\mathbb{C}^D)$. By Parseval's identity,

$$\mathbb{E}_\psi\|F(\psi) - \mathbb{E}_\psi F(\psi)\|_{\mathcal{K}}^2 = \sum_{r=1}^m \mathbb{E}_\psi\left|F_r(\psi) - \mathbb{E}_\psi F_r(\psi)\right|^2$$

Applying Theorem 3.1 to each F_r gives

$$\mathbb{E}_\psi\|F(\psi) - \mathbb{E}_\psi F(\psi)\|_{\mathcal{K}}^2 \le \frac{1}{2D-1}\sum_{r=1}^m \mathbb{E}_\psi\|\nabla F_r(\psi)\|^2. \tag{2}$$

For any real orthonormal tangent basis $\{v_a\}_a$, Parseval's identity gives

$$\sum_{r=1}^m\left|\frac{\partial F_r}{\partial v_a}(\psi)\right|^2 = \sum_{r=1}^m\left|\left\langle b_r, \frac{\partial F}{\partial v_a}(\psi)\right\rangle_{\mathcal{K}}\right|^2 = \left\|\frac{\partial F}{\partial v_a}(\psi)\right\|_{\mathcal{K}}^2.$$

Therefore, Summing the preceding identity over a and substituting into Eq. (2) proves the claim. □

3.3 Hilbert-valued functions on products of complex spheres

We now extend the preceding results to products of complex spheres. Fix the dimensions $D_1, \dots, D_L \ge 2$ and define the product sphere

$$\mathcal{M} := \prod_{j=1}^L \mathbb{S}(\mathbb{C}^{D_j}),$$

equipped with the product of the uniform measures on the individual spheres. We write a point in $\mathcal{M}$ as $\Psi = (\psi_1, \dots, \psi_L)$.

For each j, let $\{v_{j,a}\}_{a=1}^{2D_j-1}$ be a real orthonormal basis of $T_{\psi_j}\mathbb{S}(\mathbb{C}^{D_j})$. For a function $F : \mathcal{M} \to \mathcal{K}$, let $\partial_{j,a}F(\Psi)$ denote the directional derivative obtained by varying only ψ_j in the direction $v_{j,a}$. We write

$$\nabla_j F(\Psi) := \left(\partial_{j,1}F(\Psi), \dots, \partial_{j,2D_j-1}F(\Psi)\right)$$

and define

$$\|\nabla_j F(\Psi)\|_{\mathcal{K}}^2 := \sum_{a=1}^{2D_j-1}\|\partial_{j,a}F(\Psi)\|_{\mathcal{K}}^2.$$

As before, this quantity is independent of the choice of the real orthonormal tangent basis.

We first establish the Poincaré inequality for real-valued functions on the product space adapting the argument of [6, Proposition 4.3.1].

[2]This is the squared Hilbert-Schmidt norm of the differential and is denoted by $\|\mathrm{d}F_\psi\|_{\mathrm{HS}}^2$ in the literature.

Lemma 3.3 (Poincaré inequality on products of complex spheres)**.** *Let $f : \mathcal{M} \to \mathbb{R}$ be a C^1 function, where $\Psi = (\psi_1, \dots, \psi_L)$ is distributed according to the product of the uniform measures on $\mathbb{S}(\mathbb{C}^{D_1}), \dots, \mathbb{S}(\mathbb{C}^{D_L})$. Then*

$$\mathrm{Var}(f(\Psi)) \le \sum_{j=1}^{L} \frac{1}{2D_j - 1} \mathbb{E}_\Psi \|\nabla_j f(\Psi)\|^2. \tag{3}$$

Proof. We first consider the case $L = 2$, and write $f = f(\psi_1, \psi_2)$. Define $g(\psi_1) := \mathbb{E}_{\psi_2} f(\psi_1, \psi_2)$. The law of total variance gives

$$\mathrm{Var}_{\psi_1,\psi_2}(f) = \mathbb{E}_{\psi_1} \mathrm{Var}_{\psi_2}(f \mid \psi_1) + \mathrm{Var}_{\psi_1}(g). \tag{4}$$

For every fixed ψ_1, applying Theorem 3.1 to the function $\psi_2 \mapsto f(\psi_1, \psi_2)$ gives

$$\mathbb{E}_{\psi_1} \mathrm{Var}_{\psi_2}(f \mid \psi_1) \le \frac{1}{2D_2 - 1} \mathbb{E}_{\psi_1,\psi_2} \|\nabla_2 f(\psi_1, \psi_2)\|^2. \tag{5}$$

On the other hand, applying Theorem 3.1 to g gives

$$\mathrm{Var}_{\psi_1}(g) \le \frac{1}{2D_1 - 1} \mathbb{E}_{\psi_1} \|\nabla g(\psi_1)\|^2. \tag{6}$$

For every tangent direction v at ψ_1, differentiation commutes with the expectation over ψ_2, so

$$\frac{\partial g}{\partial v}(\psi_1) = \mathbb{E}_{\psi_2} \frac{\partial f}{\partial v}(\psi_1, \psi_2).$$

Therefore, Jensen's inequality gives

$$\|\nabla g(\psi_1)\|^2 \le \mathbb{E}_{\psi_2} \|\nabla_1 f(\psi_1, \psi_2)\|^2.$$

Combining this with Eqs. (4) to (6), we obtain

$$\mathrm{Var}_{\psi_1,\psi_2}(f) \le \frac{1}{2D_1 - 1} \mathbb{E}_{\psi_1,\psi_2} \|\nabla_1 f\|^2 + \frac{1}{2D_2 - 1} \mathbb{E}_{\psi_1,\psi_2} \|\nabla_2 f\|^2.$$

Iterating the same argument over the L coordinates proves Eq. (3). □

The same inequality extends to Hilbert-valued functions.

Theorem 3.4 (Hilbert-valued product-sphere Poincaré inequality)**.** *For every C^1 function $F : \mathcal{M} \to \mathcal{K}$,*

$$\mathbb{E}_\Psi \|F(\Psi) - \mathbb{E}_\Psi F(\Psi)\|_{\mathcal{K}}^2 \le \sum_{j=1}^{L} \frac{1}{2D_j - 1} \mathbb{E}_\Psi \|\nabla_j F(\Psi)\|_{\mathcal{K}}^2. \tag{7}$$

Proof. We use the variance-decomposition argument in the proof of Lemma 3.3, applying Corollary 3.2 in place of Theorem 3.1.

For $L = 2$, the Hilbert-valued variance satisfies

$$\mathrm{Var}_{\psi_1,\psi_2}(F) = \mathbb{E}_{\psi_1} \mathrm{Var}_{\psi_2}(F \mid \psi_1) + \mathrm{Var}_{\psi_1}(\mathbb{E}_{\psi_2} F).$$

Apply Corollary 3.2 to $\psi_2 \mapsto F(\psi_1, \psi_2)$ for each fixed ψ_1, and to $\psi_1 \mapsto \mathbb{E}_{\psi_2} F(\psi_1, \psi_2)$. Differentiation commutes with expectation, and Jensen's inequality gives

$$\|\nabla_1 \mathbb{E}_{\psi_2} F\|_{\mathcal{K}}^2 \le \mathbb{E}_{\psi_2} \|\nabla_1 F\|_{\mathcal{K}}^2.$$

Combining these bounds proves Eq. (7) for $L = 2$. Iterating the same argument proves it for arbitrary L. □

4 CHFS Oracle and Derivatives

We define the unitarized, or swap, common Haar-random function-like state (CHFS) oracle and study its derivatives in the context of oracle algorithms. We allow controlled access to the oracle by default and suppress the control qubit throughout.

We recall the Leibniz rules for any directional derivative ∂. Differentiation of vector- and operator-valued functions is understood entrywise, and

$$\partial |u\rangle\langle v| = |\partial u\rangle\langle v| + |u\rangle\langle \partial v|, \qquad \partial(AB) = (\partial A)B + A(\partial B), \partial(A^\dagger) = (\partial A)^\dagger,$$

where $|\partial u\rangle = \partial |u\rangle$. In particular, if $|\xi\rangle_{AB}$ is a differentiable purification of $\rho = \mathrm{Tr}_B |\xi\rangle\langle\xi|$, then $\partial\rho = \mathrm{Tr}_B (|\partial\xi\rangle\langle\xi| + |\xi\rangle\langle\partial\xi|)$.

4.1 The CHFS oracle

The CHFS oracle is parameterized by a family of states $\Phi = (\Phi_d)_{d\in\mathbb{N}}$. For every input $x \in \{0,1\}^d$ of length $d \in \mathbb{N}$, we assign a d-qubit pure state. Formally, let

$$\Phi_d = \left(|\phi_x\rangle\right)_{x\in\{0,1\}^d}, \qquad \phi_x \in \mathbb{S}(\mathbb{C}^{2^d}),$$

where all states $|\phi_x\rangle$ are sampled independently from Haar measure as an initialization, and are fixed throughout the oracle execution. For $\phi \in \mathbb{S}(\mathbb{C}^{2^d})$, we define the swap unitary associated with it by

$$S_\phi := I - \left|0^{d+1}\right\rangle\left\langle 0^{d+1}\right| - |\phi,1\rangle\langle\phi,1| + \left|0^{d+1}\right\rangle\left\langle\phi,1\right| + \left|\phi,1\right\rangle\left\langle 0^{d+1}\right| \tag{8}$$

The CHFS oracle S^Φ parameterized by $\Phi = (\Phi_d)_{d\in\mathbb{N}}$ consists of the family of oracles $(S_d^\Phi)_{d\in\mathbb{N}}$ where S_d^Φ acts on $(2d+1)$ qubits by

$$S_d^\Phi := \sum_{x\in\{0,1\}^d} |x\rangle\langle x| \otimes S_{\phi_x}.$$

We regard S_d^Φ as an operator-valued function of the state tuple

$$\Phi = (\phi_x)_{x\in\{0,1\}^d} \in \prod_{x\in\{0,1\}^d} \mathbb{S}(\mathbb{C}^{2^d}).$$

Accordingly, the tangent space becomes

$$T_\Phi\left(\prod_{x\in\{0,1\}^d} \mathbb{S}(\mathbb{C}^{2^d})\right) = \bigoplus_{x\in\{0,1\}^d} T_{\phi_x}\mathbb{S}(\mathbb{C}^{2^d}).$$

More precisely, for $\delta \in T_{\phi_x}\mathbb{S}(\mathbb{C}^{2^d})$, let $\phi_x(s)$ be any differentiable curve on $\mathbb{S}(\mathbb{C}^{2^d})$ such that $\phi_x(0) = \phi_x$ and $\frac{\mathrm{d}}{\mathrm{d}s}\phi_x(s)\big|_{s=0} = \delta$. Let $\Phi(s)$ be obtained from Φ by replacing only ϕ_x with $\phi_x(s)$. We define

$$\partial_{x,\delta} S_d^\Phi := \frac{\mathrm{d}}{\mathrm{d}s} S_d^{\Phi(s)}\Big|_{s=0}.$$

Note that the derivative $\partial_{x,\delta} S_d^\Phi$ is supported only on the input-x block since

$$S_d^\Phi = \sum_{y\in\{0,1\}^d} |y\rangle\langle y| \otimes S_{\phi_y}.$$

Note also that because of the coherent access, the phase of each state vector $|\phi\rangle$ must be considered. The phase direction is $i\phi$ and will contribute to the derivative.

4.2 CHFS oracle derivatives

This section studies the derivatives of the CHFS oracles and presents two inequalities. This section focuses on a fixed input length d. For every $x \in \{0,1\}^d$, we define for simplicity that

$$|A_x\rangle := |x\rangle \left|0^{d+1}\right\rangle, \quad |B_x\rangle := |x\rangle |\phi_x, 1\rangle, \quad |u_x\rangle := \frac{|B_x\rangle - |A_x\rangle}{\sqrt{2}}$$

so that

$$S_d^\Phi = I - 2 \sum_{x\in\{0,1\}^d} |u_x\rangle\langle u_x| . \tag{9}$$

For each x, choose a complex orthonormal basis $e_{x,1}, \ldots, e_{x,2^d-1}$ of $\phi_x^\perp$, which gives a real orthonormal basis

$$i\phi_x, e_{x,1}, ie_{x,1}, \ldots, e_{x,2^d-1}, ie_{x,2^d-1}$$

of $T_{\phi_x}\mathbb{S}(\mathbb{C}^{2^d})$ as in Eq. (1). We define the basis vectors

$$|t_{x,j}\rangle := |x\rangle |e_{x,j}\rangle |1\rangle , \qquad j \in [2^d - 1].$$

The CHFS oracle S_d^Φ can be regarded as a function on Φ, so that on ϕ_x. The following lemma gives the partial derivatives of S_d^Φ with respect to $T_{\phi_x}\mathbb{S}(\mathbb{C}^{2^d})$.

Lemma 4.1. *For every $x \in \{0,1\}^d$ and $j \in [2^d - 1]$,*

$$\begin{aligned}
\partial_{x,e_{x,j}} S_d^\Phi &= -\sqrt{2}\left(|u_x\rangle\langle t_{x,j}| + |t_{x,j}\rangle\langle u_x|\right), \\
\partial_{x,ie_{x,j}} S_d^\Phi &= i\sqrt{2}\left(|u_x\rangle\langle t_{x,j}| - |t_{x,j}\rangle\langle u_x|\right), \\
\partial_{x,i\phi_x} S_d^\Phi &= i\left(|B_x\rangle\langle A_x| - |A_x\rangle\langle B_x|\right).
\end{aligned}$$

Each derivative is zero outside the input-x block.

Proof. Fix x and suppress the subscript x. Since $|A\rangle$ is independent of ϕ, while $|B\rangle = |x\rangle|\phi, 1\rangle$, we have $\partial_{x,\delta}|u\rangle = \frac{1}{\sqrt{2}}|x\rangle|\delta, 1\rangle$ for every $\delta \in T_\phi\mathbb{S}(\mathbb{C}^{2^d})$. By Eq. (9),

$$\partial_{x,\delta} S_d^\Phi = -2\left(|\partial_{x,\delta}u\rangle\langle u| + |u\rangle\langle\partial_{x,\delta}u|\right) = -\sqrt{2}\left(|x,\delta,1\rangle\langle u| + |u\rangle\langle x,\delta,1|\right).$$

Plugging $\delta = e_j$ and $\delta = ie_j$ give the first and second derivatives. Finally, $\partial_{x,i\phi}|u\rangle = i|B\rangle/\sqrt{2}$, and hence

$$\partial_{x,i\phi} S_d^\Phi = -i\sqrt{2}\left(|B\rangle\langle u| - |u\rangle\langle B|\right) = i\left(|B\rangle\langle A| - |A\rangle\langle B|\right),$$

where we used $|u\rangle = (|B\rangle - |A\rangle)/\sqrt{2}$. Each derivative is supported only on the input-x block because only the x-th summand of S_d^Φ depends on ϕ_x. □

For each x, we define the relevant space for the CHFS oracles

$$E_x := \operatorname{span}\{A_x, B_x\}, \qquad T_x := \operatorname{span}\{t_{x,1}, \ldots, t_{x,2^d-1}\},$$

which are orthogonal to each other. Let Π_{E_x} and Π_{T_x} denote the corresponding orthogonal projections. Note that

$$\Pi_{E_x} + \Pi_{T_x} \le |x\rangle\langle x| \otimes I, \qquad \sum_{x\in\{0,1\}^d} \Pi_{E_x} + \sum_{x\in\{0,1\}^d} \Pi_{T_x} \le I$$

so that the space $E_x \oplus T_x$ are mutually orthogonal for distinct x.

Observe that they are relatively low rank compared to the ambient space. This gives the following upper bounds of the CHFS derivatives.

Lemma 4.2 (Two-sided CHFS derivative sum)**.** *Let P, Q be orthogonal projections of rank at most r on the CHFS query register together with an arbitrary ancillary register. Then*

$$\sum_{x\in\{0,1\}^d}\sum_{\delta}\left\|Q\left(\partial_{x,\delta}S_d^{\Phi}\otimes I\right)P\right\|_2^2\le 9r,$$

where, for each x, δ ranges over the basis $\{i\phi_x, e_{x,1}, ie_{x,1},\dots,e_{x,2^d-1}, ie_{x,2^d-1}\}$.

Proof. We suppress the subscript x in this proof; the distinct x's are only considered when we use the summation over $x\in\{0,1\}^d$. By Lemma 4.1 and the parallelogram identity, we have

$$\left\|Q\left(\partial_{e_j}S_d^{\Phi}\otimes I\right)P\right\|_2^2+\left\|Q\left(\partial_{ie_j}S_d^{\Phi}\otimes I\right)P\right\|_2^2=4\left\|Q\left(|u\rangle\langle t_j|\otimes I\right)P\right\|_2^2+4\left\|Q\left(|t_j\rangle\langle u|\otimes I\right)P\right\|_2^2. \tag{10}$$

Since $0\le Q\le I$, we have

$$\begin{aligned}\sum_{j=1}^{2^d-1}\left\|Q\left(|u\rangle\langle t_j|\otimes I\right)P\right\|_2^2&=\sum_{j=1}^{2^d-1}\operatorname{Tr}\left(P\left(|t_j\rangle\langle u|\otimes I\right)Q\left(|u\rangle\langle t_j|\otimes I\right)P\right)\\&\le\sum_{j=1}^{2^d-1}\operatorname{Tr}\left(P\left(|t_j\rangle\langle t_j|\otimes I\right)P\right)=\operatorname{Tr}\left(P(\Pi_T\otimes I)P\right).\end{aligned}$$

Similarly, by invariance of the Hilbert–Schmidt norm under adjoint and the same calculation with P and Q interchanged,

$$\sum_{j=1}^{2^d-1}\left\|Q\left(|t_j\rangle\langle u|\otimes I\right)P\right\|_2^2=\sum_{j=1}^{2^d-1}\left\|P\left(|u\rangle\langle t_j|\otimes I\right)Q\right\|_2^2\le\operatorname{Tr}\left(Q(\Pi_T\otimes I)Q\right).$$

Therefore, summing Eq. (10) over j gives, for each x,

$$\sum_{j=1}^{2^d-1}\left(\left\|Q\left(\partial_{e_j}S_d^{\Phi}\otimes I\right)P\right\|_2^2+\left\|Q\left(\partial_{ie_j}S_d^{\Phi}\otimes I\right)P\right\|_2^2\right)\le 4\operatorname{Tr}\left(P(\Pi_T\otimes I)P\right)+4\operatorname{Tr}\left(Q(\Pi_T\otimes I)Q\right).$$

For the phase direction, Lemma 4.1 gives $\left(\partial_{i\phi}S_d^{\Phi}\right)^{\dagger}\left(\partial_{i\phi}S_d^{\Phi}\right)=\Pi_E$ and hence,

$$\left\|Q\left(\partial_{i\phi}S_d^{\Phi}\otimes I\right)P\right\|_2^2\le\operatorname{Tr}\left(P(\Pi_E\otimes I)P\right).$$

Restoring the subscript x and summing over x, the mutual orthogonality of the spaces $E_x\oplus T_x$ gives

$$\sum_{x\in\{0,1\}^d}\sum_{\delta}\left\|Q\left(\partial_{x,\delta}S_d^{\Phi}\otimes I\right)P\right\|_2^2\le 5\operatorname{Tr}P+4\operatorname{Tr}Q\le 9r,$$

where we used $\sum_x\Pi_{E_x}\le I$, $\sum_x\Pi_{T_x}\le I$, and $\operatorname{Tr}(R\Pi R)\le\operatorname{Tr}R$ for orthogonal projections R and Π. The final inequality is due to the rank condition. □

Lemma 4.3 (One-sided CHFS derivative sum)**.** *Let I be any register. For every vector ξ on the CHFS query register together with I,*

$$\sum_{x\in\{0,1\}^d}\sum_{\delta}\left\|\left(\partial_{x,\delta}S_d^{\Phi}\otimes I_I\right)\xi\right\|^2\le 5\cdot 2^d\left\|\xi\right\|^2,$$

where, for each x, δ ranges over the basis $\{i\phi_x, e_{x,1}, ie_{x,1},\dots,e_{x,2^d-1}, ie_{x,2^d-1}\}$.

Proof. We suppress the subscript x in this proof; the distinct x's are only considered when we use the summation over $x \in \{0,1\}^d$. By Lemma 4.1, we have

$$\left(\partial_{e_j} S_d^\Phi\right)^\dagger \left(\partial_{e_j} S_d^\Phi\right) + \left(\partial_{ie_j} S_d^\Phi\right)^\dagger \left(\partial_{ie_j} S_d^\Phi\right) = 4\left(|u\rangle\langle u| + |t_j\rangle\langle t_j|\right).$$

Also, for the phase direction, we have $\left(\partial_{i\phi} S_d^\Phi\right)^\dagger \left(\partial_{i\phi} S_d^\Phi\right) = \Pi_E$. Therefore, summing over the tangent basis gives, for each x,

$$\sum_\delta \left(\partial_\delta S_d^\Phi\right)^\dagger \left(\partial_\delta S_d^\Phi\right) = 4(2^d - 1)\,|u\rangle\langle u| + 4\Pi_T + \Pi_E \le (4 \cdot 2^d - 3)\Pi_E + 4\Pi_T,$$

where we used $|u\rangle\langle u| \le \Pi_E$.

Summing over all x, the mutual orthogonality of the spaces $E_x \oplus T_x$ gives

$$\sum_{x\in\{0,1\}^d} \sum_\delta \left(\partial_{x,\delta} S_d^\Phi\right)^\dagger \left(\partial_{x,\delta} S_d^\Phi\right) \le 5 \cdot 2^d I.$$

Hence, expanding the left hand side becomes

$$\sum_{x\in\{0,1\}^d} \sum_\delta \left\|\left(\partial_{x,\delta} S_d^\Phi \otimes I_I\right)\xi\right\|^2 = \left\langle \xi, \left(\sum_{x\in\{0,1\}^d} \sum_\delta \left(\partial_{x,\delta} S_d^\Phi\right)^\dagger \left(\partial_{x,\delta} S_d^\Phi\right) \otimes I_I \right) \xi \right\rangle$$

which is bounded above by $5 \cdot 2^d \|\xi\|^2$, proving the lemma. □

4.3 Oracle algorithm derivatives

This section extends the derivative to the oracle algorithms. We assume that in each query, algorithms fix the input length classically. Given the controlled access to the oracle, this can be done by, e.g., querying $S_1^\Phi, \ldots, S_{d_{\max}}^\Phi$ instead of a single coherent query for the maximum input size $d_{\max}$ at the cost of $d_{\max}$ query number blow up. See also [14]. In our main purpose, the algorithm runs in polynomial time so that $d_{\max}$ is at most polynomial.

Oracle algorithm as a function. Consider a PRU candidate algorithm $\mathcal{A}$ having oracle access to S^Φ. Let A be the input/output register of $\mathcal{A}$, and let B be its ancillary register. After purifying intermediate measurements, the algorithm can be written as

$$W_\Phi = U_q S_{d_q}^\Phi U_{q-1} \cdots U_1 S_{d_1}^\Phi U_0, \tag{11}$$

acting on the register AB, where q is the total number of oracle queries, d_t is the classically fixed input length of the t-th query, and every unitary U_t is independent of Φ.

We now fix a query input length d and study the dependence of the algorithm only on Φ_d, while keeping all state families $\Phi_{d'}$ for $d' \ne d$ fixed. Let q_d be the number of query positions t satisfying $d_t = d$. By absorbing all gates and queries at the other input lengths into the inter-query unitaries and relabeling the length-d query positions, we can write

$$W_\Phi = U_{q_d} S_d^\Phi U_{q_d-1} S_d^\Phi \cdots U_1 S_d^\Phi U_0, \tag{12}$$

where the unitaries U_t may depend on the fixed state families $\Phi_{d'}$ for $d' \ne d$, but are independent of Φ_d. Here we reuse the same U for the query-irrelevant unitaries because they are less important in the analysis below. For the analysis below, we also drop the subscript and write Φ and S^Φ to denote Φ_d and S_d^Φ, respectively.

We consider the application of $\mathcal{A}$ to the maximally entangled state $|\Omega\rangle_{AR}$, where $R \cong A$ is a reference register not touched by the algorithm. The final states before and after tracing out B are

$$|\Xi_\Phi\rangle_{ABR} := (W_\Phi \otimes I_R)\left(|\Omega\rangle_{AR}\,|0\rangle_B\right), \qquad \rho_\Phi := \mathrm{Tr}_B\,|\Xi_\Phi\rangle\langle\Xi_\Phi|\,. \tag{13}$$

We regard this application of the oracle algorithm as a function from the product state space defining the length-d CHFS oracle to the space of Hermitian operators on $A \otimes R$. More formally, we consider the smooth function

$$F_d^{\mathcal{A}} : \Phi \longmapsto \rho_\Phi \quad \text{from} \quad \prod_{x \in \{0,1\}^d} \mathbb{S}(\mathbb{C}^{2^d}) \quad \text{to} \quad \mathrm{Herm}(A \otimes R),$$

where $\mathrm{Herm}(A \otimes R)$ is regarded as a real Hilbert space with the Hilbert-Schmidt inner product. We can therefore consider the derivatives of $F_d^{\mathcal{A}}$ with respect to the states in Φ.

Dealing with errors. Recall from Eq. (13). Because of tracing out B, ρ_Φ may be mixed even though $|\Xi_\Phi\rangle$ is pure. For each Φ, define its distance from the set of unitary Choi states by

$$\eta(\Phi) := \inf_{V \in \mathsf{U}(A)} \left\| \rho_\Phi - (V \otimes I_R) \, |\Omega\rangle\langle\Omega| \, (V \otimes I_R)^\dagger \right\|_1 .$$

By compactness of $\mathsf{U}(A)$, the infimum must be attained by some V. Choose any minimizer V_Φ and let

$$P_\Phi := (V_\Phi \otimes I_R) \, |\Omega\rangle\langle\Omega| \, (V_\Phi \otimes I_R)^\dagger, \qquad \|\rho_\Phi - P_\Phi\|_1 = \eta(\Phi).$$

Since $|\Xi_\Phi\rangle$ purifies ρ_Φ, Lemma 2.5 gives a pure state $|\chi_\Phi\rangle_B$ such that the corresponding purification of P_Φ approximates $|\Xi_\Phi\rangle$:

$$|\Xi_\Phi\rangle = \left|\Xi_\Phi^{\mathrm{id}}\right\rangle + |e_\Phi\rangle , \quad \left|\Xi_\Phi^{\mathrm{id}}\right\rangle := \left((V_\Phi \otimes I_R) \, |\Omega\rangle_{AR}\right) |\chi_\Phi\rangle_B , \quad \|e_\Phi\| \le \sqrt{\eta(\Phi)}.$$

Derivatives. Fix $t \in [q_d]$. We occasionally consider the decomposition

$$W_\Phi = W_{>t,\Phi} S^\Phi W_{<t,\Phi},$$

where $W_{<t,\Phi}$ contains everything before the t-th query and $W_{>t,\Phi}$ contains everything after it. Applying the Leibniz rule of derivatives to Eq. (12) gives

$$\partial_{x,\delta} W_\Phi = \sum_{t=1}^{q_d} W_{>t,\Phi} \big(\partial_{x,\delta} S^\Phi\big) W_{<t,\Phi}.$$

Define the t-th query part of the derivative of the purified output by

$$|\Delta_{t,x,\delta}\rangle := \big(W_{>t,\Phi} \big(\partial_{x,\delta} S^\Phi\big) W_{<t,\Phi} \otimes I_R\big) \left(|\Omega\rangle_{AR} \, |0\rangle_B\right).$$

The Leibniz rule and linearity of the partial trace give

$$\partial_{x,\delta} \rho_\Phi = \sum_{t=1}^{q_d} \mathrm{Tr}_B \left(|\Delta_{t,x,\delta}\rangle\langle\Xi_\Phi| + |\Xi_\Phi\rangle\langle\Delta_{t,x,\delta}| \right).$$

The squared gradient norm of $F_d^{\mathcal{A}}$ at Φ is

$$\|\nabla F_d^{\mathcal{A}}(\Phi)\|_{\mathrm{HS}}^2 := \sum_{x \in \{0,1\}^d} \sum_{\delta} \|\partial_{x,\delta} \rho_\Phi\|_2^2 .$$

The following lemma gives an upper bound on the gradient. The remainder of this subsection is devoted to the proof of this lemma.

Lemma 4.4. *For every CHFS state family Φ,*

$$\|\nabla F_d^{\mathcal{A}}(\Phi)\|_{\mathrm{HS}}^2 = \sum_{x \in \{0,1\}^d} \sum_{\delta} \|\partial_{x,\delta} \rho_\Phi\|_2^2 \le 72 q_d^2 \big(1 + 2^d \eta(\Phi)\big).$$

Proof. By the preceding derivative identity and Cauchy–Schwarz,

$$\|\partial_{x,\delta}\rho_\Phi\|_2^2 \le q_d \sum_{t=1}^{q_d} \|\mathrm{Tr}_B\left(|\Delta_{t,x,\delta}\rangle\langle\Xi_\Phi| + |\Xi_\Phi\rangle\langle\Delta_{t,x,\delta}|\right)\|_2^2 . \tag{14}$$

Recall $|\Xi_\Phi\rangle = \left|\Xi_\Phi^{\mathrm{id}}\right\rangle + |e_\Phi\rangle$. To deal with them separately, define

$$X_{t,x,\delta} := \mathrm{Tr}_B\left(\left|\Delta_{t,x,\delta}\right\rangle\!\left\langle\Xi_\Phi^{\mathrm{id}}\right|\right), Y_{t,x,\delta} := \mathrm{Tr}_B\left(|\Delta_{t,x,\delta}\rangle\langle e_\Phi|\right).$$

Then, we have $\mathrm{Tr}_B\left(|\Delta_{t,x,\delta}\rangle\langle\Xi_\Phi| + |\Xi_\Phi\rangle\langle\Delta_{t,x,\delta}|\right) = X_{t,x,\delta} + X_{t,x,\delta}^\dagger + Y_{t,x,\delta} + Y_{t,x,\delta}^\dagger$, and therefore

$$\|\mathrm{Tr}_B\left(|\Delta_{t,x,\delta}\rangle\langle\Xi_\Phi| + |\Xi_\Phi\rangle\langle\Delta_{t,x,\delta}|\right)\|_2^2 \le 8\|X_{t,x,\delta}\|_2^2 + 8\|Y_{t,x,\delta}\|_2^2 . \tag{15}$$

We first bound the terms $X_{t,x,\delta}$. Let $\{|a\rangle\}_{a=1}^{d_A}$ be the orthonormal basis of A used to define $|\Omega\rangle_{AR}$. Define

$$|s_a\rangle := W_{>t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)W_{<t,\Phi}\big(|a\rangle_A|0\rangle_B\big).$$

Then

$$|\Delta_{t,x,\delta}\rangle = \frac{1}{\sqrt{d_A}}\sum_{a=1}^{d_A}|s_a\rangle|a\rangle_R , \quad \left|\Xi_\Phi^{\mathrm{id}}\right\rangle = \frac{1}{\sqrt{d_A}}\sum_{b=1}^{d_A} V_\Phi|b\rangle_A|\chi_\Phi\rangle_B|b\rangle_R .$$

Therefore,

$$X_{t,x,\delta} = \frac{1}{d_A}\sum_{a,b=1}^{d_A}\left(I_A\otimes\langle\chi_\Phi|\right)|s_a\rangle\langle b|V_\Phi^\dagger\otimes|a\rangle\langle b|_R .$$

Since the elements $\{|a\rangle\langle b|_R\}_{a,b}$ can be regarded as orthonormal matrices with respect to the Hilbert-Schmidt inner product,

$$\|X_{t,x,\delta}\|_2^2 = \frac{1}{d_A^2}\sum_{a,b=1}^{d_A}\|\left(I_A\otimes\langle\chi_\Phi|\right)|s_a\rangle\|^2 .$$

Define the projections regarding the initial states and the final states and the operators $W_{<t}$ and $W_{>t}$ by

$$P_{t,\Phi} := W_{<t,\Phi}\big(I_A\otimes|0\rangle\langle 0|_B\big)W_{<t,\Phi}^\dagger, \qquad Q_{t,\Phi} := W_{>t,\Phi}^\dagger\big(I_A\otimes|\chi_\Phi\rangle\langle\chi_\Phi|_B\big)W_{>t,\Phi}.$$

Both $P_{t,\Phi}$ and $Q_{t,\Phi}$ are orthogonal projections of rank d_A. Then we can write

$$\begin{aligned}
\|X_{t,x,\delta}\|_2^2 &= \frac{1}{d_A^2}\sum_{a,b}\|\left(I_A\otimes\langle\chi_\Phi|\right)|s_a\rangle\|^2 \\
&= \frac{1}{d_A}\sum_{a}\left\|\left(I_A\otimes\langle\chi_\Phi|\right)W_{>t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)W_{<t,\Phi}\big(|a\rangle_A|0\rangle_B\big)\right\|^2 \\
&= \frac{1}{d_A}\mathrm{Tr}\left(P_{t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)^\dagger Q_{t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)P_{t,\Phi}\right) \\
&= \frac{1}{d_A}\left\|Q_{t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)P_{t,\Phi}\right\|_2^2 .
\end{aligned}$$

Since $P_{t,\Phi}$ and $Q_{t,\Phi}$ have rank d_A, Lemma 4.2 gives

$$\sum_{x,\delta}\|X_{t,x,\delta}\|_2^2 = \frac{1}{d_A}\sum_{x,\delta}\left\|Q_{t,\Phi}\big(\partial_{x,\delta}S^\Phi\big)P_{t,\Phi}\right\|_2^2 \le 9. \tag{16}$$

For the error term, $\|e_\Phi\| \le \sqrt{\eta(\Phi)}$ gives

$$\|Y_{t,x,\delta}\|_2^2 = \|\mathrm{Tr}_B\left(|\Delta_{t,x,\delta}\rangle\langle e_\Phi|\right)\|_2^2 \le \|\Delta_{t,x,\delta}\|^2\|e_\Phi\|^2 \le \eta(\Phi)\|\Delta_{t,x,\delta}\|^2$$

where we use $\|\mathrm{Tr}_B(|x\rangle\langle y|)\|_2 \le \|x\|\,\|y\|$. The pre-query state

$$(W_{<t,\Phi} \otimes I_R)\big(|\Omega\rangle_{AR}\,|0\rangle_B\big)$$

is a unit vector. Including the reference register and all other non-query registers in the ancillary register of Lemma 4.3, and using invariance of the norm under $W_{>t,\Phi}$, gives

$$\sum_{x,\delta} \|Y_{t,x,\delta}\|_2^2 \le 5 \cdot 2^d \eta(\Phi). \tag{17}$$

Summing Eq. (14) over x, δ, and using Eqs. (15) to (17), gives

$$8q_d^2\big(9 + 5 \cdot 2^d\eta(\Phi)\big) \le 72q_d^2\big(1 + 2^d\eta(\Phi)\big),$$

which proves the claim. □

4.4 CHFS oracle concentration

This section presents the concentration phenomena for the output state ρ_Φ using the perspective of $F_d^{\mathcal{A}}$. We prove the 1-norm bound instead of the 2-norm bound as the 1-norm is more convenient when discussing the distinguishability.

We first consider the CHFS states at a fixed input length d, while keeping the state families at all other input lengths fixed. Let q_d be the number of length-d queries made by the algorithm. All the expectations are over the product Haar measure on the ambient space. In particular, for $\Phi = \Phi_d$, the Haar measure is over $\prod_{x\in\{0,1\}^d} \mathbb{S}(\mathbb{C}^{2^d})$.

Lemma 4.5. *Let $\Phi = \Phi_d$ be sampled from the product Haar measure and define the expected output $\mu := \mathbb{E}_\Phi \rho_\Phi$ and the expected error bound $\bar\eta := \mathbb{E}_\Phi \eta(\Phi)$. Then, for a universal constant $C > 0$,*

$$\mathbb{E}_\Phi \|\rho_\Phi - \mu\|_1 \le C q_d \left(2^{-d/2} + \sqrt{\bar\eta}\right).$$

Proof. If $q_d = 0$, then ρ_Φ is independent of Φ, and the claim is immediate. Assume that $q_d \ge 1$.

Let Φ' be an independent copy of Φ, i.e., a Haar random element from the product space. By independence,

$$\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_2^2 = 2\mathbb{E}_\Phi \|\rho_\Phi - \mu\|_2^2 \le \frac{2}{2^{d+1}-1} \mathbb{E}_\Phi \sum_{x,\delta} \|\partial_{x,\delta}\rho_\Phi\|_2^2,$$

where we use the Poincaré inequality in Theorem 3.4. By Lemma 4.4,

$$\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_2^2 \le \frac{144q_d^2}{2^{d+1}-1}\left(1 + 2^d\bar\eta\right) \le 144q_d^2\left(2^{-d} + \bar\eta\right).$$

Jensen's inequality gives

$$\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_2 \le 12 q_d \left(2^{-d/2} + \sqrt{\bar\eta}\right).$$

Now we turn this bound into the 1-norm bound. Applying Lemma 2.6 to $\rho_\Phi, \rho_{\Phi'}$ and then averaging gives

$$\begin{aligned}
\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_1 &\le (1+\sqrt{2})\mathbb{E}_{\Phi,\Phi'}\big(\eta(\Phi) + \eta(\Phi')\big) + \sqrt{2}\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_2 \\
&= 2(1+\sqrt{2})\bar\eta + \sqrt{2}\mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_2 \\
&\le C q_d \left(2^{-d/2} + \sqrt{\bar\eta}\right).
\end{aligned} \tag{18}$$

Here $0 \le \bar\eta \le 2$, so $\bar\eta \le \sqrt{2\bar\eta}$, and $q_d \ge 1$. Finally, for fixed Φ, convexity of the trace norm gives

$$\|\rho_\Phi - \mu\|_1 = \big\|\mathbb{E}_{\Phi'}\big(\rho_\Phi - \rho_{\Phi'}\big)\big\|_1 \le \mathbb{E}_{\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_1.$$

Averaging over Φ and applying Eq. (18) proves the claim. □

Lemma 4.6. *Fix an integer $\tau \geq 0$ and the CHFS state families $\Phi_{\leq\tau}$. For each $d > \tau$, sample Φ_d independently from its product Haar measure, and let $\Phi = (\Phi_{\leq\tau}, \Phi_{>\tau})$ denote the full family of CHFS states. Let ρ_Φ and $\eta(\Phi)$ be the output state and its distance from unitary Choi states defined above. Let $q_{>\tau} := \sum_{d>\tau} q_d$ and $\mu := \mathbb{E}_{\Phi_{>\tau}} \rho_\Phi$. Then, for a universal constant $C > 0$,*

$$\mathbb{E}_{\Phi_{>\tau}} \|\rho_\Phi - \mu\|_1 \leq C q_{>\tau} \left(2^{-\tau/2} + \sqrt{\mathbb{E}_{\Phi_{>\tau}} \eta(\Phi)}\right).$$

Proof. The claim is immediate if $q_{>\tau} = 0$. Otherwise, let $d_{\max}$ be the largest input length queried by $\mathcal{A}$. Let $\Phi' = (\Phi_{\leq\tau}, \Phi'_{>\tau})$, where $\Phi'_{>\tau}$ is an independent copy of $\Phi_{>\tau}$.

For $\tau \leq j \leq d_{\max}$, let $\Phi^{(j)}$ be the family obtained from Φ by replacing Φ_d with Φ'_d for every $\tau < d \leq j$. Thus $\rho_{\Phi^{(\tau)}} = \rho_\Phi$ and $\rho_{\Phi^{(d_{\max})}} = \rho_{\Phi'}$, since $\mathcal{A}$ makes no queries of input length greater than $d_{\max}$. Every $\Phi^{(j)}$ has the same distribution as Φ, with $\Phi_{\leq\tau}$ fixed.

For each $\tau < d \leq d_{\max}$, the families $\Phi^{(d-1)}$ and $\Phi^{(d)}$ differ only at input length d. Condition on their common state families and apply Eq. (18) to the independent pair Φ_d, Φ'_d. Averaging over the common state families and applying Jensen's inequality gives

$$\mathbb{E}_{\Phi,\Phi'} \|\rho_{\Phi^{(d-1)}} - \rho_{\Phi^{(d)}}\|_1 \leq C q_d \left(2^{-d/2} + \sqrt{\mathbb{E}_{\Phi_{>\tau}} \eta(\Phi)}\right).$$

Finally, convexity of the trace norm and the triangle inequality give

$$\mathbb{E}_{\Phi_{>\tau}} \|\rho_\Phi - \mu\|_1 \leq \mathbb{E}_{\Phi,\Phi'} \|\rho_\Phi - \rho_{\Phi'}\|_1 \leq \sum_{\tau<d\leq d_{\max}} \mathbb{E}_{\Phi,\Phi'} \|\rho_{\Phi^{(d-1)}} - \rho_{\Phi^{(d)}}\|_1 .$$

Applying the preceding bound and using $2^{-d/2} \leq 2^{-\tau/2}$ and $\sum_{\tau<d\leq d_{\max}} q_d = q_{>\tau}$ proves the claim. □

4.5 CHFS oracle simulation

Finally, we recall the simulation of the CHFS algorithm using phase-invariant designs [1, 22, 10]. A distribution $\mathcal{D}$ on the unit sphere of $\mathbb{C}^D$ is a *phase-invariant δ-approximate t-design* if

$$(1-\delta) \mathop{\mathbb{E}}_{\psi\leftarrow\mathcal{H}} (|\psi\rangle\langle\psi|)^{\otimes s} \preceq \mathop{\mathbb{E}}_{\psi\leftarrow\mathcal{D}} (|\psi\rangle\langle\psi|)^{\otimes s} \preceq (1+\delta) \mathop{\mathbb{E}}_{\psi\leftarrow\mathcal{H}} (|\psi\rangle\langle\psi|)^{\otimes s},$$

for $0 \leq s \leq t$, and, for $0 \leq i, j \leq t$ with $i \neq j$, $\mathbb{E}_{\psi\leftarrow\mathcal{D}} |\psi\rangle^{\otimes i} \langle\psi|^{\otimes j} = 0$. The zeroth tensor power is the scalar 1. The construction of such a phase-invariant approximate design is given in [10] using the standard design from [12].

Lemma 4.7 (Phase-invariant designs)**.** *There is a deterministic classical algorithm, in time* $\mathrm{poly}(n, t, \log(1/\delta))$*, that, given $n, t \in \mathbb{N}$, $0 < \delta < 1$, and a seed s of length* $\mathrm{poly}(n, t, \log(1/\delta))$*, outputs a classical description of an n-qubit unitary circuit U_s of size* $\mathrm{poly}(n, t, \log(1/\delta))$*. For uniformly random s, the states $|\psi_s\rangle = U_s |0^n\rangle$ form a phase-invariant δ-approximate t-design.*

We use the following quantum-output version of [10], which can be proven in essentially the same way.

Lemma 4.8 (CHFS design simulation)**.** *Consider a q-query oracle algorithm A given a CHFS oracle S^Φ for $\Phi = (|\phi_x\rangle)_{x\in\{0,1\}^*}$ where the states in Φ are sampled independently from Haar measure. Let $\Psi = \{|\psi_x\rangle\}_{x\in\{0,1\}^*}$ where each state $|\psi_x\rangle$ is independently sampled from a distribution of phase-invariant δ-approximate $2q$-design states of the same length as $|\phi_x\rangle$. Suppose that A makes queries of input length at most L. Let ρ_Φ and ρ_Ψ be the output states of the algorithm before the measurement given oracle S^Φ and S^Ψ, respectively. Then, there are universal constants $K > 0$ and $C > 1$ such that*

$$\left\| \mathop{\mathbb{E}}_{\Phi} \rho_\Phi - \mathop{\mathbb{E}}_{\Psi} \rho_\Psi \right\|_1 \leq K q C^{q(L+1)} \delta.$$

The same bound holds when the design seeds are $2q$-wise independent. The output may include a reference register. The same conclusion holds if the oracle states at any chosen input lengths are fixed in both computations, with expectations taken over the remaining states.

Proof sketch. We follow the proof of [10, Lemma 7] that gives the advantage bound, keeping the query count q and the maximum input length L separate.

In that proof, the advantage is bounded by

$$\left|\Pr_{\Phi}[A^{S^{\Phi}} \to 1] - \Pr_{\Psi}[A^{S^{\Psi}} \to 1]\right| \leq O\left(D^{2q} \left\| \mathop{\mathbb{E}}_{\Phi} J(\mathcal{S}^{\Phi})^{\otimes q} - \mathop{\mathbb{E}}_{\Psi} J(\mathcal{S}^{\Psi})^{\otimes q} \right\|_1\right),$$

where $\mathcal{S}^{\Phi}, \mathcal{S}^{\Psi}$ are the channels induced by the finite direct sums of the full controlled CHFS gates at input lengths at most L, and D is their dimension. This uses the argument of [22, Lemmas 23 and 24], which simulates the queries using Choi states at the cost of a factor D^{2q}.

For the output states, append an arbitrary two-outcome measurement to A and apply the same bound. Taking the supremum over measurements gives

$$\left\| \mathop{\mathbb{E}}_{\Phi} \rho_{\Phi} - \mathop{\mathbb{E}}_{\Psi} \rho_{\Psi} \right\|_1 \leq O\left(D^{2q} \left\| \mathop{\mathbb{E}}_{\Phi} J(\mathcal{S}^{\Phi})^{\otimes q} - \mathop{\mathbb{E}}_{\Psi} J(\mathcal{S}^{\Psi})^{\otimes q} \right\|_1\right),$$

with the factor of two absorbed into the implicit constant. Since $D = 2^{O(L+1)}$ and the number of oracle labels is at most 2^{L+1}, the moment bound in that proof gives $O(qC^{q(L+1)}\delta)$ for a universal constant $C > 1$. Each term in the moment expansion involves at most $2q$ distinct inputs, so $2q$-wise independent design seeds suffice. □

5 Oracle Separation of PRUs from PRFSGs

This section proves our main theorem.

Theorem 5.1. *There is a unitary oracle relative to which adaptively secure, quantum-accessible PRFSGs exist, whereas PRUs do not.*

Our oracle is $\mathcal{O}^{\Phi} := (S^{\Phi}, \mathsf{QPSPACE})$, where S^{Φ} is a standard CHFS oracle. The construction of the PRFSGs is as follows: Let $\kappa, m = \omega(\log \lambda)$ be such that $\kappa + m = \lambda$. The generation algorithm on the key $k \in \{0,1\}^{\kappa}$ and input $x \in \{0,1\}^{m}$ outputs the state $|\phi_{k\|x}\rangle$, which is obtained by applying S^{Φ}_{λ} on $|k\|x, 0\rangle$. [10, 18] proved that this construction is an adaptively secure quantum-accessible PRFSG relative to the oracle.

We focus on the non-existence of PRUs in the same oracle model. Fix an oracle QPT algorithm $\mathcal{G}$ that is a candidate PRU construction, which is of the form Eq. (11) that acts on $n = n(\lambda)$ qubits for $n = \omega(\log \lambda)$ with a λ-bit key. Let q_d be the number of length-d CHFS queries and let $q = \sum_d q_d$ be the total number of queries. Suppose that $\mathcal{G}$ runs in time λ^a for a fixed constant a that depends on $\mathcal{G}$. Let $d_{\max} \leq \lceil \lambda^a \rceil$ be the maximum query length. We assume $q > 0$, otherwise $\mathcal{G}$ is insecure due to $\mathsf{QPSPACE}$.[3]

We first prove the dichotomy property of the target Choi state in Section 5.1, which presents the main object, the averaged Choi state, of the attack and shows that it satisfies a concentration bound or the construction is far from unitary. Section 5.2 shows how to simulate the averaged Choi state without using the CHFS oracle. Section 5.3 and Section 5.4 present the distinguisher and its analysis.

5.1 Dichotomy of the averaged Choi state

The main ingredient of the distinguisher is the Choi state of the candidate algorithm $\mathcal{G}$. More precisely, we consider the average Choi state where the CHFS oracle for large inputs is replaced by an independent random CHFS oracle.

[3] Alternatively, we assume that $q_d > 0$ for all $d \leq d_{\max}$ using a dummy query.

Choi states and independent Haar states. Let $\tau = \lceil (10a+19)\log_2 \lambda \rceil$ be the threshold. Let A be the n-qubit input register of $\mathcal{G}_k$, and let R be an n-qubit reference register. Recall that for a channel $\mathcal{E}$ on A

$$|\Omega\rangle_{AR} = 2^{-n/2} \sum_{u\in\{0,1\}^n} |u\rangle_A |u\rangle_R , \qquad J(\mathcal{E}) = (\mathcal{E}\otimes \mathcal{I}_R)(|\Omega\rangle\langle\Omega|).$$

Thus, one copy of the Choi state $J(\mathcal{E})$ is prepared by applying $\mathcal{E}$ to the A register of $|\Omega\rangle_{AR}$. Any ancillary registers of $\mathcal{E}$ are traced out. We are mainly interested in the Choi state

$$\rho_{k,\Phi} = J(\mathcal{G}_k^{\mathcal{O}^\Phi})$$

that can be obtained by applying $\mathcal{G}$ once.

For families of pure states defining the CHFS oracles Φ, we write $\Phi_{\le\tau} = (\Phi_d)_{d\le\tau}$ and $\Phi_{>\tau} = (\Phi_d)_{d>\tau}$. We consider another family of states that are sampled from the identical distribution $\Phi'_{>\tau} = (\Phi'_d)_{d>\tau}$: for each $d > \tau$ and $x \in \{0,1\}^d$, independently sample a d-qubit Haar state $\left|\phi'_{d,x}\right\rangle$ and use these states to define $S_d^{\Phi'}$ by the same CHFS construction. These states are independent of the original oracle Φ.

The state $\rho_{k,(\Phi_{\le\tau},\Phi'_{>\tau})}$ denotes the Choi state for $\mathcal{G}$ using the CHFS oracle with $\Phi_{\le\tau}, \Phi'_{>\tau}$, i.e., one with the original S_d^Φ for $d \le \tau$ and $S_d^{\Phi'}$ for $d > \tau$. Let

$$\mu_k := \mathop{\mathbb{E}}_{\Phi'_{>\tau}} \left[\rho_{k,(\Phi_{\le\tau},\Phi'_{>\tau})}\right]$$

be the average of the states $\rho_{k,(\Phi_{\le\tau},\Phi'_{>\tau})}$ over (product) Haar random $\Phi'_{>\tau}$. In particular, μ_k depends on the original oracle only through $\Phi_{\le\tau}$.

Dichotomy. The next lemma shows that either the Choi states have noticeable average mixedness over the key, or they are close on average to the states μ_k.

Lemma 5.2 (Mixedness or concentration)**.** *With probability one over Φ, for all sufficiently large λ, either*

$$\mathop{\mathbb{E}}_k[1 - \mathrm{Tr}(\rho_{k,\Phi}^2)] \ge \lambda^{-(4a+9)} \quad \text{or} \quad \mathop{\mathbb{E}}_k \|\rho_{k,\Phi} - \mu_k\|_1 \le \lambda^{-1}.$$

Proof. Fix $\Phi_{\le\tau}$. If $d_{\max} \le \tau$, then $\rho_{k,\Phi} = \mu_k$ and the claim is immediate. Otherwise, consider two choices $\Phi_{>\tau}$ and $\Phi'_{>\tau}$.

Observe that for unit vectors ϕ, ϕ', the definition in Eq. (8) gives

$$\|S_\phi - S_{\phi'}\|_\infty \le \||\phi,1\rangle\langle\phi,1| - |\phi',1\rangle\langle\phi',1|\|_\infty + 2\|\phi-\phi'\| \le 4\|\phi-\phi'\|.$$

Since S_d^Φ is block diagonal in the input x,

$$\left\|S_d^\Phi - S_d^{\Phi'}\right\|_\infty = \max_{x\in\{0,1\}^d} \left\|S_{\phi_{d,x}} - S_{\phi'_{d,x}}\right\|_\infty \le 4 \max_{x\in\{0,1\}^d} \left\|\phi_{d,x} - \phi'_{d,x}\right\|.$$

The same operator-norm bound holds for the controlled gates. The diamond distance between the channels induced by two unitaries is at most twice their operator-norm distance. Thus a hybrid argument over the queries gives

$$\begin{aligned}
\left\|\rho_{k,\Phi} - \rho_{k,(\Phi_{\le\tau},\Phi'_{>\tau})}\right\|_1 &\le 2 \sum_{\tau<d\le d_{\max}} q_d \left\|S_d^\Phi - S_d^{\Phi'}\right\|_\infty \\
&\le 8 \sum_{\tau<d\le d_{\max}} q_d \max_{x\in\{0,1\}^d} \left\|\phi_{d,x} - \phi'_{d,x}\right\| \\
&\le 8q \left(\sum_{\tau<d\le d_{\max}} \sum_{x\in\{0,1\}^d} \left\|\phi_{d,x} - \phi'_{d,x}\right\|^2\right)^{1/2}.
\end{aligned}$$

Since $|\operatorname{Tr}(\rho^2) - \operatorname{Tr}(\sigma^2)| \le 2\|\rho - \sigma\|_1$, the function $\Phi_{>\tau} \mapsto \mathbb{E}_k[1 - \operatorname{Tr}\left(\rho_{k,\Phi}^2\right)]$ is $16q$-Lipschitz in the product Euclidean distance.

For each (d, x) with $\tau < d \le d_{\max}$, write $\phi_{d,x} = U_{d,x}\left|0^d\right\rangle$ with independent Haar-random unitaries $U_{d,x}$. Since $\|U\left|0\right\rangle - V\left|0\right\rangle\| \le \|U - V\|_2$, this Lipschitz bound also holds in the product Frobenius norm. We can therefore apply [22, Theorem 10], with minimum unitary dimension at least $2^{\tau+1}$.

If $\mathbb{E}_{k,\Phi_{>\tau}}[1 - \operatorname{Tr}\left(\rho_{k,\Phi}^2\right)] \ge 2\lambda^{-(4a+9)}$, the lower-tail bound in [22, Theorem 10] gives

$$\mathbb{P}_{\Phi_{>\tau}}\left[\mathop{\mathbb{E}}_{k}[1 - \operatorname{Tr}(\rho_{k,\Phi}^2)] < \lambda^{-(4a+9)}\right] \le \exp\left(-\frac{(2^{\tau+1}-2)\lambda^{-2(4a+9)}}{24(16q)^2}\right) = \exp(-\Omega(\lambda)).$$

If instead $\mathbb{E}_{k,\Phi_{>\tau}}[1 - \operatorname{Tr}\left(\rho_{k,\Phi}^2\right)] < 2\lambda^{-(4a+9)}$, Lemma 2.7 and Jensen's inequality give

$$\mathop{\mathbb{E}}_{k,\Phi_{>\tau}} e(\mathcal{G}_k^{\mathcal{O}^\Phi}) \le 5\sqrt{\mathop{\mathbb{E}}_{k,\Phi_{>\tau}}[1 - \operatorname{Tr}\left(\rho_{k,\Phi}^2\right)]} < 5\sqrt{2\lambda^{-(4a+9)}}.$$

Applying the CHFS oracle concentration lemma (Lemma 4.6) for each key and then Jensen's inequality yields

$$\mathop{\mathbb{E}}_{k,\Phi_{>\tau}} \|\rho_{k,\Phi} - \mu_k\|_1 \le Cq\left(2^{-\tau/2} + \sqrt{\mathop{\mathbb{E}}_{k,\Phi_{>\tau}} e(\mathcal{G}_k^{\mathcal{O}^\Phi})}\right) \le C_1 q\left(2^{-\tau/2} + \lambda^{-(4a+9)/4}\right).$$

Markov's inequality therefore gives

$$\mathbb{P}_{\Phi_{>\tau}}\left[\mathop{\mathbb{E}}_{k}\|\rho_{k,\Phi} - \mu_k\|_1 > \lambda^{-1}\right] \le C_1\lambda q\left(2^{-\tau/2} + \lambda^{-(4a+9)/4}\right) = O(\lambda^{-5/4}).$$

In either case, the probability that both alternatives fail is $O(\lambda^{-5/4})$, uniformly over $\Phi_{\le\tau}$. Averaging over $\Phi_{\le\tau}$ gives the same bound over Φ. These failure probabilities are summable in λ, so the Borel–Cantelli lemma proves the claim. □

5.2 Simulating $\mathcal{G}$

The distinguisher uses the approximation of μ_k. We explain that, after preprocessing on $\Phi_{\le\tau}$ with tomography, the approximation of μ_k can be efficiently prepared for each k without querying S^Φ at all.

Tomography of the small layers. For every $d \le \tau = \lceil(10a + 19)\log_2\lambda\rceil$, the distinguisher performs process tomography in Theorem 2.1 on the controlled gate $|0\rangle\langle 0|_C \otimes I + |1\rangle\langle 1|_C \otimes S_d^\Phi$ with diamond-norm accuracy $1/400q$ and failure probability $2^{-2\lambda}$. Let $\widetilde{S}_d$ be the unitary specified by the classical output of tomography, and let $\widetilde{S}_{\le\tau} = (\widetilde{S}_d)_{d\le\tau}$ denote the collection of these descriptions.

Each reconstructed gate acts on $2d + 2 = O(\log\lambda)$ qubits. By Theorem 2.1, the procedure runs in polynomial time and uses polynomially many queries to S^Φ. A union bound shows that all reconstructed gates have the stated accuracy except with probability at most $(\tau + 1)2^{-2\lambda}$.

Preparing the simulated Choi state. Given $\widetilde{S}_{\le\tau}$ and $\Phi'_{>\tau}$, let $\widetilde{\mathcal{G}}_k^{\Phi'_{>\tau}}$ be the channel obtained from $\mathcal{G}_k$ by replacing each controlled CHFS query of length $d \le \tau$ by $\widetilde{S}_d$, and each query of length $d > \tau$ by the corresponding controlled gate for $S_d^{\Phi'}$. All other operations and queries to $\mathsf{QPSPACE}$ are unchanged.

We approximate the average state $\mathbb{E}_{\Phi'_{>\tau}} J(\widetilde{\mathcal{G}}_k^{\Phi'_{>\tau}})$ using phase-invariant designs, following the simulation in [10]. Let $\mathcal{T}$ be the deterministic classical algorithm of Lemma 4.7. More precisely, on input $(d, 2q, \delta; s)$, where s is a design seed, $\mathcal{T}$ outputs a classical description of a d-qubit unitary circuit $U_{d,s}$ of size $\operatorname{poly}(d, q, \log(1/\delta))$. For uniform s, the states $U_{d,s}\left|0^d\right\rangle$ form a phase-invariant δ-approximate $2q$-design. Choose δ according to Lemma 4.8 so that Eq. (19) holds. Note that $\log(1/\delta) = O(\lambda^{2a})$.

Let $\{s_z\}_{z\in\{0,1\}^\ell}$ be an efficiently computable family of $2q$-wise independent functions on the inputs (d,x) with $\tau < d \le d_{\max}$ and $x \in \{0,1\}^d$ with $\ell = \mathrm{poly}(\lambda)$. The output length of $s_z(d,x)$ equals the seed length for $\mathcal{T}(d,2q,\delta;\cdot)$. For each $\tau < d \le d_{\max}$ and $x \in \{0,1\}^d$, let $|\psi_{d,x,z}\rangle$ be the state $U_{d,s_z(d,x)}\left|0^d\right\rangle$, where $\mathcal{T}(d,2q,\delta;s_z(d,x))$ outputs the description of $U_{d,s_z(d,x)}$. We consider, for each z, the simulation of $\widetilde{\mathcal{G}}_k^{\Phi'_{>\tau}}$ by replacing the state $\left|\phi'_{d,x}\right\rangle$ for the CHFS oracle by $|\psi_{d,x,z}\rangle$. Let $\widehat{\mu}_k$ be the Choi state of this simulation, averaged over uniform z.

Circuit for $\widehat{\mu}_k$. Given λ, k and $\widetilde{S}_{\le\tau}$, the circuit prepares a purification of $\widehat{\mu}_k$ as follows.

1. Prepare an ℓ-qubit register Z in $H^{\otimes\ell}\left|0^\ell\right\rangle$ and prepare $|\Omega\rangle_{AR}$.
2. Run the purified computation of $\mathcal{G}_k$ on A. For each CHFS query with $d \le \tau$, apply $\widetilde{S}_d$. For each CHFS query with $d > \tau$ and input x, perform the following operations controlled on Z:
   (a) Compute the design seed $s_z(d,x)$.
   (b) Run $\mathcal{T}(d,2q,\delta;s_z(d,x))$ to compute the circuit description of $U_{d,s_z(d,x)}$.
   (c) Use this circuit and its inverse to implement the controlled CHFS gate $S_{\psi_{d,x,z}}$ on the query target register, with the original query control.
   (d) Uncompute the circuit description and the design seed.

   The seed and circuit description are computed reversibly. All other gates and QPSPACE queries are unchanged.
3. Output all registers.

For fixed z, this is the simulation defining $\widehat{\mu}_k$. Tracing out Z and all ancillary registers gives its Choi state averaged over uniform z, namely $\widehat{\mu}_k$. Thus the output is a purification of $\widehat{\mu}_k$.

For the chosen δ, both $s_z(d,x)$ and $\mathcal{T}(d,2q,\delta;s_z(d,x))$ can be computed in polynomial time. Since $\mathcal{T}$ outputs circuits of polynomial size, the corresponding unitaries, their inverses, and the controlled CHFS gates can also be implemented in polynomial time. Thus the preparation procedure is a QPT algorithm with access to QPSPACE.

The following lemma ensures the accuracy of the above circuit.

Lemma 5.3. *Given λ, k and the descriptions $\widetilde{S}_{\le\tau}$, the above procedure prepares a purification of $\widehat{\mu}_k$. Except with probability at most $(\tau+1)2^{-2\lambda}$ over tomography, for every key k, $\|\widehat{\mu}_k - \mu_k\|_1 \le 1/200$.*

Proof. The preparation claim follows from the construction above. Fix the tomography descriptions, which are used for all keys. Apply Lemma 4.8 to $\widetilde{\mathcal{G}}_k$ acting on one half of $|\Omega\rangle$ for the CHFS oracles with input length $d > \tau$. The computation makes at most q CHFS oracle queries, each of length at most $\lceil\lambda^a\rceil$. Thus

$$\left\| \widehat{\mu}_k - \mathop{\mathbb{E}}_{\Phi'_{>\tau}} J(\widetilde{\mathcal{G}}_k^{\Phi'_{>\tau}}) \right\|_1 \le KqC^{q(\lceil\lambda^a\rceil+1)}\delta \le \frac{1}{400}, \tag{19}$$

where K, C are the universal constants in that lemma. Since $q \le \lambda^a$, this choice requires only $\log(1/\delta) = O(\lambda^{2a})$.

Now suppose that all tomography estimates have error at most $1/400q$. For every key k and every $\Phi'_{>\tau}$, a hybrid argument over the queries of size $\le\tau$ gives

$$\left\| J(\widetilde{\mathcal{G}}_k^{\Phi'_{>\tau}}) - \rho_{k,(\Phi_{\le\tau},\Phi'_{>\tau})} \right\|_1 \le \sum_{d\le\tau} \frac{q_d}{400q} \le \frac{1}{400}.$$

Averaging over $\Phi'_{>\tau}$ and combining the two bounds gives $\|\widehat{\mu}_k - \mu_k\|_1 \le 1/200$. A union bound over all the process tomography procedures gives failure probability at most $(\tau+1)2^{-2\lambda}$. □

5.3 The distinguisher

Now we are ready to describe the distinguisher. The distinguisher is given oracle access to a channel $\mathcal{V}$, where either $\mathcal{V} = \mathcal{G}_{k^\star}^{\mathcal{O}^\Phi}$ for a uniformly random key $k^\star \in \{0,1\}^\lambda$, or $\mathcal{V}$ is an n-qubit Haar-random unitary channel.

Let $r = \lceil 143\lambda \rceil$. The distinguisher uses the following binary measurement M_k for each k that is defined as follows.

Input: r copies of $J(\mathcal{V})$.

Preparation: Prepare r independent copies of $\widehat{\mu}_k$, using fresh design seeds for each copy and the same tomography descriptions $\widetilde{S}_{\leq\tau}$ for all copies.

Measurement: Apply a swap test to each pair and accept if at least $2r/3$ tests pass.

It is clearly a QPT algorithm with access to $\mathsf{QPSPACE}$ given k and $\widetilde{S}_{\leq\tau}$, given that the preparation of $\widehat{\mu}_k$ is done in QPT with access to $\mathsf{QPSPACE}$.

Our distinguisher is defined as follows.

Distinguisher $\mathcal{D}^{\mathcal{V},\mathcal{O}^\Phi}$.

1. Run the process tomography procedure of Section 5.2 to obtain $\widetilde{S}_{\leq\tau}$.
2. Prepare $r+2$ copies of $|\Omega\rangle_{AR}$ and apply $\mathcal{V}$ to obtain $J(\mathcal{V})^{\otimes(r+2)}$.
3. Apply the purity test of Lemma 2.4 to the first two copies of $J(\mathcal{V})$. If it fails, output 1 and abort.
4. Otherwise, apply the quantum OR test to $\{M_k : k \in \{0,1\}^\lambda\}$ on the remaining r copies of $J(\mathcal{V})$. Output 1 if it accepts and 0 otherwise.

Lemma 5.4 (QPT implementation)**.** *The distinguisher is a uniform QPT algorithm with oracle access to $\mathcal{V}$ and $\mathcal{O}^\Phi = (S^\Phi, \mathsf{QPSPACE})$.*

Proof. The tomography procedure runs in polynomial time as shown in Section 5.2. For fixed tomography descriptions, the preparation for $\widehat{\mu}_k$ and the $r = O(\lambda)$ swap tests give unitary QPT implementations of the measurements M_k with access to $\mathsf{QPSPACE}$, whose circuits can be generated uniformly from k and $\widetilde{S}_{\leq\tau}$. There are 2^λ measurements, so Corollary 2.12 gives a QPT implementation of the quantum OR test with access to $\mathsf{QPSPACE}$. All remaining operations have polynomial size. □

We prove the following lemmas in the next subsection.

Lemma 5.5 (Haar case)**.** *If $\mathcal{V}$ is a Haar-random unitary, the distinguisher outputs 1 with negligible probability.*

Lemma 5.6 (Candidate case)**.** *Fix an oracle Φ for which the conclusion of Lemma 5.2 holds. If $\mathcal{V} = \mathcal{G}_{k^\star}^{\mathcal{O}^\Phi}$ for a uniformly random $k^\star$, then the distinguisher outputs 1 with non-negligible probability. More precisely, $\mathbb{P}[\mathcal{D}^{\mathcal{V},\mathcal{O}^\Phi} = 1] \geq \frac{1}{2}\lambda^{-(4a+9)} - \mathrm{negl}(\lambda)$ for all sufficiently large λ.*

We are now ready to prove our main theorem.

Proof of Theorem 5.1. Enumerate the candidate PRUs $(\mathcal{G}_i)_{i\in\mathbb{N}}$, which form a countable set because uniform oracle machines with explicit polynomial time bounds form a countable set. For each fixed $\mathcal{G}_i$, the dichotomy in Lemma 5.2 holds with probability one over Φ. By countable subadditivity, with probability one over the CHFS oracle, the dichotomy holds for all candidates and the PRFSG security (which happens with probability 1 [10]) holds for the construction described in the introduction of this section. Fix Φ so that this holds.

It suffices to discuss the non-existence of secure PRUs. For every candidate $\mathcal{G}$, the oracle QPT distinguisher $\mathcal{D}$ distinguishes the candidate from a Haar-random unitary channel with inverse-polynomial advantage by Lemmas 5.5 and 5.6. Hence PRUs do not exist relative to $\mathcal{O}^\Phi$. □

5.4 Analysis

Proof of Lemma 5.5. Let $\mathcal{V} = \mathcal{U}$ for a Haar-random unitary U. The purity test always passes because $J(\mathcal{U})$ is pure. It therefore suffices to bound the acceptance probability of the quantum OR test.

Fix the original oracle and the tomography outcomes. Every $\widehat{\mu}_k$ is then fixed independently of U. For each key k, a swap test on $J(\mathcal{U})$ and $\widehat{\mu}_k$ passes with probability $(1 + \mathrm{Tr}(J(\mathcal{U})\widehat{\mu}_k))/2$. Haar invariance gives

$$\mathop{\mathbb{E}}_{U} J(\mathcal{U}) = \frac{I_{AR}}{2^{2n}}, \qquad \mathop{\mathbb{E}}_{U}\left[\frac{1 + \mathrm{Tr}(J(\mathcal{U})\widehat{\mu}_k)}{2}\right] = \frac{1}{2} + \frac{1}{2^{2n+1}}.$$

This swap test uses one query to U. Hence Lemma 2.2 gives, for each k,

$$\mathbb{P}_U\left[\frac{1 + \mathrm{Tr}(J(\mathcal{U})\widehat{\mu}_k)}{2} > \frac{7}{12}\right] \leq \exp(-\Omega(2^n)).$$

A union bound over the keys shows that, except with probability at most $2^\lambda \exp(-\Omega(2^n)) = \mathrm{negl}(\lambda)$, the swap-test acceptance probability is at most $7/12$ for every key. We only consider this case. Then, by Chernoff inequality (Lemma 2.3) and $r = \lceil 143\lambda \rceil$, each M_k accepts with probability at most $2^{-3\lambda}$. Thus the quantum OR test accepts and the distinguisher outputs 1 with probability at most $4 \cdot 2^\lambda \cdot 2^{-3\lambda}$ by Lemma 2.11. Overall, the distinguisher outputs 1 with negligible probability. □

Proof of Lemma 5.6. We first consider the mixedness case, i.e., it holds that $\mathbb{E}_k[1 - \mathrm{Tr}\left(\rho_{k,\Phi}^2\right)] \geq \lambda^{-(4a+9)}$. In this case, the purity test alone outputs 1 with probability at least $\frac{1}{2}\lambda^{-(4a+9)}$.

Otherwise, Lemma 5.2 gives $\mathbb{E}_k \left\|\rho_{k,\Phi} - \mu_k\right\|_1 \leq \lambda^{-1}$. By Markov's inequality, all but a $100\lambda^{-(4a+9)} + 100\lambda^{-1}$ fraction of keys satisfy the near-purity and closeness conditions simultaneously:

$$1 - \mathrm{Tr}\left(\rho_{k,\Phi}^2\right) \leq \frac{1}{100}, \qquad \left\|\rho_{k,\Phi} - \mu_k\right\|_1 \leq \frac{1}{100}.$$

Condition on successful tomography. By Lemma 5.3, every such key satisfies $\left\|\rho_{k,\Phi} - \widehat{\mu}_k\right\|_1 \leq 3/200$, and consequently

$$\mathrm{Tr}(\rho_{k,\Phi}\widehat{\mu}_k) \geq \mathrm{Tr}\left(\rho_{k,\Phi}^2\right) - \left\|\rho_{k,\Phi} - \widehat{\mu}_k\right\|_1 \geq \frac{39}{40}.$$

For each such key k, every swap test in M_k passes with probability at least $79/80$. Since M_k rejects only if fewer than $2r/3$ of its r independent swap tests pass, the Chernoff-Hoeffding bound (Lemma 2.3) gives a rejection probability of at most $2^{-3\lambda}$. For the actual key $k^\star$ satisfying both the near-purity and closeness conditions, the quantum OR test (Lemma 2.11) thus accepts with probability at least $(1 - 2^{-3\lambda})^2/7$.

The purity test acts on separate copies of $J(\mathcal{V})$, and its failure also causes the distinguisher to output 1. Therefore, averaging over the key and accounting for tomography failure gives

$$\mathbb{P}[\mathcal{D}^{\mathcal{V},\mathcal{O}^\Phi} = 1] \geq \left(1 - 100\lambda^{-(4a+9)} - \frac{100}{\lambda}\right)\frac{(1 - 2^{-3\lambda})^2}{7} - (\tau + 1)2^{-2\lambda}$$

where $(\tau + 1)2^{-2\lambda}$ is due to the tomography failure probability from Lemma 5.3. This is bounded below by a positive constant for all sufficiently large λ, proving the claim. □

References

[1] Andris Ambainis and Joseph Emerson. Quantum t-designs: t-wise independence in the quantum world. In *IEEE Conference on Computational Complexity (CCC)*, 2007.

[2] Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 226–254. Springer, Cham, May 2024.

[3] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. Cryptography in the common Haar state model: Feasibility results and separations. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 94–125. Springer, Cham, December 2024.

[4] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 237–265. Springer, Cham, November 2022.

[5] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Cham, August 2022.

[6] Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and Geometry of Markov Diffusion Operators*, volume 348 of *Grundlehren der Mathematischen Wissenschaften*. Springer, 2014.

[7] Mohammed Barhoush. Separating pseudorandom generators from logarithmic pseudorandom states. In *Advances in Cryptology – ASIACRYPT 2026*, 2026. To appear.

[8] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. One-way functions imply secure computation in a quantum world. In *Advances in Cryptology – CRYPTO 2021*, 2021.

[9] Amit Behera, Giulio Malavolta, Tomoyuki Morimae, Tamer Mour, and Takashi Yamakawa. A new world in the depths of microcrypt: Separating OWSGs and quantum money from QEFID. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 23–52. Springer, Cham, May 2025.

[10] Samuel Bouaziz-Ermann, Minki Hhan, Garazi Muguruza, and Quoc-Huy Vu. On limits on the provable consequences of quantum pseudorandomness. In *Advances in Cryptology – ASIACRYPT 2026*, 2026. To appear.

[11] Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 417–440. Springer, Cham, August 2020.

[12] Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346(2):397–434, 2016.

[13] Boyang Chen, Andrea Coladangelo, and Or Sattath. The power of a single haar random state: Constructing and separating quantum pseudorandomness. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 108–137. Springer, Cham, May 2025.

[14] Jelle Don, Serge Fehr, and Yu-Hsuan Huang. Adaptive versus static multi-oracle algorithms, and quantum security of a split-key PRF. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part I*, volume 13747 of *Lecture Notes in Computer Science*, pages 33–51. Springer, 2022.

[15] Ben Foxman, Alex Lombardi, Fermi Ma, Barak Nehoran, and John Wright. Quantum lazy sampling and path recording for any group, 2026.

[16] Eli Goldin and Mark Zhandry. Translating between the common haar random state model and the unitary model. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 269–300. Springer, Cham, August 2025.

[17] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, 1986.

[18] Aditya Gulati, Yao-Ting Lin, Tomoyuki Morimae, and Shogo Yamada. Black-box separation between pseudorandom unitaries, pseudorandom isometries, and pseudorandom function-like states. Cryptology ePrint Archive, Report 2025/1864, 2025.

[19] Jeongwan Haah, Robin Kothari, Ryan O'Donnell, and Ewin Tang. Query-optimal estimation of unitary channels in diamond distance. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, 2023.

[20] Aram Wettroth Harrow, Cedric Yen-Yu Lin, and Ashley Montanaro. Sequential measurements, disturbance and property testing. In Philip N. Klein, editor, *28th SODA*, pages 1598–1611. ACM-SIAM, January 2017.

[21] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Cham, August 2018.

[22] William Kretschmer. Quantum pseudorandomness and classical complexity. In *Theory of Quantum Computation, Communication and Cryptography (TQC)*, 2021. See the updated full version at https://arxiv.org/abs/2103.09320.

[23] William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In Barna Saha and Rocco A. Servedio, editors, *55th ACM STOC*, pages 1589–1602. ACM Press, June 2023.

[24] William Kretschmer, Luowen Qian, and Avishay Tal. Quantum-computable one-way functions without one-way functions. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 189–200. ACM Press, June 2025.

[25] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 979–990. ACM Press, June 2024.

[26] Chuhan Lu, Minglong Qin, Fang Song, Penghui Yao, and Mingnan Zhao. Quantum pseudorandom scramblers. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part II*, volume 15365 of *LNCS*, pages 3–35. Springer, Cham, December 2024.

[27] Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.

[28] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 806–809. ACM Press, June 2025.

[29] Elizabeth S. Meckes. *The Random Matrix Theory of the Classical Compact Groups*, volume 218 of *Cambridge Tracts in Mathematics*. Cambridge University Press, 2019.

[30] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries. In *65th FOCS*, pages 485–492. IEEE Computer Society Press, October 2024.

[31] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Cham, August 2022.

[32] Alexander Poremba, Seyoon Ragavan, and Vinod Vaikuntanathan. Cloning games, black holes and cryptography. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 109:1–109:21, 2026.

[33] Gregory Rosenthal. Efficient quantum state synthesis with one query. In David P. Woodruff, editor, *35th SODA*, pages 2508–2534. ACM-SIAM, January 2024.

[34] Mark Zhandry. How to model unitary oracles. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 237–268. Springer, Cham, August 2025.