Pseudorandom (Function-Like) Quantum State Generators: New Definitions and Applications

Prabhanjan Ananth[*]
UCSB

Aditya Gulati[†]
UCSB

Luowen Qian[‡]
Boston University

Henry Yuen[§]
Columbia University

Abstract

Pseudorandom quantum states (PRS) are efficiently constructible states that are computationally indistinguishable from being Haar-random, and have recently found cryptographic applications. We explore new definitions, new properties and applications of pseudorandom states, and present the following contributions:

1. **New Definitions**: We study variants of pseudorandom *function-like* state (PRFS) generators, introduced by Ananth, Qian, and Yuen (CRYPTO'22), where the pseudorandomness property holds even when the generator can be queried adaptively or in superposition. We show feasibility of these variants assuming the existence of post-quantum one-way functions.
2. **Classical Communication**: We show that PRS generators with logarithmic output length imply commitment and encryption schemes with *classical communication*. Previous constructions of such schemes from PRS generators required quantum communication.
3. **Simplified Proof**: We give a simpler proof of the Brakerski–Shmueli (TCC'19) result that polynomially-many copies of uniform superposition states with random binary phases are indistinguishable from Haar-random states.
4. **Necessity of Computational Assumptions**: We also show that a secure PRS with output length logarithmic, or larger, in the key length necessarily requires computational assumptions.

[*] prabhanjan@cs.ucsb.edu

[†] adityagulati@ucsb.edu

[‡] luowenq@bu.edu. Supported by DARPA under Agreement No. HR00112020023.

[§] hyuen@cs.columbia.edu. Supported by AFOSR award FA9550-21-1-0040 and NSF CAREER award CCF-2144219.

Contents

1 Introduction

The study of pseudorandom objects is central to the foundations of cryptography. After many decades, cryptographers have developed a deep understanding of the zoo of pseudorandom primitives such as one-way functions (OWF), pseudorandom generators (PRG), and pseudorandom functions (PRF) [GGM86, HILL99].

The study of pseudorandomness in the quantum setting, on the other hand, is just getting started. Objects such as state and unitary k-designs have been studied extensively, but these are best thought of as quantum analogues of k-wise independent hash functions [AE07, DCEL09]. There are unconditional constructions of state and unitary designs and they do not imply any computational assumptions [AE07, RS09].

Quantum pseudorandomness requiring computational assumptions, in contrast, has been studied much less. Ji, Liu, and Song introduced the notion of *pseudorandom quantum states (PRS)* and *pseudorandom quantum unitaries (PRU)* [JLS18]. At a high level, these are efficiently sampleable distributions over states/unitaries that are computationally indistinguishable from being sampled from the Haar distribution (i.e., the uniform measure over the space of states/unitaries). Ji, Liu, and Song as well as Brakerski and Shmueli have presented constructions of PRS that are based on quantum-secure OWFs [JLS18, BS19, BS20]. Kretschmer showed, however, that PRS do not necessarily imply OWFs; there are oracles relative to which PRS exist but OWFs don't [Kre21]. This was followed by recent works that demonstrated the cryptographic utility of PRS: basic cryptographic tasks such as bit commitment, symmetric-key encryption, and secure multiparty computation can be accomplished using only PRS as a primitive [AQY21, MY21]. It is an intriguing research direction to find more cryptographic applications of PRS and PRU.

The key idea in [AQY21] that unlocked the aforementioned applications was the notion of a *pseudorandom function-like state (PRFS) generator*. To explain this we first review the definition of PRS generators. A quantum polynomial-time (QPT) algorithm G is a PRS generator if for a uniformly random key $k \in \{0,1\}^\lambda$ (with λ being the security parameter), polynomially-many copies of the state $|\psi_k\rangle = G(k)$ is indistinguishable from polynomially-many copies of a state $|\vartheta\rangle$ sampled from the Haar measure by all QPT algorithms. One can view this as a quantum analogue of classical PRGs. Alternately, one could consider a version of PRS where the adversary only gets one copy of the state. However, as we will see later, the multi-copy security of PRS will play a crucial role in our applications.

The notion of PRFS generator introduced by [AQY21] is a quantum analogue of PRF (hence the name *function-like*): in addition to taking in a key k, the generator G also takes an *input* x (just like a PRF takes a key k and an input x). Let $|\psi_{k,x}\rangle = G(k,x)$. The pseudorandomness property of G is that for all sequences of inputs $(x_1, \ldots, x_s)$ for polynomially large s, averaged over the key k, the collection of states $|\psi_{k,x_1}\rangle^{\otimes t}, \ldots, |\psi_{k,x_s}\rangle^{\otimes t}$ for polynomially large t is computationally indistinguishable from $|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t}$ where the $|\vartheta_i\rangle$'s are sampled independently from the Haar measure. In other words, while PRS generators look like (to a computationally bounded distinguisher) they are sampling a *single* state from the Haar measure, PRFS generators look like they are sampling *many* (as compared to the key length) states from the Haar measure. Importantly, this still holds true even when the distinguisher is given the inputs $x_1, \ldots, x_s$.

As mentioned, this (seemingly) stronger notion of quantum pseudorandomness provided a useful conceptual tool to perform cryptographic tasks (encryption, commitments, secure computation, etc) using pseudorandom states alone. Furthermore, [AQY21] showed that for a number of applications, PRFS generators with logarithmic input length suffices and furthermore such objects can be constructed in a black-box way from PRS generators.[1]

Despite exciting progress in this area in the last few years, there is still much to understand about the properties, relationships, and applications of pseudorandom states. In this paper we explore a number of natural questions about pseudorandom states:

- *Feasibility of Stronger Definitions of PRFS*: In the PRFS definition of [AQY21], it was assumed that the set of inputs on which the adversary obtains the outputs are determined ahead of time. Moreover, the adversary could obtain the output of PRFS on only classical inputs. This is often referred to as *selective security* in the cryptography literature. For many interesting applications, this definition is insufficient[2]. This leads us to ask: *is it feasible to obtain strengthened versions of PRFS that maintain security in the presence of adaptive and superposition queries?*

- *Necessity of Assumptions*: In the classical setting, essentially all cryptographic primitives require computational assumptions, at the very least $\mathsf{P} \neq \mathsf{NP}$. What computational assumptions are required by pseudorandom quantum states? The answer appears to depend on the output length of the PRS generator. Brakerski and Shmueli [BS20] constructed PRS generators with output length $c \log \lambda$ for some $c > 0$ satisfying statistical security (in other words, the outputs are statistically close to being Haar-random). On the other hand, Kretschmer showed that the existence of PRS generators with output length λ implies that $\mathsf{BQP} \neq \mathsf{PP}$ [Kre21]. This leads to an intriguing question: *is it possible to unconditionally show the existence of $n(\lambda)$-length output PRS, for some $n(\lambda) \geq \log(\lambda)$?*

- *Necessity of Quantum Communication*: A common theme in all the different PRS-based cryptographic constructions of [AQY21, MY21] is that the parties involved in the system perform quantum communication. Looking forward, it is conceivable that quantum communication will be a much more expensive resource than having access to a quantum computer. Achieving quantum cryptography with classical communication has been an important direction, dating back to Gavinsky [Gav12]. We ask the following question: *is quantum communication inherent in the cryptographic constructions based on PRS*?

1.1 Our Results

We explore the aforementioned questions. Our results include the following.

Adaptive-Secure and Quantum-Accessible PRFS. As mentioned earlier, the notion of PRFS given by [AQY21] has *selective security*, meaning that the inputs $x_1, \ldots, x_s$ are fixed ahead of time. Another way of putting it is, the adversary can only make non-adaptive, classical queries to the PRFS generator (where by query we mean, submit an input x to the generator and receive $|\psi_{k,x}\rangle = G(k, x)$ where k is the hidden, secret key).

We study the notion of *adaptively secure PRFS*, in which the security holds with respect to adversaries that can make queries to the generator adaptively. We consider two variants of this: one where the adversary is restricted to making classical queries to the generator (we call this

[1] However, unlike the equivalence between PRG and PRF in the classical setting [GGM86], it is not known whether *every* PRFS generator can be constructed from PRS generators in a black-box way.

[2] For example, the application of private-key encryption from PRFS as described in [AQY21] is only selectively secure. This is due to the fact that the underlying PRFS is selectively secure.

a *classically-accessible adaptively secure PRFS*), and one where there are no restrictions at all; the adversary can even query the generator on a *quantum superposition of inputs* (we call this a *quantum-accessible adaptively secure PRFS*). These definitions can be found in Section 4.

We then show feasibility of these definitions by constructing classically- and quantum-accessible adaptively secure PRFS generators from the existence of post-quantum one-way functions. These constructions are given in Section 5.2 and Section 5.3 respectively.

A Sharp Threshold For Computational Assumptions. In Section 6 we show that there is a sharp threshold between when computational assumptions are required for the existence of PRS generators: we give a simple argument that demonstrates that PRS generators with $\log \lambda$-length outputs require computational assumptions on the adversary[3]. This complements the aforementioned result of Brakerski and Shmueli [BS20] that shows $c \log \lambda$-length PRS for some $c > 0$ do not require computational assumptions. We also note that the calculations of [Kre21] can be refined to show that the existence of $(1+\epsilon)\log\lambda$-length PRS for all $\epsilon > 0$ implies that $\mathsf{BQP} \neq \mathsf{PP}$.

PRS-Based Constructions With Classical Communication. We show that bit commitments and pseudo one-time pad schemes can be achieved using only classical communication based on the existence of PRS with λ-bit keys and $O(\log(\lambda))$-output length. This improves upon the previous result of [AQY21] who achieved bit commitments and pseudo one-time pad schemes from PRS using quantum communication. However, we note that [AQY21] worked with a wider range of parameters, while our constructions are based on PRS with $O(\log(\lambda))$-output length.

En route, we use quantum state tomography (or tomography for short), a well studied concept in quantum information. Roughly speaking, tomography, allows for obtaining a classical string u that captures some properties of an unknown quantum state ρ, given many copies of this state.

We develop a new notion called *verifiable tomography* that might particularly be useful in cryptographic settings. Verifiable tomography allows for verifying whether a given string u is consistent (according to some prescribed verification procedure) with a quantum state ρ. We present the definition and instantiations of verifiable tomography in Section 7. In Section 8, we use verifiable tomography to achieve the aforementioned applications. At a high level, our constructions are similar to the ones in [AQY21], except that verifiable tomography is additionally used to make the communication classical.

A Simpler Analysis of Binary-Phase PRS. Consider the following construction of PRS. Let $\{F_k : \{0,1\}^n \to \{0,1\}\}_{k\in\{0,1\}^\lambda}$ denote a (quantum-secure) pseudorandom function family. Then $\{|\psi_k\rangle\}_k$ forms a PRS, where $|\psi_k\rangle$ is defined as

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

In other words, the pseudorandom states are *binary phase states* where the phases are given by a pseudorandom function. This is a simpler construction of PRS than the one originally given

[3] We also note that there is a much more roundabout argument for a quantitatively weaker result: [AQY21] constructed bit commitment schemes from $O(\log\lambda)$-length PRS. If such PRS were possible to construct unconditionally, this would imply information-theoretically secure bit commitment schemes in the quantum setting. However, this contradicts the famous results of [LC97, May97], which rules out this possibility. Our calculation, on the other hand, directly shows that $\log\lambda$ (without any constants in front) is a sharp threshold.

by [JLS18], where the phases are pseudorandomly chosen N-th roots of unity with $N = 2^n$. Ji, Liu, and Song conjectured that the binary phase construction should also be pseudorandom, and this was confirmed by Brakerski and Shmueli [BS19].

We give a simpler proof of this in Section 5.1, which may be of independent interest.

2 Technical Overview

2.1 Threshold For Computational Assumptions

We show that PRS generators with λ-bit keys and $\log \lambda$-length outputs cannot be statistically secure. To show this we construct an inefficient adversary, given polynomially many copies of a state, can distinguish whether the state was sampled from the output distribution of a $\log \lambda$-length PRS generator or sampled from the Haar distribution on $\log \lambda$-qubit states with constant probability.

Simple Case: PRS output is always pure. Let us start with a simple case when the PRS generator is such that each possible PRS state is pure. Consider the subspace spanned by all possible PRS outputs. The dimension of the subspace spanned by these states is atmost 2^λ: the reason being that there are at most 2^λ keys. Now, consider the subspace spanned by t-copies of PRS states. The dimension of this subspace is still at most 2^λ and in particular, independent of t. Define $P^{(t)}$ to be a projector (which could have an inefficient implementation) onto this subspace. By definition, the measurement of t copies of the output of a PRS generator with respect to $P^{(t)}$ always succeeds.

Recall that the subspace spanned by t-copies of states sampled from the Haar distribution (of length $\log \lambda$) is a symmetric subspace of dimension $\binom{2^\lambda+t-1}{t}$. By choosing t as an appropriate polynomial (in particular, set $t \gg \lambda$), we can make $\binom{2^\lambda+t-1}{t} \gg 2^\lambda$, such that a measurement with $P^{(t)}$ on t copies of states sampled from the Haar distribution fails with constant probability. Hence, an adversary, who just runs P, can successfully distinguish between t copies of the output of a $\log \lambda$-length PRS generator and t copies of a sample from a Haar distribution with constant probability.

General Case. Now let us focus on the case when the PRS generator can also output mixed states. Then we have 2 cases:

- *The majority of outputs of the PRS generator are negligibly close to a pure state:* In this case, we show that the previous approach still works. We replace the projector $P^{(t)}$ with a projection onto the space spanned by states closest to the output states of the PRS generator and we can show that modified projector still succeeds with constant probability.

- *The majority of outputs of the PRS generator are not negligibly close to a pure state:* In this case, most PRS outputs have purity[4] non-negligibly away from 1. Thus, we can violate the security of PRS as follows: run polynomially (in λ) many SWAP tests to check if the state is mixed or not. When the input state is from a Haar distribution, the test will always determine the input state to be pure. On the other hand, if the input state is the output of a PRS generator, the test will determine the input to be pure with probability that is non-negligibly bounded away from 1. Thus, this case cannot happen if the PRS generator is secure.

[4] A density matrix ρ has purity p if $\mathrm{Tr}(\rho^2) = p$.

Details can be found in Section 6.

2.2 Cryptographic Applications With Classical Communication

We show how to construct bit commitments and pseudo one-time encryption schemes from $O(\log(\lambda))$-output PRS with classical communication. Previously, [AQY21] achieved the same result for a wider range of parameters. In this overview, we mainly focus on bit commitments since the main techniques used in constructing commitments will be re-purposed for designing pseudo one-time encryption schemes.

We use the construction of bit commitments from [AQY21] as a starting point. Let $d = O(\log \lambda)$, $n = O(\log \lambda)$ and G is a (d, n)-PRFS generator[5]. The commitment scheme from [AQY21] is as follows:

- In the commit phase, the receiver sends a random $2^d n$-qubit Pauli $P = P_1 \otimes P_2 \otimes \cdots \otimes P_{2^d-1}$ to the sender, where each P_i is an n-qubit Pauli. The sender on input bit b, samples a key k uniformly at random from $\{0,1\}^\lambda$. The sender then sends the state $\rho = \bigotimes_{x \in [2^d]} P_x^b \sigma_{k,x} P_x^b$, where $\sigma_{k,x} = G(k,x)$ to the receiver.
- In the reveal phase, the sender sends (k, b) to the receiver. The receiver accepts if $P^b \rho P^b$ is a tensor product of the PRFS evaluations of (k, x), for all $x = 0, \ldots, 2^d - 1$.

To convert this scheme into one that only has classical comunication, we need a mechanism to generate classical information c from ρ, where ρ is generated from (k, b) as above, that have the following properties:

1. *Classical Description*: c can be computed efficiently and does not leak any information about b.
2. *Correctness*: (k, b) is accepted as a valid opening for c,
3. *Binding*: (k', b'), for $b \neq b'$, is rejected as an opening for c

State Tomography. To design such a mechanism, we turn to quantum state tomography. Quantum state tomography is a process that takes as input multiple copies of a quantum state σ and outputs a string u that is close (according to some distance metric) to a classical description of the state σ. In general, tomography procedures require exponential in d number of copies of a state and also run in time exponential in d, where d is the dimension of the state. Since the states in question are $O(\log(\lambda))$-output length PRFS states, all the algorithms in the commitment scheme would still be efficient.

Since performing tomography on a PRFS state does not violate its pseudorandomness property, the hiding property is unaffected. For achieving correctness and binding properties, we need to also equip the tomography process with a verification algorithm, denoted by Verify. A natural verification algorithm that can be associated with the tomography procedure is the following: to check if u is a valid classical description of a state σ, simply run the above tomography procedure on many copies of σ and check if the output density matrix is close to u.

More formally, we introduce a new tomography called verifiable tomography and we present a generic transformation that converts a specific tomography procedure into one that is also verifiable.

[5] This in turn can be built from $O(\log(\lambda))$-output PRS as shown in [AQY21].

We will see how verifiable tomography helps us achieve both correctness and binding. Before we dive into the new notion and understand its properties, we will first discuss the specific tomography procedure that we consider.

Instantiation. We develop a tomography procedure based on [Low21] that outputs a denisity matrix close (constant distance away) to the input with $1 - \mathsf{negl}(\lambda)$ probability. This is an upgrade to the tomography procedure in [Low21], the expected distance of whose output was a constant. To achieve this, we make use of the fact that if we repeat [Low21]'s tomography procedure polynomially many times, most output states cluster around the input at a constant distance with $1 - \mathsf{negl}(\lambda)$ probability. We believe this procedure might be of independent interest. Details about this procedure can be found in Section 7.2.

Verifiable Tomography. Verifiable tomography is a pair of efficient algorithms $(\mathsf{Tomography}, \mathsf{Verify})$ associated with a family of channels Φ_λ such that the following holds:

- *Same-input correctness:* Let $u_1 = \mathsf{Tomography}(\Phi_\lambda(x))$, then $\mathsf{Verify}(u_1, \Phi_\lambda(x))$ accepts with high probability.
- *Different-input correctness:* Let $u_1 = \mathsf{Tomography}(\Phi_\lambda(x_1))$, and $x_1 \neq x_2$, then $\mathsf{Verify}(u_1, \Phi_\lambda(x_2))$ rejects with high probability.

The family of channels we consider corresponds to the PRFS state generation. That is, $\Phi_\lambda(x = (k, i))$ outputs $G(k, i)$. As mentioned earlier, we can generically convert the above instantiation into a verifiable tomography procedure. Let us see how the generic transformation works.

For simplicity, consider the case when the underlying PRFS has perfect state generation, i.e., the output of PRFS is always a pure state. In this case, the verification algorithm is the canonical one that we described earlier: on input u and PRFS key k, input i, it first performs tomography on many copies of $G(k, i)$ to recover u' and then checks if u is close to u' or not. The same-input correctness follows from the tomography guarantee of the instantiation. To prove the different-input correctness, we use the fact that PRFS outputs are close to uniformly distributed and the following fact [AQY21, Fact 6.9]: for two arbitrary n-qubit states $|\psi\rangle$ and $|\phi\rangle$,

$$\mathop{\mathbb{E}}_{P \xleftarrow{\$} \mathcal{P}_n} \left[\left| \langle\psi| P |\phi\rangle \right|^2 \right] = 2^{-n}.$$

Thus, if $x_1 \neq x_2$ then u_1 and u_2 are most likely going to be far and thus, differing-input correctness property is satisfied as well.

The proofs get more involved when the underlying PRFS does not satisfy perfect state generation. We consider PRFS generators that satisfy recognisable abort; we note that this notion of PRFS can be instantiated from PRS, also with $O(\log(\lambda))$ outpout length, using [AQY21]. A $(d(\lambda), n(\lambda))$-PRFS generator G has the *strongly recognizable abort property* if its output can be written as follows: $G_\lambda(k, x) = Tr_{\mathcal{A}} \left(\eta |0\rangle\langle 0| \otimes |\psi\rangle\langle\psi| + (1 - \eta) |\bot\rangle\langle\bot| \right)$, where $\mathcal{A}$ is the register with the first qubit. Moreover, $|\bot\rangle$ is of the form $|1\rangle |\widehat{\bot}\rangle$ for some $n(\lambda)$-qubit state state $|\widehat{\bot}\rangle$ so that, $(\langle 0| \otimes \langle\psi|)(|\bot\rangle) = 0$. The same-input correctness essentially follows as before; however arguing differing-input correctness property seems more challenging.

Consider the following degenerate case: suppose k be a key and x_1, x_2 be two inputs such that PRFS on input (k, x_1) and PRFS on (k, x_2) abort with very high probability (say, close to 1). Note that the recognizable abort property does not rule out this degenerate case. Then, it holds that the outputs $u_1 = \mathsf{Tomography}(\Phi_\lambda(x_1))$ and $u_2 = \mathsf{Tomography}(\Phi_\lambda(x_2))$ are close. $\mathsf{Verify}(u_1, \Phi_\lambda(x_2))$ accepts and thus, the different-input correctness is not satisfied. To handle such degenerate cases, we incorporate the following into the verification procedure: on input (u_1, u_2), reject if either u_1 or u_2 is close to an abort state. Checking whether a classical description of a state is close to an abort state can be done efficiently.

From Verifiable Tomography to Commitments. Incorporating verifiable tomography into the commitment scheme, we have the following:

- The correctness follows from the same-input correctness of the tomography procedure.
- The binding property follows from the different-input correctness of the tomography procedure.
- The hiding property follows from the fact that the output of a PRFS generator is indistinguishable from Haar random, even given polynomially many copies of the state.

3 Preliminaries

We present the preliminaries in this section. We use λ to denote the security parameter. We use the notation $\mathsf{negl}(\cdot)$ to denote a negligible function.

We refer the reader to [NC10] for a comprehensive reference on the basics of quantum information and quantum computation. We use I to denote the identity operator. We use $\mathcal{D}(\mathcal{H})$ to denote the set of density matrices on a Hilbert space $\mathcal{H}$.

Haar Measure. The Haar measure over $\mathbb{C}^d$, denoted by $\mathscr{H}(\mathbb{C}^d)$ is the uniform measure over all d-dimensional unit vectors. One useful property of the Haar measure is that for all d-dimensional unitary matrices U, if a random vector $|\psi\rangle$ is distributed according to the Haar measure $\mathscr{H}(\mathbb{C}^d)$, then the state $U|\psi\rangle$ is also distributed according to the Haar measure. For notational convenience we write $\mathscr{H}_m$ to denote the Haar measure over m-qubit space, or $\mathscr{H}((\mathbb{C}^2)^{\otimes m})$.

Fact 3.1. *We have*

$$\mathop{\mathbb{E}}_{|\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^d)} |\psi\rangle\langle\psi| = \frac{I}{d} \ .$$

3.1 Distance Metrics and Matrix Norms

Trace Distance. Let $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ be density matrices. We write $\mathrm{TD}(\rho, \sigma)$ to denote the trace distance between them, i.e.,

$$\mathrm{TD}(\rho, \sigma) = \frac{1}{2}\|\rho - \sigma\|_1$$

where $\|X\|_1 = \mathrm{Tr}(\sqrt{X^\dagger X})$ denotes the trace norm. We denote $\|X\| := \sup_{|\psi\rangle}\{\langle\psi|X|\psi\rangle\}$ to be the operator norm where the supremum is taken over all unit vectors. For a vector x, we denote its Euclidean norm to be $\|x\|_2$.

Frobenius Norm. The Frobenius norm of a matrix M is

$$\|M\|_F = \sqrt{\sum_{i,j} |M_{i,j}|^2} = \sqrt{\mathrm{Tr}\left(MM^\dagger\right)},$$

where $M_{i,j}$ denotes the $(i,j)^{th}$ entry of M.

We state some useful facts about Frobenius norm below.

Fact 3.2. *For all matrices A, B we have $\|A - B\|_F^2 = \|A\|_F^2 + \|B\|_F^2 - 2\mathrm{Tr}(A^\dagger B)$.*

Fact 3.3. *Let M_0, M_1 be density matricies and $|\psi\rangle$ be a pure state such that $\langle\psi| M_0 |\psi\rangle \leq \alpha$ and $\|M_0 - M_1\|_F^2 \leq \beta$, where $\beta + 2\alpha < 1$ then*

$$\langle\psi| M_1 |\psi\rangle \leq \alpha + \sqrt{\beta} + \sqrt{(2 - 2\alpha)\,\beta}.$$

Proof. From fact 3.2, we have the following:

$$\begin{aligned}\|M_0 - |\psi\rangle\langle\psi|\,\|_F &= \sqrt{\|M_0\|_F^2 + \|\,|\psi\rangle\langle\psi|\,\|_F^2 - 2\mathrm{Tr}(M_0^\dagger\,|\psi\rangle\langle\psi|)} \\ &= \sqrt{\|M_0\|_F^2 + 1 - 2\,\langle\psi| M_0 |\psi\rangle} \\ &\geq \sqrt{\|M_0\|_F^2 + 1 - 2\alpha}.\end{aligned}$$

By triangle inequality, we know

$$\|M_1\|_F \leq \|M_0\|_F + \|M_0 - M_1\|_F \leq ||M_0||_F + \sqrt{\beta}.$$

Similarly by fact 3.2,

$$\|M_1 - |\psi\rangle\langle\psi|\,\|_F = \sqrt{1 + \|M_1\|_F^2 - 2\,\langle\psi| M_1 |\psi\rangle} \leq \sqrt{1 + \left(\|M_0\|_F + \sqrt{\beta}\right)^2 - 2\,\langle\psi| M_1 |\psi\rangle}.$$

By triangle inequality, we know $\|M_0 - |\psi\rangle\langle\psi|\,\|_F \leq \|M_1 - |\psi\rangle\langle\psi|\,\|_F + \|M_0 - M_1\|_F$. Hence,

$$\sqrt{1 + \|M_0\|_F^2 - 2\alpha} \leq \sqrt{1 + \left(\|M_0\|_F + \sqrt{\beta}\right)^2 - 2\,\langle\psi| M_1 |\psi\rangle} + \sqrt{\beta}.$$

By rearranging the terms, we get

$$\langle\psi| M_1 |\psi\rangle \leq \alpha + \|M_0\|_F^2\sqrt{\beta} + \sqrt{\left(1 + \|M_0\|_F^2 - 2\alpha\right)\beta} \leq \alpha + \sqrt{\beta} + \sqrt{(2 - 2\alpha)\,\beta}.$$

□

Fact 3.4. *For any $0 \leq \varepsilon \leq 1$,*

$$\mathsf{Pr}_{|\psi_1\rangle,|\psi_2\rangle \leftarrow \mathscr{H}_n}\left[\|\,|\psi_1\rangle\langle\psi_1| - |\psi_2\rangle\langle\psi_2|\,\|_F^2 \leq \varepsilon\right] \leq \frac{1}{e^{2^n\left(1 - \frac{\varepsilon}{2}\right)}}.$$

Proof. From Fact 3.2,

$$\begin{aligned} \| \, |\psi_1\rangle\langle\psi_1| - |\psi_2\rangle\langle\psi_2| \, \|_F^2 &= \| \, |\psi_1\rangle\langle\psi_1| \, \|_F^2 + \| \, |\psi_2\rangle\langle\psi_2| \, \|_F^2 - 2\mathrm{Tr}\left(|\psi_1\rangle\langle\psi_1| \, |\psi_2\rangle\langle\psi_2| \right) \\ &= 2 - 2|\langle\psi_1|\psi_2\rangle|^2 \end{aligned}$$

Thus, we have the following:

$$\begin{aligned} \mathsf{Pr}_{|\psi_1\rangle,|\psi_2\rangle \leftarrow \mathscr{H}_n}\left[\| \, |\psi_1\rangle\langle\psi_1| - |\psi_2\rangle\langle\psi_2| \, \|_F^2 \leq \varepsilon \right] &= \mathsf{Pr}_{|\psi_1\rangle,|\psi_2\rangle \leftarrow \mathscr{H}_n}\left[|\langle\psi_1|\psi_2\rangle|^2 \geq 1 - \frac{\varepsilon}{2} \right] \\ &\leq \frac{1}{e^{2^n\left(1-\frac{\varepsilon}{2}\right)}}, \end{aligned}$$

where the last inequality was shown in [BHH16] (Equation 14). □

3.2 Quantum Algorithms

A quantum algorithm A is a family of generalized quantum circuits $\{A_\lambda\}_{\lambda\in\mathbb{N}}$ over a discrete universal gate set (such as $\{CNOT, H, T\}$). By generalized, we mean that such circuits can have a subset of input qubits that are designated to be initialized in the zero state, and a subset of output qubits that are designated to be traced out at the end of the computation. Thus a generalized quantum circuit A_λ corresponds to a *quantum channel*, which is a is a completely positive trace-preserving (CPTP) map. When we write $A_\lambda(\rho)$ for some density matrix ρ, we mean the output of the generalized circuit A_λ on input ρ. If we only take the quantum gates of A_λ and ignore the subset of input/output qubits that are initialized to zeroes/traced out, then we get the *unitary part* of A_λ, which corresponds to a unitary operator which we denote by $\hat{A}_\lambda$. The *size* of a generalized quantum circuit is the number of gates in it, plus the number of input and output qubits.

We say that $A = \{A_\lambda\}_\lambda$ is a *quantum polynomial-time (QPT) algorithm* if there exists a polynomial p such that the size of each circuit A_λ is at most $p(\lambda)$. Furthermore we say that A is *uniform* if there exists a deterministic polynomial-time Turing machine M that on input 1^n outputs the description of A_λ.

We also define the notion of a *non-uniform* QPT algorithm A that consists of a family $\{(A_\lambda, \rho_\lambda)\}_\lambda$ where $\{A_\lambda\}_\lambda$ is a polynomial-size family of circuits (not necessarily uniformly generated), and for each λ there is additionally a subset of input qubits of A_λ that are designated to be initialized with the density matrix ρ_λ of polynomial length. This is intended to model non-uniform quantum adversaries who may receive quantum states as advice. Nevertheless, the reductions we show in this work are all uniform.

The notation we use to describe the inputs/outputs of quantum algorithms will largely mimic what is used in the classical cryptography literature. For example, for a state generator algorithm G, we write $G_\lambda(k)$ to denote running the generalized quantum circuit G_λ on input $|k\rangle\langle k|$, which outputs a state ρ_k.

Ultimately, all inputs to a quantum circuit are density matrices. However, we mix-and-match between classical, pure state, and density matrix notation; for example, we may write $A_\lambda(k, |\theta\rangle, \rho)$ to denote running the circuit A_λ on input $|k\rangle\langle k| \otimes |\theta\rangle\langle\theta| \otimes \rho$. In general, we will not explain all the input and output sizes of every quantum circuit in excruciating detail; we will implicitly assume that a quantum circuit in question has the appropriate number of input and output qubits as required by context.

3.3 Pseudorandomness Notions

Next, we recall the different notions of pseudorandomness. First, in Section 3.3.1, we recall (classical) pseudorandom functions (prfs) and consider two notions of security associated with it. Then in Section 3.3.2, we define pseudorandom quantum state (PRS) generators, which are a quantum analogue of pseudorandom generators (PRGs). Finally in Section 3.3.3, we define pseudorandom function-like quantum state (PRFS) generators, which are a quantum analogue of pseudorandom functions. To make it less confusing to the reader, we use the abbreviation "prfs" (small letters) for classical pseudorandom functions and "PRFS" (all caps) for pseudorandom function-like states.

3.3.1 Pseudorandom Functions

We present two security notions of pseudorandom functions. First, we consider the notion of post-quantum security, defined below.

Definition 3.5 (Post-quantum pseudorandom functions)**.** *We say that a deterministic polynomial-time algorithm $F : \{0,1\}^\lambda \times \{0,1\}^{d(\lambda)} \to \{0,1\}^{n(\lambda)}$ is a* post-quantum secure pseudorandom function (pq-prf) *if for all QPT (non-uniform) distinguishers $A = (A_\lambda, \rho_\lambda)$ there exists a negligible function $\varepsilon(\cdot)$ such that the following holds:*

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda^{\mathcal{O}_{\mathsf{prf}}(k,\cdot)}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\mathsf{Rand}}} \left[A_\lambda^{\mathcal{O}_{\mathsf{Rand}}(\cdot)}(\rho_\lambda) = 1 \right] \right| \leq \varepsilon(\lambda),$$

where:

- $\mathcal{O}_{\mathsf{prf}}(k,\cdot)$*, modeled as a classical algorithm, on input $x \in \{0,1\}^{d(\lambda)}$, outputs $F(k,x)$.*
- $\mathcal{O}_{\mathsf{Rand}}(\cdot)$*, modeled as a classical algorithm, on input $x \in \{0,1\}^{d(\lambda)}$, outputs y_x, where $y_x \leftarrow \{0,1\}^{n(\lambda)}$.*

Moreover, the adversary A_λ only has classical access to $\mathcal{O}_{\mathsf{prf}}(k,\cdot)$ and $\mathcal{O}_{\mathsf{Rand}}(\cdot)$. That is, any query made to the oracle is measured in the computational basis.

We also say that F is a $(d(\lambda), n(\lambda))$-pq-prf to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$.

Next, we consider the quantum-query security, as considered by Zhandry [Zha12a]. In this security notion, the adversary has superposition access to either $\mathcal{O}_{\mathsf{prf}}$ or $\mathcal{O}_{\mathsf{Rand}}$. By definition, quantum-query security implies post-quantum security.

Unlike all the other pseudorandom notions considered in this section, we are going to use a different convention and allow the key length to be a polynomial in λ, instead of it being just λ. We also parameterize the advantage of the adversary. The motivation behind these changes in the definition will become clear in Section 5.3.

Definition 3.6 (Quantum-query secure pseudorandom functions)**.** *We say that a deterministic polynomial-time algorithm $F : \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{d(\lambda)} \to \{0,1\}^{n(\lambda)}$ is a* quantum-query ε-secure pseudorandom function (qprf) *if for all QPT (non-uniform) distinguishers $A = (A_\lambda, \rho_\lambda)$ there exists a function $\varepsilon(\cdot)$ such that the following holds:*

$$\left| \Pr_{k \leftarrow \{0,1\}^{\ell(\lambda)}} \left[A_\lambda^{|\mathcal{O}_{\mathsf{prf}}(k,\cdot)\rangle}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\mathsf{Rand}}} \left[A_\lambda^{|\mathcal{O}_{\mathsf{Rand}}(\cdot)\rangle}(\rho_\lambda) = 1 \right] \right| \leq \varepsilon(\lambda),$$

where:

- *$\mathcal{O}_{\mathsf{prf}}(k,\cdot)$ on input a $(d+n)$-qubit state on registers $\mathbf{X}$ (first d qubits) and $\mathbf{Y}$, applies an $(n+d)$-qubit unitary U described as follows: $U\left|x\right\rangle\left|a\right\rangle = \left|x\right\rangle\left|a \oplus F(k,x)\right\rangle$. It sends back the registers $\mathbf{X}$ and $\mathbf{Y}$.*
- *$\mathcal{O}_{\mathsf{Rand}}(\cdot)$ on input a $(d+n)$-qubit state on registers $\mathbf{X}$ (first d qubits) and $\mathbf{Y}$, applies an $(n+d)$-qubit unitary R described as follows: $R\left|x\right\rangle\left|a\right\rangle = \left|x\right\rangle\left|a \oplus y_x\right\rangle$, where $y_x \leftarrow \{0,1\}^{n(\lambda)}$. It sends back the registers $\mathbf{X}$ and $\mathbf{Y}$.*

Moreover, A_λ has superposition access to $\mathcal{O}_{\mathsf{prf}}(k,\cdot)$ and $\mathcal{O}_{\mathsf{Rand}}(\cdot)$. We denote the fact that A_λ has quantum access to an oracle $\mathcal{O}$ by $A_\lambda^{|\mathcal{O}\rangle}$.

We also say that F is a $(\ell(\lambda), d(\lambda), n(\lambda), \varepsilon)$-qprf to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$. When $\ell(\lambda) = \lambda$, we drop $\ell(\lambda)$ from the notation. Similarly, when $\varepsilon(\lambda)$ can be any negligible function, we drop $\varepsilon(\lambda)$ from the notation.

Zhandry [Zha12a] presented a construction of quantum-query secure pseudorandom functions from one-way functions.

Lemma 3.7 (Zhandry [Zha12a])**.** *Assuming post-quantum one-way functions, there exists quantum-query secure pseudorandom functions.*

Useful Lemma. We will use the following lemma due to Zhandry [Zha12b]. The lemma states that any q-query algorithm cannot distinguish (quantum) oracle access to a random function versus a $2q$-wise independent hash function. We restate the lemma using our notation.

Lemma 3.8 ([Zha12b, Theorem 3.1])**.** *Let A be a q-query algorithm. Then, for any $d, n \in \mathbb{N}$, every $2q$-wise independent hash function $H : \{0,1\}^{\ell(q)} \times \{0,1\}^d \to \{0,1\}^n$ satisfies the following:*

$$\left| \Pr_{k \leftarrow \{0,1\}^{\ell(q)}} \left[A_\lambda^{|\mathcal{O}_{\mathsf{H}}(k,\cdot)\rangle}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\mathsf{Rand}}} \left[A_\lambda^{|\mathcal{O}_{\mathsf{Rand}}(\cdot)\rangle}(\rho_\lambda) = 1 \right] \right| = 0,$$

where $\mathcal{O}_{\mathsf{Rand}}$ is as defined in Definition 3.6 and $\mathcal{O}_{\mathsf{H}}$ is defined similarly to $\mathcal{O}_{\mathsf{prf}}$ except that the unitary U uses H instead of F.

3.3.2 Pseudorandom Quantum State Generators

We move onto the pseudorandom notions in the quantum world. The notion of pseudorandom states were first introduced by Ji, Liu, and Song in [JLS18]. We reproduce their definition here:

Definition 3.9 (PRS Generator [JLS18])**.** *We say that a QPT algorithm G is a* pseudorandom state (PRS) generator *if the following holds.*

1. ***State Generation**. For all λ and for all $k \in \{0,1\}^\lambda$, the algorithm G behaves as*

   $$G_\lambda(k) = |\psi_k\rangle\langle\psi_k| .$$

   for some $n(\lambda)$-qubit pure state $|\psi_k\rangle$.

2. ***Pseudorandomness**. For all polynomials $t(\cdot)$ and QPT (nonuniform) distinguisher A there exists a negligible function $\varepsilon(\cdot)$ such that for all λ, we have*

   $$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda(G_\lambda(k)^{\otimes t(\lambda)}) = 1 \right] - \Pr_{|\vartheta\rangle \leftarrow \mathscr{H}_{n(\lambda)}} \left[A_\lambda(|\vartheta\rangle^{\otimes t(\lambda)}) = 1 \right] \right| \leq \varepsilon(\lambda) .$$

We also say that G is a $n(\lambda)$-PRS generator to succinctly indicate that the output length of G is $n(\lambda)$.

Ji, Liu, and Song showed that post-quantum one-way functions can be used to construct PRS generators.

Theorem 3.10 ([JLS18, BS20])**.** *If post-quantum one-way functions exist, then there exist PRS generators for all polynomial output lengths.*

3.3.3 Pseudorandom Function-Like State (PRFS) Generators

In this section, we recall the definition of pseudorandom function-like state (PRFS) generators by Ananth, Qian and Yuen [AQY21]. PRFS generators generalize PRS generators in two ways: first, in addition to the secret key k, the PRFS generator additionally takes a (classical) input x. The second way in which this definition generalizes the definition of PRS generators is that the output of the generator need not be a pure state.

However, they considered the weaker selective security definition (stated below) where the adversary needs to choose all the inputs to be queried to the PRFS ahead of time. Later we will introduce the stronger and the more useful definition of adaptive security.

Definition 3.11 (Selectively Secure PRFS generator)**.** *We say that a QPT algorithm G is a (selectively secure)* pseudorandom function-like state (PRFS) generator *if for all polynomials $s(\cdot), t(\cdot)$, QPT (nonuniform) distinguishers A and a family of indices $\left(\{x_1, \ldots, x_{s(\lambda)}\} \subseteq \{0,1\}^{d(\lambda)}\right)_\lambda$, there exists a negligible function $\varepsilon(\cdot)$ such that for all λ,*

$$\Bigg| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda(x_1, \ldots, x_{s(\lambda)}, G_\lambda(k, x_1)^{\otimes t(\lambda)}, \ldots, G_\lambda(k, x_{s(\lambda)})^{\otimes t(\lambda)}) = 1 \right] \\ - \Pr_{|\vartheta_1\rangle, \ldots, |\vartheta_{s(\lambda)}\rangle \leftarrow \mathscr{H}_{n(\lambda)}} \left[A_\lambda(x_1, \ldots, x_{s(\lambda)}, |\vartheta_1\rangle^{\otimes t(\lambda)}, \ldots, |\vartheta_{s(\lambda)}\rangle^{\otimes t(\lambda)}) = 1 \right] \Bigg| \leq \varepsilon(\lambda) \ .$$

We say that G is a $(d(\lambda), n(\lambda))$-PRFS generator to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$.

Our notion of security here can be seen as a version of *(classical) selective security*, where the queries to the PRFS generator are fixed before the key is sampled.

State Generation Guarantees. Towards capturing a natural class of PRFS generators, [AQY21] introduced the concept of *recognizable abort*. At a high level, recognizable abort is the property that the output of PRFS can be written as a convex combination of a pure state and a known abort state, denoted by $|\bot\rangle$. In more detail, the PRFS generator works in two stages. In the first stage it either generates a valid PRFS state $|\psi\rangle$ or it aborts. If it outputs a valid PRFS state then the first qubit is set to $|0\rangle$ and if it aborts, the entire state is set to $|\bot\rangle$. We have the guarantee that $|0\rangle|\psi\rangle$ is orthogonal to $|\bot\rangle$. In the next stage, the PRFS generator traces out the first qubit and outputs the resulting state. Our definition could be useful to capture many generators that don't always succeed in generating the pseudorandom state; for example, Brakerski and Shmueli [BS20] design generators that doesn't always succeed in generating the state.

We formally define the notion of recognizable abort[6] below.

[6] We note that [AQY21] define a slightly weaker definition of recognizable abort. However, the definitions and results considered in [AQY21] also work with our (stronger) definition of recognizable abort.

Definition 3.12 (Recognizable abort)**.** *A $(d(\lambda), n(\lambda))$-PRFS generator G has the* strongly recognizable abort property *if there exists an algorithm $\widehat{G}$ and a special $(n(\lambda)+1)$-qubit state $|\bot\rangle$ such that $G_\lambda(k,x)$ has the following form: it takes as input $k \in \{0,1\}^\lambda$, $x \in \{0,1\}^{d(\lambda)}$ and does the following,*

- *Compute $\widehat{G}_\lambda(k,x)$ to obtain an output of the form $\eta\left(|0\rangle\langle 0| \otimes |\psi\rangle\langle\psi|\right) + (1-\eta)\,|\bot\rangle\langle\bot|$ and moreover, $|\bot\rangle$ is of the form $|1\rangle\,|\widehat{\bot}\rangle$ for some $n(\lambda)$-qubit state state $|\widehat{\bot}\rangle$. As a consequence, $(\langle 0| \otimes \langle\psi|)(|\bot\rangle) = 0$.*
- *Trace out the first bit of $\widehat{G}_\lambda(k,x)$ and output the resulting state.*

As observed by [AQY21], the definition alone does not have any constraint on η being close to 1. The security guarantee of a PRFS generator implies that η will be negligibly close to 1 with overwhelming probability over the choice of k [AQY21, Lemma 3.6].

4 Adaptive Security

The previous work by [AQY21] only considers PRFS that is selectively secure. That is, the adversary needs to declare the input queries ahead of time. For many applications, selective security is insufficient. For example, in the application of PRFS to secret-key encryption (satisfying multi-message security), the resulting scheme was also only proven to be selectively secure, whereas one could ask for security against adversaries that can make *adaptive* queries to the PRFS generator. Another drawback of the notion considered by [AQY21] is the assumption that the adversary can make classical queries to the challenger who either returns PRFS states or independent Haar random states, whereas one would ideally prefer security against adversaries that can make *quantum superposition* queries.

In this work, we consider stronger notions of security for PRFS. We strengthen the definitions of [AQY21] in two ways. First, we allow the the adversary to make adaptive queries to the PRFS oracle, and second, we allow the adversary to make *quantum* queries to the oracle. The oracle model we consider here is slightly different from the usual quantum query model. In the usual model, there is an underlying function f and the oracle is modelled as a unitary acting on two registers, a *query* register $\mathbf{X}$ and an *answer* register $\mathbf{Y}$ mapping basis states $|x\rangle_{\mathbf{X}} \otimes |y\rangle_{\mathbf{Y}}$ to $|x\rangle_{\mathbf{X}} \otimes |y \oplus f(x)\rangle_{\mathbf{Y}}$ (in other words, the function output is XORed with answer register in the standard basis). The query algorithm also acts on the query and answer registers; indeed, it is often useful in quantum algorithms to initialize the answer register to something other than all zeroes.

In the PRS/PRFS setting, however, there is no underlying classical function: the output of the PRFS generator G could be an entangled pseudorandom state far from any standard basis state; it seems unnatural to XOR the pseudorandom the state with a standard basis state. Instead we consider a model where the query algorithm submits a query register $\mathbf{X}$ to the oracle, and the oracle returns the query register $\mathbf{X}$ as well as an answer register $\mathbf{Y}$. If the algorithm submits query $|x\rangle_{\mathbf{X}}$, then the joint state register $\mathbf{XY}$ after the query is $|x\rangle_{\mathbf{X}} \otimes |\psi_x\rangle_{\mathbf{Y}}$ for some pure state $|\psi_x\rangle$. Each time the algorithm makes a query, the oracle returns a fresh answer register. Thus, the number of qubits that the query algorithm acts on grows with the number of queries.[7]

[7] Alternatively, one can think of answer registers $\mathbf{Y}_1, \mathbf{Y}_2, \ldots$ as being initialized in the zeroes state at the beginning, and the query algorithm is only allowed to act nontrivially on $\mathbf{Y}_i$ after the i'th query.

How the oracle behaves when the query algorithm submits a superposition $\sum_x \alpha_x |x\rangle_{\mathbf{X}}$ in the query register is a further modeling choice. In the most general setting, the oracle behaves as a unitary on registers $\mathbf{XY}$,[8] and the resulting state of the query and answer registers is $\sum_x \alpha_x |x\rangle_{\mathbf{X}} \otimes |\psi_x\rangle_{\mathbf{Y}}$. That is, queries are answered in superposition. We call such an oracle *quantum-accessible*.

We also consider the case where the queries are forced to be *classical*, which may already be useful for some applications. Here, the oracle is modeled as a channel (instead of a unitary) that first measures the query register in the standard basis before returning the corresponding state $|\psi_x\rangle$. In other words, if the query is $\sum_x \alpha_x |x\rangle_{\mathbf{X}}$, then the resulting state becomes the mixed state $\sum_x |\alpha_x|^2 |x\rangle\langle x|_{\mathbf{X}} \otimes |\psi_x\rangle\langle\psi_x|_{\mathbf{Y}}$. This way, the algorithm cannot take advantage of quantum queries – but it can still make queries adaptively. We call such an oracle *classically-accessible*.

To distinguish between classical and quantum access to oracles, we write $A^{\mathcal{O}}$ to denote a quantum algorithm that has classical access to the oracle $\mathcal{O}$, and $A^{|\mathcal{O}\rangle}$ to denote a quantum algorithm that has quantum access to the oracle $\mathcal{O}$.

4.1 Classical Access

We define adaptively secure PRFS, where the adversary is given *classical access* to the PRFS/Haar-random oracle.

Definition 4.1 (Adaptively-Secure PRFS)**.** *We say that a QPT algorithm G is an* adaptively secure pseudorandom function-like state (APRFS) *generator if for all QPT (non-uniform) distinguishers A, there exists a negligible function ε, such that for all λ, the following holds:*

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda^{\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\mathsf{Haar}}} \left[A_\lambda^{\mathcal{O}_{\mathsf{Haar}}(\cdot)}(\rho_\lambda) = 1 \right] \right| \leq \varepsilon(\lambda),$$

where:

- $\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)$*, on input* $x \in \{0,1\}^{d(\lambda)}$*, outputs* $G_\lambda(k,x)$*.*
- $\mathcal{O}_{\mathsf{Haar}}(\cdot)$*, on input* $x \in \{0,1\}^{d(\lambda)}$*, outputs* $|\vartheta_x\rangle$*, where, for every* $y \in \{0,1\}^{d(\lambda)}$*,* $|\vartheta_y\rangle \leftarrow \mathscr{H}_{n(\lambda)}$*.*

Moreover, the adversary A_λ has classical access to $\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)$ and $\mathcal{O}_{\mathsf{Haar}}(\cdot)$. That is, we can assume without loss of generality that any query made to either oracle is measured in the computational basis.

We say that G is a $(d(\lambda), n(\lambda))$-APRFS generator to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$.

Some remarks are in order.

Instantiation. For the case when $d(\lambda) = O(\log(\lambda))$, selectively secure PRFS is equivalent to adaptively secure PRFS. The reason being that we can assume without loss of generality, the selective adversary can query on all possible inputs (there are only polynomially many) and use the outputs to simulate the adaptive adversary. As a consequence of the result that log-input selectively-secure PRFS can be built from PRS [AQY21], we obtain the following.

Lemma 4.2. *For $d = O(\log(\lambda))$ and $n = d + \omega(\log\log\lambda)$, assuming the existence of $(d+n)$-PRS, there exists a (d,n)-APRFS.*

In the case when $d(\lambda)$ is an arbitrary polynomial in λ, we present a construction of APRFS from post-quantum one-way functions in Section 5.2.

[8] Alternatively, one can think of the oracle as an *isometry* mapping register $\mathbf{X}$ to registers $\mathbf{XY}$.

Test procedure. It was shown by [AQY21] that a PRFS admits a Test procedure (See Section 3.3 in [AQY21]). The goal of a Test procedure is to determine whether the given state is a valid PRFS state or not. Having a Test procedure is useful in applications. For example, [AQY21] used a Test procedure in the construction of a bit commitment scheme. We note that the same Test procedure also works for adaptively secure PRFS.

Multiple copies. In the definition of PRS (Definition 3.9) and selectively-secure PRFS (Definition 3.11), the adversary is allowed to obtain multiple copies of the same pseudorandom (or haar random) quantum state. While we do not explicitly state it, even in Definition 4.1, the adversary can indeed obtain multiple copies of a (pseudorandom or haar random) quantum state. To obtain t copies of the output of $G_\lambda(k, x)$ (or $|\vartheta_x\rangle$), the adversary can query the same input x, t times, to the oracle $\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$ (or $\mathcal{O}_{\mathsf{Haar}}(\cdot)$).

4.2 Quantum Access

We further strengthen our notion of adaptively secure PRFS by allowing the adversary to make superposition queries to either $\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$ or $\mathcal{O}_{\mathsf{Haar}}(\cdot)$. Providing superposition access to the adversary not only makes the definition stronger[9] than Definition 4.1 but is also arguably more useful for a larger class of applications. To indicate quantum query access, we put the oracle inside the ket notation: $A^{|\mathcal{O}\rangle}$ (whereas for classical query access we write $A^{\mathcal{O}}$).

We provide the formal definition below.

Definition 4.3 (Quantum-accessible Adaptively-secure PRFS)**.** *We say that a QPT algorithm G is a quantum-accessible adaptively secure pseudorandom function-like state (QAPRFS) generator if for all QPT (non-uniform) distinguishers A if there exists a negligible function ε, such that for all λ, the following holds:*

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda^{|\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)\rangle}(\rho_\lambda) = 1 \right] - \Pr_{\mathcal{O}_{\mathsf{Haar}}} \left[A_\lambda^{|\mathcal{O}_{\mathsf{Haar}}(\cdot)\rangle}(\rho_\lambda) = 1 \right] \right| \leq \varepsilon(\lambda),$$

where:

- *$\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$, on input a d-qubit register $\mathbf{X}$, does the following: it applies a channel that controlled on the register $\mathbf{X}$ containing x, it creates and stores $G_\lambda(k, x)$ in a new register $\mathbf{Y}$. It outputs the state on the registers $\mathbf{X}$ and $\mathbf{Y}$.*
- *$\mathcal{O}_{\mathsf{Haar}}(\cdot)$, modeled as a channel, on input a d-qubit register $\mathbf{X}$, does the following: it applies a channel that controlled on the register $\mathbf{X}$ containing x, stores $|\vartheta_x\rangle\langle\vartheta_x|$ in a new register $\mathbf{Y}$, where $|\vartheta_x\rangle$ is sampled from the Haar distribution. It outputs the state on the registers $\mathbf{X}$ and $\mathbf{Y}$.*

Moreover, A_λ has superposition access to $\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$ and $\mathcal{O}_{\mathsf{Haar}}(\cdot)$.

We say that G is a $(d(\lambda), n(\lambda))$-QAPRFS generator to succinctly indicate that its input length is $d(\lambda)$ and its output length is $n(\lambda)$.

[9] It is stronger in the sense that an algorithm that has quantum query access to the oracle can simulate an algorithm that only has classical query access.

We present a construction satisfying the above definition in Section 5.3.

Unlike Definition 4.1, it is not without loss of generality that A_λ can get multiple copies of a quantum state. To illustrate, consider an adversary that submits a state of the form $\sum_x \alpha_x |x\rangle$ to the oracle. It then gets back $\sum_x \alpha_x |x\rangle |\psi_x\rangle$ (where $|\psi_x\rangle$ is either the output of PRFS[10] or it is Haar random) instead of $\sum_x \alpha_x |x\rangle |\psi_x\rangle^{\otimes t}$, for some polynomial t. On the other hand, if the adversary can create multiple copies of $\sum_x \alpha_x |x\rangle$, the above definition allows the adversary to obtain $(\sum_x \alpha_x |x\rangle |\psi_x\rangle)^{\otimes t}$ for any polynomial $t(\cdot)$ of its choice.

5 Constructions of Pseudorandom Primitives

We present improved and/or new constructions of pseudorandom primitives.

1. In Section 5.1, we present a simpler proof of binary phase PRS. The current known proof of binary phase PRS by Brakerski and Shmueli [BS19] is arguably more involved.

2. In Section 5.2, we present a construction of APRFS (Definition 4.1) from post-quantum one-way functions. Recall that in this definition, the adversary only has classical access to the oracle.

3. In Section 5.3, we present two constructions of QAPRFS (Definition 4.3) from post-quantum one-way functions. Recall that in this definition, the adversary has quantum access to the oracle.

5.1 Simpler Analysis of Binary Phase PRS

In this section we give a simpler analysis of the binary phase PRS construction suggested by [JLS18] and later analyzed by [BS19].

Theorem 5.1. *Let $F : \{0,1\}^\lambda \times \{0,1\}^n \to \{0,1\}$ is a quantum-query secure PRF (Definition 3.6), then $G : \{0,1\}^\lambda \to \mathbb{C}^{2^n}$ defined by $G(k) = |\psi_k\rangle = 2^{-n/2} \sum_x (-1)^{F(k,x)} |x\rangle$ is a n-qubit PRS generator.*

To prove this, we establish some basic facts about the *symmetric subspace*.

The Symmetric Subspace and Its Properties. The symmetric subspace of the tensor product space $(\mathbb{C}^N)^{\otimes t}$ is the subspace of states invariant under all permutations of the t tensor factors. We write $\Pi^{N,t}_{\mathsf{sym}}$ to denote the projector onto the symmetric subspace of $(\mathbb{C}^N)^{\otimes t}$. The next fact gives an equivalent definition of the projector. For proofs of these facts, please see [Har13].

Fact 5.2. *For a permutation $\sigma \in S_t$, let $P_N(\sigma)$ denote the permutation on $(\mathbb{C}^N)^{\otimes t}$ that permutes the t tensor factors according to σ. Hence,*

$$P_N(\sigma) = \sum_{x_1,\ldots,x_t \in [N]} |x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(t)}\rangle\langle x_1, \ldots, x_t| \,.$$

Then we have

$$\Pi^{N,t}_{\mathsf{sym}} = \frac{1}{t!} \sum_{\sigma \in S_t} P_N(\sigma) \ .$$

[10] In this illustration, we are pretending that the PRFS satisfies perfect state generation property. That is, the output of PRFS is always a pure state.

Fact 5.3 (Average of copies of Haar-random states)**.** *For all $N, t \in \mathbb{N}$, we have*

$$\mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathscr{H}(\mathbb{C}^N)} |\vartheta\rangle\langle\vartheta|^{\otimes t} = \frac{\Pi_{\mathsf{sym}}^{N,t}}{\mathrm{Tr}(\Pi_{\mathsf{sym}}^{N,t})} \ .$$

Fact 5.4. *Let ρ_1, ρ_2 be density matrices such that $\rho_2 = \alpha\rho_1 + \beta\rho_1^{\perp}$ where $\rho_1\rho_1^{\perp} = 0$, $\rho_1^{\perp}\rho_1 = 0$ and $\alpha, \beta \in [0,1]$, $\alpha + \beta = 1$, then*

$$\mathrm{TD}\left(\rho_1, \rho_2\right) = \beta.$$

Proof.

$$\begin{aligned}
\mathrm{TD}\left(\rho_1, \rho_2\right) &= \frac{1}{2}\mathrm{Tr}\left(\sqrt{\left(\rho_1 - \rho_2\right)^2}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\sqrt{\left(\rho_1 - (\alpha\rho_1 + \beta\rho_1^{\perp})\right)^2}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\sqrt{\left((1-\alpha)\rho_1 - \beta\rho_1^{\perp}\right)^2}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\sqrt{(1-\alpha)^2\rho_1^2 + \beta^2\left(\rho_1^{\perp}\right)^2}\right)
\end{aligned}$$

Here, the first equality is from definition of trace distance, the second equality is by definition of ρ_2, the third equality is by simplification, and the fourth equality is because $\rho_1\rho_1^{\perp} = 0$, $\rho_1^{\perp}\rho_1 = 0$.

Since, $\rho_1\rho_1^{\perp} = 0$, $\rho_1^{\perp}\rho_1 = 0$, we can add $(1-\alpha)\beta\rho_1\rho_1^{\perp} + (1-\alpha)\beta\rho_1^{\perp}\rho_1$ without changing the value.

$$\begin{aligned}
\mathrm{TD}\left(\rho_1, \rho_2\right) &= \frac{1}{2}\mathrm{Tr}\left(\sqrt{(1-\alpha)^2\rho_1^2 + \beta^2\left(\rho_1^{\perp}\right)^2 + (1-\alpha)\beta\rho_1\rho_1^{\perp} + (1-\alpha)\beta\rho_1^{\perp}\rho_1}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left(\sqrt{\left((1-\alpha)\rho_1 + \beta\rho_1^{\perp}\right)^2}\right) \\
&= \frac{1}{2}\mathrm{Tr}\left((1-\alpha)\rho_1 + \beta\rho_1^{\perp}\right) \\
&= \frac{1}{2}\left((1-\alpha)\mathrm{Tr}(\rho_1) + \beta\mathrm{Tr}\left(\rho_1^{\perp}\right)\right) \\
&= \frac{1}{2}\left((1-\alpha) + \beta\right) \\
&= \beta
\end{aligned}$$

Here, the first line is since $\rho_1\rho_1^{\perp} = 0$, $\rho_1^{\perp}\rho_1 = 0$, the second line is by simplification, the third equality is since $(1-\alpha)\rho_1 + \beta\rho_1^{\perp}$ is positive semidefinite, the fourth equality is by linearity of trace, the fifth equality is by $\mathrm{Tr}(\rho) = 1$ and $\mathrm{Tr}(\rho^{\perp}) = 1$, and the last equality is by $\alpha + \beta = 1$. □

We define some new notation to look at a different charecterisation of the symmetric space.

Definition 5.5. *Let $v \in [N]^t$ for some $N, t \in \mathbb{N}$, then define type(v) to be a vector in $[t+1]^N$ where the i^{th} entry in type(v) denotes the frequency of i in v.*

Definition 5.6. *Let $T \in [t+1]^N$ for some $N, t \in \mathbb{N}$, then define*

$$|type_T\rangle = \beta \sum_{\substack{v \in [N]^t \\ type(v) = T}} |v\rangle \,,$$

where $\beta \in \mathbb{R}$ is an appropriately chosen constant. Similarly, for $T \in \{0,1\}^N$, define

$$|bintype_T\rangle = \beta \sum_{\substack{v \in [N]^t \\ type(v) \pmod 2 = T}} |v\rangle \,,$$

where $\beta \in \mathbb{R}$ is an appropriately chosen constant.

Note that if $hamming(T) = t$ and $T \in \{0,1\}^N$, then $|type_T\rangle = |bintype_T\rangle$.

Lemma 5.7. *For all $N, t \in \mathbb{N}$, we have*

$$\mathrm{TD}\left(\mathop{\mathbb{E}}_{\substack{T \leftarrow \{0,1\}^N \\ hamming(T) = t}} |type_T\rangle\langle type_T| \,, \mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathscr{H}(\mathbb{C}^N)} |\vartheta\rangle\langle\vartheta|^{\otimes t} \right) \leq O\left(\frac{t^2}{N}\right) .$$

Proof. From Fact 5.3,

$$\begin{aligned}
\rho = \mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathscr{H}(\mathbb{C}^N)} |\vartheta\rangle\langle\vartheta|^{\otimes t} &= \frac{\Pi_{\mathsf{sym}}^{N,t}}{\mathrm{Tr}(\Pi_{\mathsf{sym}}^{N,t})} \\
&= \frac{1}{t!\mathrm{Tr}(\Pi_{\mathsf{sym}}^{N,t})} \sum_{\sigma \in S_t} P_N(\sigma) \\
&= \frac{1}{t!\mathrm{Tr}(\Pi_{\mathsf{sym}}^{N,t})} \sum_{\sigma \in S_t} \sum_{x_1, \ldots, x_t \in [N]} |x_{\sigma^{-1}(1)}, \ldots, x_{\sigma^{-1}(t)}\rangle\langle x_1, \ldots, x_t|
\end{aligned}$$

where the second and the third line follow from Fact 5.2. From Definition 5.6,

$$\begin{aligned}
\sigma = \mathop{\mathbb{E}}_{\substack{T\leftarrow\{0,1\}^N\\ hamming(T)=t}} |type_T\rangle\langle type_T| &= \mathop{\mathbb{E}}_{\substack{T\leftarrow\{0,1\}^N\\ hamming(T)=t}} \left(\frac{1}{\sqrt{t!}}\sum_{\substack{v\in[N]^t\\ type(v)=T}} |v\rangle\right)\left(\frac{1}{\sqrt{t!}}\sum_{\substack{v'\in[N]^t\\ type(v')=T}} \langle v'|\right)\\
&= \frac{1}{t!}\mathop{\mathbb{E}}_{\substack{T\leftarrow\{0,1\}^N\\ hamming(T)=t}} \left(\sum_{\substack{v,v'\in[N]^t\\ type(v)=type(v')=T}} |v\rangle\langle v'|\right)\\
&= \frac{1}{t!}\mathop{\mathbb{E}}_{\substack{T\leftarrow\{0,1\}^N\\ hamming(T)=t}} \left(\sum_{\substack{v\in[N]^t\\ type(v)=T}} \sum_{\substack{\sigma\in S_t\\ v'=\sigma(v)}} |v\rangle\langle v'|\right)\\
&= \frac{1}{t!\binom{N}{t}} \sum_{\substack{v\in[N]^t\\ type(v)\in\{0,1\}^N}} \sum_{\sigma\in S_t} |v\rangle\langle\sigma(v)|\\
&= \frac{1}{t!\binom{N}{t}} \sum_{\substack{x_1,\ldots,x_t\in[N]\\ x_1,\ldots,x_t \text{ are distinct}}} \sum_{\sigma\in S_t} |x_1,\ldots,x_t\rangle\langle x_{\sigma(1)},\ldots,x_{\sigma(t)}|
\end{aligned}$$

where the third line follows by re-interpreting vector of same type as permutation of each other, the fourth line follows by taking expectation (since there are a total of $\binom{N}{t}$ strings of hamming weight of t in $\{0,1\}^N$). The fifth line follows since $type(v)\in\{0,1\}^N$, hence all elements of v are distinct.

Define

$$\sigma^\perp = \frac{1}{t!\left(\binom{N+t-1}{t}-\binom{N}{t}\right)} \sum_{\substack{x_1,\ldots,x_t\in[N]\\ x_1,\ldots,x_t \text{ are not distinct}}} \sum_{\sigma\in S_t} |x_1,\ldots,x_t\rangle\langle x_{\sigma(1)},\ldots,x_{\sigma(t)}|$$

Hence, $\rho = \alpha\sigma + \beta\sigma^\perp$, where β is probability of picking $x_1,\ldots,x_t\in[N]$ such that there is a colision which is less than $O\left(\frac{t^2}{N}\right)$.

Hence, from Fact 5.4, we see

$$\mathrm{TD}\left(\mathop{\mathbb{E}}_{\substack{T\leftarrow\{0,1\}^N\\ hamming(T)=t}} |type_T\rangle\langle type_T|, \mathop{\mathbb{E}}_{|\vartheta\rangle\leftarrow\mathscr{H}(\mathbb{C}^N)} |\vartheta\rangle\langle\vartheta|^{\otimes t}\right) \le O\left(\frac{t^2}{N}\right).$$

□

We prove Theorem 5.1 via a hybrid argument. Let t be a polynomial in λ.

Hybrid 1. Sample a random $k\leftarrow\{0,1\}^\lambda$. Let $|\psi\rangle = |\psi_k\rangle$ as defined in Equation (1) and output $|\psi\rangle^{\otimes t}$.

Hybrid 2. For all $x \in \{0,1\}^n$, sample $\alpha_x \in \{\pm 1\}$ uniformly at random. Let

$$|\psi\rangle = 2^{-n/2} \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle \ .$$

Output $|\psi\rangle^{\otimes t}$.

Hybrid 3. Sample $w \in [N]^t$ uniformly at random. Let $T = \mathit{type}(w) \pmod 2$. Output $|\mathit{bintype}_T\rangle^{\otimes t}$.

Hybrid 4. Sample $T \in \{0,1\}^N$ with $\mathit{hamming}(T) = t$ uniformly at random. Output $|\mathit{type}_T\rangle^{\otimes t}$.

Hybrid 5. Sample a Haar-random n-qubit state $|\psi\rangle$. Output $|\psi\rangle^{\otimes t}$.

The Hybrids 1 and 2 are computationally indistinguishable since the PRF F cannot be distinguished by a quantum adversary (that can make superposition queries to $F_k(\cdot) = F(k,\cdot)$) from a random function, from the quantum query security in Definition 3.6.

The rest of the proof can be seen as the following lemma.

Lemma 5.8. *Fix $t, n \in \mathbb{N}$. Let $\{\alpha_x\}_x$ be independent and uniformly random ± 1 values. Define*

$$|\psi\rangle = 2^{-n/2} \sum_{x \in \{0,1\}^n} \alpha_x \ |x\rangle \ .$$

Then

$$\mathrm{TD}\left(\mathop{\mathbb{E}} |\psi\rangle\langle\psi|^{\otimes t} \ , \mathop{\mathbb{E}} |\vartheta\rangle\langle\vartheta|^{\otimes t} \ \right) \leq O(t^2/2^n) \tag{2}$$

where $|\vartheta\rangle$ is a Haar-random n-qubit state.

A version of this Lemma is also proven by Brakerski and Shmueli [BS19], we give an alternate and a more straightforward proof.

Lemma 5.9. *Hybrids 2 and 3 (from above) are identical.*

Proof. The output of Hybrid 2 is as follows:

$$\begin{aligned}
\rho &= \mathop{\mathbb{E}}_{\alpha_x} \left(\frac{1}{N^{t/2}} \sum_{x_1,\ldots,,x_t \in [N]} \alpha_{x_1} \ldots \alpha_{x_t} |x_1, \ldots, x_t\rangle \right) \left(\frac{1}{N^{t/2}} \sum_{y_1,\ldots,y_t \in [N]} \alpha_{y_1} \ldots \alpha_{y_t} \langle y_1, \ldots, y_t| \right) \\
&= \frac{1}{N^t} \mathop{\mathbb{E}}_{\alpha_x} \left(\sum_{\substack{x_1,\ldots,x_t \in [N] \\ y_1,\ldots,y_t \in [N]}} \alpha_{x_1} \ldots \alpha_{x_t} \alpha_{y_1} \ldots \alpha_{y_t} |x_1, \ldots, x_t\rangle\langle y_1, \ldots, y_t| \right) \\
&= \frac{1}{N^t} \left(\sum_{\substack{x_1,\ldots,x_t \in [N] \\ y_1,\ldots,y_t \in [N]}} \sum_{\mathit{type}(x) \bmod 2 = \mathit{type}(y) \bmod 2} |x_1, \ldots, x_t\rangle\langle y_1, \ldots, y_t| \right)
\end{aligned}$$

where the second line follows by rearranging terms and third line follows from the fact that taking expectation over any unpaired α_x results in zero.

The output of Hybrid 3 is as follows:

$$\begin{aligned}\rho' &= \mathop{\mathbb{E}}_{w\in[N]^t}\left(\sum_{\substack{v\in[N]^t\\ \mathit{type}(v) \bmod 2=\mathit{type}(w) \bmod 2}} |v\rangle\right)\left(\sum_{\substack{v'\in[N]^t\\ \mathit{type}(v') \bmod 2=\mathit{type}(w) \bmod 2}} \langle v'|\right)\\ &= \frac{1}{N^t}\left(\sum_{\substack{v,v'\in[N]^t\\ \mathit{type}(v) \bmod 2=\mathit{type}(v') \bmod 2}} |v\rangle\langle v'|\right)\end{aligned}$$

where the second line follows because there are N^t options for w. From above we can see that the output of Hybrid 2 and Hybrid 3 are exactly the same. □

Let $w \in [N]^t$ be sampled uniformly at random and $T = \mathit{type}(w) \pmod 2$, then $\mathit{hamming}(T) = t$ with probability $1 - t^2/N$ as w has no collisions with probabilty $1 - t^2/N$. Hence, the trace distance between the outputs of Hybrids 3 and 4 is $O(t^2/N)$.

By Lemma 5.7, the trace distance between the outputs of Hybrids 4 and 5 is $O(t^2/N)$.

Combining our bounds on above, we get that

$$\mathrm{TD}\left(|\psi\rangle\langle\psi|^{\otimes t}, \mathbb{E}\,|\vartheta\rangle\langle\vartheta|^{\otimes t}\right) \le O(t^2/N)$$

as desired.

5.2 (Classically-Accessible) Adaptively-Secure PRFS

We construct an adaptively secure PRFS where the adversary is only allowed classical access in the security experiment. While the quantum-accessible APRFS constructed and analyzed in Section 5.3 are also classically-accessible; their analyses are much more involved. Here we present a construction and analysis of a classically-accessible APRFS that is simpler, and may be a helpful starting point for applications that only require classically-accessible PRFS.

To construct a $(d(\lambda), n(\lambda))$-APRFS G, we start with the following two primitives:

- $(d(\lambda), \lambda)$-post-quantum secure pseudorandom function F (Definition 3.5) and,
- $n(\lambda)$-pseudorandom state generator g (Definition 3.9).

Construction. We describe G_λ as follows: on input $k \in \{0,1\}^\lambda$, $x \in \{0,1\}^{d(\lambda)}$,

- $k_x \leftarrow F(k, x)$,
- $\rho \leftarrow g(k_x)$,
- Output ρ.

Lemma 5.10. *Assuming the post-quantum security of the pseudorandom function F, the QPT procedure G described above is a secure $(d(\lambda), n(\lambda))$-APRFS (Definition 4.1).*

Proof. We prove this by a standard hybrid argument. Let A_λ be a QPT distinguisher that distinguishes the oracles $\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)$, where $k \leftarrow \{0,1\}^\lambda$, and $\mathcal{O}_{\mathsf{Haar}}(\cdot)$ with probability ε. We prove that ε is negligible. Let q be the number of queries made by A_λ. For simplicity, we assume below that all the queries are distinct. The same proof generalizes to the setting when the queries are not distinct.

Hybrid H_1. Output $A_\lambda^{\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$.

Hybrid H_2. Output $A_\lambda^{\mathcal{O}_{H_2}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$, where $\mathcal{O}_{H_2}(k,\cdot)$ is defined as follows.

On input $x \in \{0,1\}^{d(\lambda)}$, compute $\rho \leftarrow g(k_x)$, where k_x is sampled uniformly at random. Output ρ.

The computational indistinguishability of H_1 and H_2 follows from the post-quantum security of F.

Hybrid $H_{3.i}$ for $i \in [q]$. Output $A_\lambda^{\mathcal{O}_{H_{3.i}}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$, where $\mathcal{O}_{H_{3.i}}(k,\cdot)$ is defined as follows.

On input $x \in \{0,1\}^{d(\lambda)}$, do the following: if this is the j^{th} query, where $j \leq i$, then output $|\vartheta_x\rangle$, where $|\vartheta_x\rangle$ is sampled from the Haar distribution and if $j > i$, output $\rho \leftarrow g(k_x)$, where k_x is sampled uniformly at random.

Claim 1. *Assuming the security of PRS generator g, for every $i \in [q]$, A_λ can distinguish hybrids $H_{3.i-1}$ (we set $H_{3.0} = H_2$) and $H_{3.i}$ only with negligible probability.*

Proof. Suppose A_λ can distinguish the hybrids $H_{3.i-1}$ and $H_{3.i}$ with probability $\varepsilon_3(\lambda)$. We construct an adversary $\mathcal{B}$ that violates the security of g with probability $\varepsilon_3(\lambda)$. Then we invoke the security of g to conclude that $\varepsilon_3(\lambda)$ has to be negligible in λ. First, we let $\mathcal{B}$ be inefficient and later, we remark how to make $\mathcal{B}$ efficient.

$\mathcal{B}$ gets as input a state ρ. $\mathcal{B}$ then runs A_λ while simulating the oracle A_λ has access to. For the j^{th} query, say $x \in \{0,1\}^{d(\lambda)}$, where $j < i$, it outputs $|\vartheta_x\rangle\langle\vartheta_x|$, where $|\vartheta_x\rangle \leftarrow \mathscr{H}_{n(\lambda)}$. For the i^{th} query, it outputs ρ. For the j^{th} query, say $x \in \{0,1\}^{d(\lambda)}$, where $j > i$, it outputs ρ_{k_x}, where $\rho_{k_x} \leftarrow g(k_x)$ and $k_x \leftarrow \{0,1\}^\lambda$. The output of $\mathcal{B}$ is set to be the output of A_λ.

If $\rho \leftarrow g(k)$, for some $k \leftarrow \{0,1\}^\lambda$ then the output distribution of $\mathcal{B}$ is precisely the output distribution of $H_{3.i-1}$ and if $\rho = |\psi\rangle\langle\psi|$, where $|\psi\rangle \leftarrow \mathscr{H}_{n(\lambda)}$, then the output distribution of $\mathcal{B}$ is precisely the output distribution of $H_{3.i}$. Thus, the probability that A_λ distinguishes the hybrids $H_{3.i-1}$ and $H_{3.i}$ is precisely the distinguishing probability of $\mathcal{B}$.

Note that $\mathcal{B}$ is not efficient since it needs to sample Haar random states. Instead of sampling Haar random states, $\mathcal{B}$ instead uses $q(\lambda)$-state designs, where $q(\lambda)$ is the maximum number of queries A_λ can make. $q(\lambda)$-state designs can be efficiently generated (in time polynomial in $q(\lambda)$) [AE07, DCEL09].

Since $\mathcal{B}$ is QPT, $\varepsilon_3(\lambda)$, which is the probability that $\mathcal{B}$ distinguishes PRS from Haar random, has to be negligible in λ. □

Hybrid H_4. Output $A_\lambda^{\mathcal{O}_{\mathsf{Haar}}(\cdot)}$.

The hybrids $H_{3.q}$ and H_4 are identical.

Thus, it follows that A_λ can distinguish H_1 and H_4 at most with negligible probability. This completes the proof.

□

5.3 (Quantum-Accessible) Adaptively-Secure PRFS

We now focus on achieving adaptive quantum query security (Definition 4.3).

We present two constructions and two proofs. The first construction requires the existence of subexponentially-secure post-quantum one-way functions, and we use a technique called *complexity leveraging* to prove security of the construction. The second construction only requires the existence of polynomially-secure post-quantum one-way functions, but the analysis requires more sophisticated tools, namely an extension of Zhandry's small-range distribution technique to handle random unitary oracles. We believe that presenting two different ways to analyze PRFS may be helpful for future works.

5.3.1 First Construction

Our first construction is similar to the construction in Section 5.2 except that we use the binary phase PRS construction of [JLS18, BS19] (and which we analyze in Section 5.1).

To construct a $(d(\lambda), n(\lambda))$-APRFS G, where $d(\lambda) = \lambda^\gamma$ with $1 > \gamma > 0$ and $\frac{n(\lambda)}{2} - d(\lambda) = \omega(\log(\lambda))$, we start with two pseudorandom functions:

- $(d(\lambda), \ell(\lambda))$-post-quantum secure pseudorandom function F_1 (Definition 3.5),
- $(\ell(\lambda), n(\lambda), 1, \frac{1}{2^{\lambda^\delta}})$-quantum-query secure pseudorandom function F_2 (Definition 3.6), where $1 > \delta > \gamma > 0$.

We note that the same construction also works for any $d(\lambda) < \lambda^\gamma$.

Construction. We present the construction of G_λ in Figure 1.

Input: $k \in \{0,1\}^\lambda$, $x \in \{0,1\}^{d(\lambda)}$.

- Initialize $|x\rangle_{\mathbf{X}} |0\rangle_{\mathbf{K}}$.
- Apply a unitary V_{F_1} on $|x\rangle_{\mathbf{X}} |0\rangle_{\mathbf{K}}$, where V_{F_1} is defined as follows: $V_{F_1} |a\rangle |b\rangle = |a\rangle |F_1(k,a) \oplus b\rangle$. Let $k_x = F_1(k,x)$.
- Apply a unitary U_{F_2} on $|k_x\rangle_{\mathbf{K}} |0\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ to obtain $|k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$, where $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y \in \{0,1\}^{n(\lambda)}} \frac{(-1)^{F_2(k_x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$, and U_{F_2} is described as follows:
    - Apply $H^{\otimes n(\lambda)}$ on $\mathbf{Y}$.
    - Apply a unitary that maps $|a\rangle_{\mathbf{K}} |b\rangle_{\mathbf{Y}} |c\rangle_{\mathbf{Anc}}$ to $|a\rangle_{\mathbf{K}} |b\rangle_{\mathbf{Y}} |c \oplus F_2(a,b)\rangle_{\mathbf{Anc}}$.
- The resulting state is $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$. Trace out the register **Anc**.
- Apply the unitary V_{F_1} on the **X** and **K** registers (again) to obtain $|x\rangle_{\mathbf{X}} |0\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Anc}}$.
- Output $|\psi_{k_x}\rangle$.

Figure 1: Construction of G_λ.

We prove the security of the above construction in the lemma below.

Lemma 5.11. *Assuming the quantum-query $\nu(\lambda)$-security of the pseudorandom function F_1, where $\nu(\lambda)$ is a negligible function, the quantum-query $\frac{1}{2^{\lambda^\delta}}$-security of F_2, the QPT procedure G described above is a secure $(d(\lambda), n(\lambda))$-APRFS (Definition 4.1).*

Proof. We prove this by a standard hybrid argument. Let A_λ be a QPT distinguisher that distinguishes the oracles $\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$, where $k \leftarrow \{0,1\}^\lambda$, and $\mathcal{O}_{\mathsf{Haar}}$ with probability ε. We prove that ε is negligible. Let q be the number of queries made by A_λ.

Hybrid H_1. Output $A_\lambda^{\mathcal{O}_{\mathsf{PRFS}}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$. Observe that $\mathcal{O}_{\mathsf{PRFS}}(k, \cdot)$ can be implemented as follows.

On input a d-qubit register $\mathbf{X}$, do the following:

1. Initialize a register $\mathbf{K}$ with $|0\rangle$.
2. Apply the unitary V_{F_1} on the registers $\mathbf{X}$ and $\mathbf{K}$, where V_{F_1} is defined as follows: $V_{F_1} |x\rangle |a\rangle = |x\rangle |F_1(k, x) \oplus a\rangle$.
3. Initialize two registers $\mathbf{Y}$ and $\mathbf{Anc}$ with $|0\rangle$ and $|-\rangle$ respectively.
4. Apply the unitary U_{F_2} (described in the construction) on the registers $\mathbf{K}$, $\mathbf{Y}$ and $\mathbf{Anc}$.
5. Apply the unitary V_{F_1} on the registers $\mathbf{X}$ and $\mathbf{K}$.
6. Output the state in the registers $\mathbf{X}$ and $\mathbf{Y}$.

Hybrid H_2. Output $A_\lambda^{\mathcal{O}_{H_2}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$, where O_{H_2} is defined below.

Sample $\widehat{F_1} : \{0,1\}^{d(\lambda)} \rightarrow \{0,1\}^{\ell(\lambda)}$ uniformly at random. On input a d-qubit register $\mathbf{X}$, do the following:

1. Initialize a register $\mathbf{K}$ with $|0\rangle$.
2. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$, where $V_{\widehat{F_1}}$ is defined as follows: $V_{\widehat{F_1}} |x\rangle |a\rangle = |x\rangle |\widehat{F_1}(x) \oplus a\rangle$.
3. Initialize two registers $\mathbf{Y}$ and $\mathbf{Anc}$ with $|0\rangle$ and $|-\rangle$ respectively..
4. Apply the unitary U_{F_2} (described in the construction) on the registers $\mathbf{K}$, $\mathbf{Y}$ and $\mathbf{Anc}$.
5. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$.
6. Output the state in the registers $\mathbf{X}$ and $\mathbf{Y}$.

Claim 2. *Assuming the quantum-query security of F_1, A_λ can distinguish the hybrids H_1 and H_2 with probability at most $\nu(\lambda)$.*

Proof. Suppose A_λ can distinguish H_1 and H_2 with probability $> \nu(\lambda)$. We show that there exists a QPT adversary $\mathcal{B}$ that violates the quantum-query security of F_1 with probability $> \nu(\lambda)$.

$\mathcal{B}$ runs A_λ by simulating the oracle that A_λ has access to. For every query made by A_λ, $\mathcal{B}$ with oracle access to $\mathcal{O}$ (which is either $\mathcal{O}_{\mathsf{prf}}(k,\cdot)$ or $\mathcal{O}_{\mathsf{Rand}}$; Definition 3.6), does the following: it performs 6 steps, where steps 1,3,4 and 6 are the same as in H_1 (also identical to the steps 1,3,4 and 6 of H_2). We describe the steps 2 and 5 below.

2. Send the registers $\mathbf{X}$ and $\mathbf{K}$ to $\mathcal{O}$.

5. Send the registers $\mathbf{X}$ and $\mathbf{K}$ to $\mathcal{O}$.

If $\mathcal{B}$ has oracle access to $\mathcal{O}_{\mathsf{prf}}(k,\cdot)$ then the output distribution of A_λ is identical to the output distribution of A_λ in H_1. If $\mathcal{B}$ has oracle access to $\mathcal{O}_{\mathsf{Rand}}$ then the output distribution of A_λ is identical to the output of A_λ in H_2. Thus, the distinguishing probability of $\mathcal{B}$ is $> \nu(\lambda)$, which is a contradiction to the quantum-query security of F_1. □

We fix some lexicographic ordering on $\{0,1\}^{d(\lambda)}$ and we use this ordering implicitly in the next few hybrids.

Hybrid $H_{3.i}$ for $i \in \{0,1\}^{d(\lambda)}$. Output $A_\lambda^{\mathcal{O}_{H_{3.i}}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$, where $O_{H_{3.i}}$ is defined below.

Sample $\widehat{F_1} : \{0,1\}^{d(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ uniformly at random. Also, sample $\widehat{F_2} : \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}$ uniformly at random. On input a d-qubit register $\mathbf{X}$, do the following:

1. Initialize a register $\mathbf{K}$ with $|0\rangle$.

2. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$, where $V_{\widehat{F_1}}$ is defined as follows: $U_{\widehat{F_1}} |x\rangle |a\rangle = |x\rangle |\widehat{F_1}(x) \oplus a\rangle$.

3. Initialize two registers $\mathbf{Y}$ and $\mathbf{Anc}$ with $|0\rangle$ and $|-\rangle$ respectively.

4. Apply the unitary U^{Hyb} on the registers $\mathbf{X}, \mathbf{K}, \mathbf{Y}$ and $\mathbf{Anc}$, where U^{Hyb} maps $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |0\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ to $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ and $|\psi_{k_x}\rangle$ is defined below:
   - If $x \le i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y \in \{0,1\}^{n(\lambda)}} \frac{(-1)^{\widehat{F_2}(x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$,
   - If $x > i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y \in \{0,1\}^{n(\lambda)}} \frac{(-1)^{F_2(k_x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$

5. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$.

6. Output the state in the registers $\mathbf{X}$ and $\mathbf{Y}$.

Claim 3. *For every $i \in \{0,1\}^{d(\lambda)}$, assuming the quantum-query $\frac{1}{2^{\lambda^\delta}}$-security of F_2, A_λ can distinguish the hybrids $H_{3.i}$ and $H_{3.i-1}$ only with probability at most $\frac{1}{2^{\lambda^\delta}}$.*

Proof. Suppose A_λ can distinguish $H_{3.i}$ and $H_{3.i-1}$ with probability $> \frac{1}{2^{\lambda^\delta}}$. We show that there exists a QPT adversary $\mathcal{B}$ that violates the quantum-query security of F_2 with probability $> \frac{1}{2^{\lambda^\delta}}$. We will first see how to construct an inefficient $\mathcal{B}$ and later, we will see how to make $\mathcal{B}$ run in quantum poly-time.

$\mathcal{B}$ runs A_λ by simulating the oracle A_λ has access to. It first samples $\widehat{F_1} : \{0,1\}^{d(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ uniformly at random. For every query made by A_λ, $\mathcal{B}$ with oracle access to $\mathcal{O}$ (which is either $\mathcal{O}_{\mathsf{prf}}(k_i^*, \cdot)$, where $k_i^* \xleftarrow{\$} \{0,1\}^{\ell(\lambda)}$ or $\mathcal{O}_{\mathsf{Rand}}$; Definition 3.6), does the following: it performs 6 steps, where steps 1,2,3,5 and 6 are the same as in H_1 (in turn identical to the steps 1,2,3,5 and 6 of H_2). We describe the step 4 below.

4. Apply the following unitary on the registers $\mathbf{X}, \mathbf{K}$, $\mathbf{Y}$ and $\mathbf{Anc}$, where the unitary maps $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |0\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ to $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ and $|\psi_{k_x}\rangle$ is defined below[11]:
   - If $x < i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y\in\{0,1\}^{n(\lambda)}} \frac{(-1)^{\widehat{F_2}(x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$
   - If $x = i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}}$ is obtained by querying $\mathcal{O}$ on $\sum_{y\in\{0,1\}^{n(\lambda)}} \frac{1}{\sqrt{2^{n(\lambda)}}} |y\rangle |-\rangle$ and then trace out the last qubit.
   - If $x > i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y\in\{0,1\}^{n(\lambda)}} \frac{(-1)^{F_2(k_x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$

If $\mathcal{B}$ has oracle access to $\mathcal{O}_{\mathsf{prf}}(k_i^*, \cdot)$ then the output distribution of A_λ is identical to the output distribution of A_λ in $H_{3.i-1}$. If $\mathcal{B}$ has oracle access to $\mathcal{O}_{\mathsf{Rand}}$ then the output distribution of A_λ is identical to the output of A_λ in $H_{3.i}$. Thus, the distinguishing probability of $\mathcal{B}$ is at least $\frac{1}{2^{\lambda^\delta}}$.

Since $\mathcal{B}$ samples $\widehat{F_1}$ and $\widehat{F_2}$ uniformly at random, $\mathcal{B}$ does not run in polynomial time. However, using Lemma 3.8, we can replace both $\widehat{F_1}$ and $\widehat{F_2}$ with $2q(\lambda)$-wise independent hash functions, where $q(\lambda)$ is the number of queries made by A_λ, without changing the distinguishing probability. Once this change is made, $\mathcal{B}$ is now a QPT algorithm. This contradicts the quantum-query security of F_2. □

Similarly, the following proof also holds. Let $i_{\min}$ be the minimum element in $\{0,1\}^{d(\lambda)}$ according to the lexicographic ordering described earlier.

Claim 4. *For every $i \in \{0,1\}^{d(\lambda)}$, assuming the quantum-query $\frac{1}{2^{\lambda^\delta}}$-security of F_2, A_λ can distinguish the hybrids H_2 and $H_{3.i_{\min}}$ only with probability at most $\frac{1}{2^{\lambda^\delta}}$.*

Hybrid H_4. Output $A_\lambda^{\mathcal{O}_{H_4}(k,\cdot)}$, where $k \leftarrow \{0,1\}^\lambda$, where O_{H_4} is defined below.

Sample $\widehat{F_1} : \{0,1\}^{d(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ uniformly at random. Also, sample $\widehat{F_2} : \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{n(\lambda)} \to \{0,1\}$ uniformly at random. On input a d-qubit register $\mathbf{X}$, do the following:

1. Initialize a register $\mathbf{K}$ with $|0\rangle$.

2. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$, where $V_{\widehat{F_1}}$ is defined as follows: $U_{\widehat{F_1}} |x\rangle |a\rangle = |x\rangle |\widehat{F_1}(x) \oplus a\rangle$.

[11] To see how to implement this unitary, let us take an example where A_λ queries a pure state $\sum_{x\in\{0,1\}^{d(\lambda)}} \alpha_x |x\rangle_{\mathbf{X}}$. This argument can be naturally generalized to the case when A_λ queries a mixed state. Using an appropriately defined controlled unitary, first create the state $\sum_{x\neq i} \alpha_x |x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} + \alpha_i |i\rangle_{\mathbf{X}} |k_i\rangle_{\mathbf{K}} |0\rangle_{\mathbf{Y}}$, where $|\psi_{k_x}\rangle$ for $x \neq i$ is computed as mentioned in the bullets in the proof of Claim 3. Then, using the oracle $\mathcal{O}$, create the state $|\psi_i\rangle$. Using $|\psi_i\rangle$ and the controlled SWAP operation, create the state $\sum_{x\neq i} \alpha_x |x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |\psi_i\rangle_{\mathbf{Z}} + \alpha_i |i\rangle_{\mathbf{X}} |k_i\rangle_{\mathbf{K}} |\psi_i\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Z}}$, where $\mathbf{Z}$ is some new register. Using the oracle $\mathcal{O}$ again, we can create the state $\sum_{x\neq i} \alpha_x |x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Z}} |0\rangle_{\mathbf{Anc}} + \alpha_i |i\rangle_{\mathbf{X}} |k_i\rangle_{\mathbf{K}} |\psi_i\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Z}} |\mathcal{O}(0)\rangle_{\mathbf{Anc}}$. Finally, after querying $\mathcal{O}$ on 0, we can suitably modify the previous state to obtain $\sum_{x\neq i} \alpha_x |x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Z}} |-\rangle_{\mathbf{Anc}} + \alpha_i |i\rangle_{\mathbf{X}} |k_i\rangle_{\mathbf{K}} |\psi_i\rangle_{\mathbf{Y}} |0\rangle_{\mathbf{Z}} |-\rangle_{\mathbf{Anc}}$. We can trace out the $\mathbf{Z}$ register to obtain the desired outcome.

3. Initialize two registers $\mathbf{Y}$ and $\mathbf{Anc}$ with $|0\rangle$ and $|-\rangle$ respectively.
4. Apply the unitary $U_{\widehat{F_2}}$ on the registers $\mathbf{X}, \mathbf{K}$, $\mathbf{Y}$ and $\mathbf{Anc}$, where $U_{\widehat{F_2}}$ is defined the same way as U_{F_2} except that $\widehat{F_2}$ is used instead of F_2.
5. Apply the unitary $V_{\widehat{F_1}}$ on the registers $\mathbf{X}$ and $\mathbf{K}$.
6. Output the state in the registers $\mathbf{X}$ and $\mathbf{Y}$.

The following claim follows from the descriptions of $H_{3.i_{\max}}$, where $i_{\max}$ is the maximum element in $\{0,1\}^{d(\lambda)}$ according to the lexicographic ordering considered earlier, and H_4.

Claim 5. *The hybrids $H_{3.i_{\max}}$ and H_4 are identically distributed.*

Hybrid H_5. Output $A_\lambda^{\mathcal{O}_{\mathsf{Haar}}(\cdot)}$.

Claim 6. *Suppose q be the number of queries made by A_λ. Then, A_λ can distinguish the hybrids H_4 and H_5 with probability at most $O\left(\frac{q \cdot 2^{d(\lambda)}}{2^{\frac{n(\lambda)}{2}}}\right)$.*

Proof. We can define a sequence of $2^{d(\lambda)}$ intermediate hybrids. Hybrid $H_{4.i}$, for $i \in \{0,1\}^{d(\lambda)}$, is defined as follows: it behaves like hybrid H_4 except in Step 4.

4. Apply the following unitary on the registers $\mathbf{X}, \mathbf{K}$, $\mathbf{Y}$ and $\mathbf{Anc}$, where the unitary maps $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |0\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ to $|x\rangle_{\mathbf{X}} |k_x\rangle_{\mathbf{K}} |\psi_{k_x}\rangle_{\mathbf{Y}} |-\rangle_{\mathbf{Anc}}$ and $|\psi_{k_x}\rangle$ is defined below:
   - If $x \leq i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} = \sum_{y \in \{0,1\}^{n(\lambda)}} \frac{(-1)^{\widehat{F_2}(x,y)}}{\sqrt{2^{n(\lambda)}}} |y\rangle$
   - If $x > i$, then $|\psi_{k_x}\rangle_{\mathbf{Y}} \leftarrow \mathscr{H}_{n(\lambda)}$.

Hybrids $H_{4.i-1}$ and $H_{4.i}$ are $O(\frac{q}{2^{n(\lambda)frm-e}})$-close from Lemma 5.8. Moreover, for the same reason, Hybrids H_4 and $H_{4.1}$ are also $O(\frac{q}{2^{n(\lambda)/2}})$-close. Finally, hybrids $H_{4.i_{\max}}$, where $i_{\max}$ is the maximum element, and H_5 are identically distributed.

From this, it follows that the hybrids H_4 and H_5 are $O\left(\frac{q \cdot 2^{d(\lambda)}}{2^{n(\lambda)/2}}\right)$-statistically close. □

By applying triangle inequality to Claim 2, Claim 3, Claim 4, Claim 5 and Claim 6, the following holds:

$$\varepsilon = \nu(\lambda) + \frac{2^{d(\lambda)}}{2^{\lambda^\delta}} + O\left(\frac{q \cdot 2^{d(\lambda)}}{2^{\frac{n(\lambda)}{2}}}\right)$$

Using the facts that $d(\lambda) = \lambda^\gamma$, $\gamma < \delta$ and $\frac{n(\lambda)}{2} - d(\lambda) = \omega(\log(\lambda))$, it follows that $\varepsilon(\lambda)$ is a negligible function in λ. □

5.3.2 Second Construction

We now present our second construction of QAPRFS. The advantage of this construction and analysis is that the security of Figure 1 can be based only on the existence of polynomially-secure post-quantum one-way functions (as opposed to the sub-exponentially secure ones needed by the first construction). The security proof is a bit more involved and uses a quantum unitary version of Zhandry's small range distribution theorem, proven in Appendix A, which could be of independent interest.

The construction is actually nearly identical, except we set the pseudorandom functions F_1, F_2 to the *same* pseudorandom function f. In the previous subsection, we had to treat F_1 and F_2 separately since we required different levels of security from both.

Theorem 5.12. *Assuming the existence of quantum query secure one-way functions, Figure 1 with $F_1 = F_2 = f$ is a $(d(\lambda), n(\lambda))$-QAPRFS.*

Proof. Assume for contradiction that there is an adversary that distinguishes the real world oracle from the ideal world oracle with noticeable advantage $\alpha(\lambda)$ with $q(\lambda)$ queries. We carry out the following hybrids, with $r = 1200q^3/\alpha$.

Hybrid 0 This is the real world oracle.

Hybrid 1 We change f_k in the construction to a random function f, i.e. the oracle is now

$$|x\rangle \mapsto |x\rangle \otimes \sum_{y\in\{0,1\}^d} (-1)^{f(x,y)} |y\rangle .$$

This is computationally indistinguishable to the last hybrid by the quantum query security of PRF.

Hybrid 2 We interpret the random function as $\{0,1\}^d \to (\{0,1\}^n \to \{0,1\}^n)$ instead of $\{0,1\}^{d+n} \to \{0,1\}^n$, and change the random function f to instead be sampled by a small-range distribution $\mathsf{SR}_r^{\{0,1\}^n\to\{0,1\}^n}(\{0,1\}^d)$, i.e. we sample r random functions $\{0,1\}^n \to \{0,1\}^n$ and use a random one for every prefix $f(x,\cdot)$, where $x \in \{0,1\}^d$. By Theorem A.6, this change is statistically indistinguishable except with $300q^3/r$ advantage.

Hybrid 3.0 − 3.r Hybrid 3.0 is the same as Hybrid 2; or rather, we will design it in a way so that they are identically distributed. In particular, we are going to consider Hybrid 3.0 to be sampled equivalently (as Hybrid 2) as follows: for each $i \in [r]$, sample a function g_i from $\{0,1\}^n \to \{0,1\}^n$; and upon invoking at $x \in \{0,1\}^d$, we consider invoking the isometry U_{i_x}, where isometry U_i outputs $\sum_{y\in\{0,1\}^d} (-1)^{g_i(y)} |y\rangle$ for $i \in [r]$.

For $i = 1, ..., r$, Hybrid 3.i is the same as Hybrid 3.$(i-1)$ except that the i-th entry of the small range distribution U_i is changed from above to an isometry that outputs a Haar random state. (More formally, we are going to sample a Haar random state $|\psi_i\rangle$ and when U_i is invoked, it simply outputs $|\psi_i\rangle$.) By Lemma 5.8, this change is statistically indistinguishable except with a negligible $O(q/2^{n/2})$ advantage.

Hybrid 4 We change the unitary distribution from $\mathsf{SR}_r^{U_n}(\{0,1\}^d)$ to $U_n^{\{0,1\}^d}$, where U_n denote the isometry for outputting n-qubit Haar random states as specified before. This is equivalent to the ideal world oracle, and it is statistically indistinguishable from Hybrid 3.r except with $300q^3/r$ advantage again by Theorem A.6.

By triangle inequality, we conclude that no efficient adversary is able to distinguish Hybrid 0 from Hybrid 4 with advantage more than $\frac{600q^3}{r} + \varepsilon = \alpha/2 + \varepsilon$ for some negligible quantity ε. This is a contradiction as α is noticeable. □

QAPRFS from PRUs. We present another construction of QAPRFS from pseudorandom unitaries. We first recall the definition of PRUs [JLS18], with added parameters about the size of the unitary.

Definition 5.13. *Let $\mathcal{H}(\lambda)$ be the Hilbert space over $n(\lambda)$ qubits. A family of unitary operators $\{U_k \in U(\mathcal{H})\}_{k\in\mathcal{K}}$ is pseudorandom, if two conditions hold:*

1. *(Efficient generation) There is an efficient quantum algorithm Q, such that for all $k \in \{0,1\}^\lambda$ and any $|\psi\rangle \in S(\mathcal{H})$, $Q(k, |\psi\rangle) = U_k |\psi\rangle$.*
2. *(Pseudorandomness) U_k for a random k, given as an oracle, is computationally indistinguishable from a Haar random unitary operator. More precisely, for any efficient quantum algorithm $\mathcal{A}$, there exists a negligible function $\varepsilon(\cdot)$ such that*

$$\left| \Pr_{k\leftarrow\{0,1\}^\lambda}[A^{U_k}(1^\lambda) = 1] - \Pr_{U\leftarrow\mu}[A^U(1^\lambda) = 1] \right| \leq \varepsilon(\lambda).$$

Theorem 5.14. *Assuming $n(\lambda)$-qubit pseudorandom unitaries (PRUs) with $n = \omega(\log \lambda)$, $(d(\lambda), n(\lambda))$-QAPRFS exist for any $d \leq n$.*

Proof. We are going to consider the input length to be n instead as $d < n$ can be handled by padding zeroes.

Let Q_λ be a pseudorandom unitary. The construction of $(d(\lambda), n(\lambda))$-QAPRFS G_λ is defined as follows: $G(k, x) = Q_\lambda(k, x)$.

More formally, this is implemented as follows. We will abuse the notation and denote the circuit computing $G(k, \cdot)$ as $G_k(\cdot)$. Similarly, let U_k be the unitary associated with $Q_\lambda(k, \cdot)$. We implement G_k as follows: G_k is an isometry that maps $|x\rangle$ to $|x\rangle \otimes U_k |x\rangle$.

We argue that G_λ is a quantum accessible-secure PRFS. Assume for contradiction that there is an adversary that distinguishes the real world oracle from the ideal world oracle with noticeable advantage $\alpha(\lambda)$ with $q(\lambda)$ queries. We carry out the following hybrids, with $r = 1200q^3/\alpha$.

Hybrid 0 This is the real world oracle.

Hybrid 1 We change U_k in the construction to a Haar-random unitary U, i.e. the oracle is now

$$|x\rangle \mapsto |x\rangle \otimes U |x\rangle .$$

This is computationally indistinguishable to the last hybrid by security of PRU.

Hybrid 2 Instead of applying U, we instead first apply a random permutation $\Pi : \{0,1\}^n \to \{0,1\}^n$ followed by a Haar random unitary U'. More formally, the oracle now computes the following:

$$|x\rangle \mapsto |x\rangle \otimes U'\Pi |x\rangle = |x\rangle \otimes U' |\Pi(x)\rangle .$$

By unitary invariance of Haar, Hybrid 2 is identically distributed as Hybrid 1.

Hybrid 3 We replace Π with a random function $f : \{0,1\}^n \to \{0,1\}^n$. In other words, the oracle is now

$$|x\rangle \mapsto |x\rangle \otimes U' |f(x)\rangle .$$

Using (quantum) collision resistance of random functions [Zha15], we know that Hybrids 2 and 3 are statistically indistinguishable except with a negligible $O(q^3/2^n)$ advantage.

Hybrid 4 We change the random function f to instead be sampled by a $\mathsf{SR}_r^{\{0,1\}^n \to \{0,1\}^n}(\{0,1\}^d)$. By Theorem A.6, this change is statistically indistinguishable except with $300q^3/r$ advantage.

Hybrid 5.1 − 5.r Hybrid 5.1 is the same as Hybrid 4. However, we are going to reinterpret the oracle as a unitary oracle. Formally, we are going to consider Hybrid 5.1 to be sampled equivalently (as Hybrid 4) as follows: for each $i \in [r]$, sample a Haar random state $|\psi_i\rangle$ that is orthogonal to the span of $\{|\psi_j\rangle\}_{j=1,\dots,i-1}$; and upon invoking at $x \in \{0,1\}^d$, we consider invoking the isometry U_{i_x}, where isometry U_i outputs $|\psi_i\rangle$ for $i \in [r]$.

For $i = 2, ..., r$, Hybrid 5.i is the same as Hybrid 5.$(i-1)$ except that the $|\psi_i\rangle$ is sampled from Haar random instead of Haar random in a subspace. To compare the difference between these two hybrids, we can consider $|\psi_i\rangle$ to be first sampled from Haar random for Hybrid 5.i, and then project it to the orthogonal subspace to produce $|\psi_i'\rangle$ for Hybrid 5.$(i-1)$. Note that any projector is 2-Lipschitz, and thus by Lévy's lemma [AQY21, Fact 2.2], the probability that the overlap of a Haar random state with a subspace of polynomial dimension is more than $\frac{\log \lambda}{2^{n/2}}$ (which is negligible in λ) is negligible. Therefore, we establish that these hybrids are statistically indistinguishable except with a negligible probability by Lemma A.3.

Hybrid 6 We change the unitary distribution from $\mathsf{SR}_r^{U_n}(\{0,1\}^d)$ to $U_n^{\{0,1\}^d}$, where U_n denote the unitary for n-qubit Haar random states. This is equivalent to the ideal world oracle, and it is statistically indistinguishable except with $300q^3/r$ advantage again by Theorem A.6.

By triangle inequality, we conclude that no efficient adversary is able to distinguish Hybrid 0 from Hybrid 6 with advantage more than $\frac{600q^3}{r} + \varepsilon$ for some negligible quantity ε, which is negligibly close to $\alpha/2$ by our choice of parameters. This is a contradiction as α is noticeable. □

We conclude by noting that both constructions above (a) use PRF/PRU as a black-box (and their proofs relativize to the setting of either random oracles or families of Haar unitary oracles as considered by Kretschmer [Kre21]); and (b) only get super-logarithmic output length. We leave as future work to construct these for logarithmic output length.

6 On the Necessity of Computational Assumptions

The following lemma shows that the security guarantee of a PRS generator (and thus of PRFS generators) can only hold with respect to computationally bounded distinguishers, provided that the output length is at least $\log \lambda$.

Lemma 6.1. *Let G be a PRS generator with output length $n(\lambda) \geq \log \lambda$. Then there exists a polynomial $t(\lambda)$ and a quantum algorithm A (not efficient in general) such that*

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[A_\lambda \left(G_\lambda(k)^{\otimes t(\lambda)} \right) = 1 \right] - \Pr_{|\vartheta\rangle \leftarrow \mathscr{H}_{n(\lambda)}} \left[A_\lambda \left(|\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)} \right) = 1 \right] \right| \geq \frac{1}{3}$$

for all sufficiently large λ.

Proof. For notational convenience we abbreviate $n = n(\lambda)$ and $t = t(\lambda)$. We split the proof into two cases.

Case 1: if there does not exist a negligible function $\nu(\cdot)$ such that

$$\Pr_k \left[\min_{|\theta\rangle} \mathrm{TD}(G_\lambda(k), |\theta\rangle\langle\theta|) \le \nu(\lambda) \right] \ge \frac{1}{2}. \tag{3}$$

Then there exists some non-negligible function $\kappa(\cdot)$ such that with probability at least $\frac{1}{2}$ over the choice of k, $\min_{|\theta\rangle} \mathrm{TD}(G_\lambda(k), |\theta\rangle\langle\theta|) \ge \kappa(\lambda)$. Let $\nu_{k,1} \ge ... \ge \nu_{k,2^n}$ and $|\alpha_{k,1}\rangle, ..., |\alpha_{k,2^n}\rangle$ be eigenvalues and eigenvectors for $G_\lambda(k)$. Then $\kappa \le \mathrm{TD}(G_\lambda(k), |\alpha_{k,1}\rangle\langle\alpha_{k,1}|) = \frac{1}{2}(1 - \nu_{k,1} + \nu_{k,2} + \cdots + \nu_{k,2^n}) = 1 - \nu_{k,1}$. Thus by Hölder's inequality, $\mathrm{Tr}(G_\lambda(k)^2) \le 1 - \kappa$. Therefore, a purity test using $t = O(1/\kappa(\lambda))$ copies will correctly reject PRS states with probability at least $\frac{1}{3}$ but never incorrectly reject any Haar random state.

Case 2: if there exists a negligible function $\nu(\cdot)$ such that (3) holds. There exists a polynomial $t(\lambda)$ such that

$$2^\lambda \le \frac{1}{6} \cdot \dim \Pi_{\mathsf{sym}}^{2^n,t} = \frac{1}{6} \cdot \binom{2^n + t - 1}{t}$$

for all sufficiently large λ. This is because by setting $t = \lambda + 1$, we can lower bound the dimension of $\Pi_{\mathsf{sym}}^{2^n,t}$ by $\binom{2\lambda}{\lambda+1}$ and

$$\binom{2\lambda}{\lambda} \ge \frac{\lambda}{\lambda+1} \frac{4^\lambda}{\sqrt{\pi\lambda}} \left(1 - \frac{1}{8\lambda}\right)$$

which is much larger than $6 \cdot 2^\lambda$ for all sufficiently large λ.

Let $g \subseteq \{0,1\}^\lambda$ be the set of k's such that $\min_{|\theta\rangle} \mathrm{TD}(G_\lambda(k), |\theta\rangle\langle\theta|) \le \nu(\lambda)$. Note that 2^λ is an upper bound on the rank of the density matrix

$$\mathop{\mathbb{E}}_{k \leftarrow g} |\psi_k\rangle\langle\psi_k|^{\otimes t}, \tag{4}$$

where $|\psi_k\rangle = \arg\min_{|\theta\rangle} \mathrm{TD}(G_\lambda(k), |\theta\rangle\langle\theta|)$. Note that by Fact 5.3 the rank of the density matrix

$$\mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathscr{H}_{n(\lambda)}} |\vartheta\rangle\langle\vartheta|^{\otimes t} = \frac{\Pi_{\mathsf{sym}}^{2^n,t}}{\dim \Pi_{\mathsf{sym}}^{2^n,t}} \tag{5}$$

is equal to $\dim \Pi_{\mathsf{sym}}^{2^n,t}$.

For all λ, define the quantum circuit A_λ that, given a state on tn qubits, performs the two-outcome measurement $\{P, I - P\}$ where P is the projector onto the support of $\mathbb{E}_{k \leftarrow g} |\psi_k\rangle\langle\psi_k|^{\otimes t}$, and accepts if the P outcome occurs.

By assumption of case 2, given the density matrix (4) the circuit A_λ will accept with probability at least $\frac{1}{2}$. On the other hand, given the density matrix (5) the circuit A_λ will accept with probability

$$\mathrm{Tr}\left(P \cdot \frac{\Pi_{\mathsf{sym}}^{2^n,t}}{\dim \Pi_{\mathsf{sym}}^{2^n,t}}\right) \le \mathrm{Tr}\left(\frac{P}{\dim \Pi_{\mathsf{sym}}^{2^n,t}}\right) = \frac{\mathrm{rank}(P)}{\dim \Pi_{\mathsf{sym}}^{2^n,t}} \le \frac{1}{6}\ .$$

Letting $A = \{A_\lambda\}_\lambda$ we obtained the desired Lemma statement. □

We remark that the attack given in Lemma 6.1 cannot be used on smaller output length, up to additive factors of superpolynomially smaller order in the output length. Suppose $n = \log\lambda - \omega(\log\log\lambda)$ and for any $t = \lambda^{O(1)}$,

$$\begin{aligned}\log\binom{2^n+t-1}{t} &\leq 2^n\cdot\log\frac{e(2^n+t-1)}{2^n-1}\\ &= \frac{\lambda}{\omega(\log\lambda)}\cdot O(\log\lambda).\end{aligned}$$

This means that $\binom{2^n+t-1}{t} = 2^{\lambda/\omega(\log\lambda)} \lll 2^\lambda$ and therefore the attack above does not necessarily apply. Indeed, Brakerski and Shmueli [BS20] have shown that PRS generators with output length $n(\lambda) \leq c\log\lambda$ for some $c > 0$ can be achieved with statistical security.

We conclude the section by remarking that the result of Kretschmer [Kre21] can be easily generalized so that PRS generators with output length at least $\log\lambda + c$ (for some small constant $0 < c < 2$) imply $\mathsf{BQP} \neq \mathsf{PP}$ as well[12].

7 Tomography with Verification

Quantum state tomography (or just *tomography* for short) is a process that takes as input multiple copies of a quantum state ρ and outputs a string u that is a classical description of the state ρ; for example, u can describe an approximation of the density matrix ρ, or it could be a a more succinct description such as a *classical shadow* in the sense of [HKP20]. In this paper, we use tomography as a tool to construct protocols based on pseudorandom states with only *classical* communication.

For our applications, we require tomography procedures satisfying a useful property called verification. Suppose we execute a tomography algorithm on multiple copies of a state to obtain a classical string u. The verification algorithm, given u and the algorithm to create this state, checks if u is consistent with this state or not. Verification comes in handy when tomography is used in cryptographic settings, where we would like to make sure that the adversary has generated the classical description associated with a quantum state according to some prescribed condition (this will be implictly incorporated in the verification algorithm).

Verifiable Tomography. Let $\mathcal{C} = \{\Phi_\lambda : \lambda \in \mathbb{N}\}$ be a family of channels where each channel Φ_λ takes as input $\ell(\lambda)$ qubits for some polynomial $\ell(\cdot)$. A *verifiable tomography scheme* associated with $\mathcal{C}$ is a pair $(\mathsf{Tomography}, \mathsf{Verify})$ of QPT algorithms, which have the following input/output behavior:

- $\mathsf{Tomography}$: given as input a quantum state $\rho^{\otimes L}$ for some density matrix ρ and some number L, output a classical string u (called a *tomograph* of ρ).
- Verify: given as input a pair of classical strings $(\mathbf{x}, u)$ where $\mathbf{x}$ has length $\ell(\lambda)$, output Valid or $\mathsf{Invalid}$.

We would like $(\mathsf{Tomography}, \mathsf{Verify})$ to satisfy correctness which we describe next.

[12] For readers familiar with [Kre21], it can be verified that a sufficient condition for that proof to go through is if $2^\lambda \cdot e^{-2^n/3}$ is negligible, which is satisfied if $n \geq \log\lambda + 2$.

7.1 Correctness Notions for Verifiable Tomography

We can consider two types of correctness. The first type of correctness, referred to as *same-input correctness*, states that $\mathsf{Verify}(\mathbf{x}, u)$ outputs Valid if u is obtained by running the $\mathsf{Tomography}$ procedure on copies of the output of $\Phi_\lambda(\mathbf{x})$. The second type of correctness, referred to as *different-input correctness*, states that $\mathsf{Verify}(\mathbf{x}', u)$ outputs $\mathsf{Invalid}$ if u is obtained by applying tomography to $\Phi_\lambda(\mathbf{x})$, where $(\mathbf{x}', \mathbf{x})$ do not satisfy a predicate Π.

Same-Input Correctness. Consider the following definition.

Definition 7.1 (Same-Input Correctness). *We say that* $(\mathsf{Tomography}, \mathsf{Verify})$ *satisfies* L-same-input correctness*, for some polynomial* $L(\cdot)$*, such that for every* $\mathbf{x} \in \{0,1\}^{\ell(\lambda)}$*, if the following holds:*

$$\mathsf{Pr}\left[\mathsf{Valid} \leftarrow \mathsf{Verify}\left(\mathbf{x}, \mathsf{Tomography}\left((\Phi_\lambda(\mathbf{x}))^{\otimes L(\lambda)}\right)\right)\right] \geq 1 - \mathsf{negl}(\lambda),$$

For some applications, it suffices to consider a weaker definition. Instead of requiring the correctness guarantee to hold for every input, we instead require that it holds over some input distribution.

Definition 7.2 (Distributional Same-Input Correctness). *We say that* $(\mathsf{Tomography}, \mathsf{Verify})$ *satisfies* $(L, \mathcal{D})$-distributional same-input correctness*, for some polynomial* $L(\cdot)$ *and distribution* $\mathcal{D}$ *on* $\ell(\lambda)$*-length strings, if the following holds:*

$$\mathsf{Pr}\left[\mathsf{Valid} \leftarrow \mathsf{Verify}\left(\mathbf{x}, \mathsf{Tomography}\left((\Phi_\lambda(\mathbf{x}))^{\otimes L(\lambda)}\right)\right) \ : \ \mathbf{x} \leftarrow \mathcal{D}\right] \geq 1 - \mathsf{negl}(\lambda)$$

Different-Input Correctness. Ideally, we would require that $\mathsf{Verify}(\mathbf{x}, u)$ outputs $\mathsf{Invalid}$ if u is produced by tomographing $\Phi_\lambda(\mathbf{x}')$, and $\mathbf{x}'$ is any string such that $\mathbf{x}' \neq \mathbf{x}$. However, for applications, we only require that this be the case when the pair $(\mathbf{x}, \mathbf{x}')$ satisfy a relation defined by a predicate Π. In other words, we require $\mathsf{Verify}(\mathbf{x}, u)$ outputs $\mathsf{Invalid}$ only when u is a tomograph of $\Phi_\lambda(\mathbf{x}')$ and $\Pi(\mathbf{x}', \mathbf{x}) = 0$.

We define this formally below.

Definition 7.3 (Different-Input Correctness). *We say that* $(\mathsf{Tomography}, \mathsf{Verify})$ *satisfies* (L, Π)-different-input correctness*, for some polynomial* $L(\cdot)$ *and predicate* $\Pi : \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)} \to \{0,1\}$*, such that for every* $\mathbf{x}, \mathbf{x}' \in \{0,1\}^{\ell(\lambda)}$ *satisfying* $\Pi(\mathbf{x}, \mathbf{x}') = 0$*, if the following holds:*

$$\mathsf{Pr}\left[\mathsf{Invalid} \leftarrow \mathsf{Verify}\left(\mathbf{x}', \mathsf{Tomography}\left((\Phi_\lambda(\mathbf{x}))^{\otimes L(\lambda)}\right)\right)\right] \geq 1 - \mathsf{negl}(\lambda)$$

Analogous to Definition 7.2, we correspondingly define below the notion of $(L, \mathcal{D}, \Pi)$-different-input correctness.

Definition 7.4 (Distributional Different-Input Correctness). *We say that* $(\mathsf{Tomography}, \mathsf{Verify})$ *satisfies* $(L, \Pi, \mathcal{D})$-distributional different-input correctness*, for some polynomial* $L(\cdot)$*, predicate* $\Pi : \{0,1\}^\lambda \times \{0,1\}^\lambda \to \{0,1\}$ *and distribution* $\mathcal{D}$ *supported on* $(\mathbf{x}, \mathbf{x}') \in \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)}$ *satisfying* $\Pi(\mathbf{x}, \mathbf{x}') = 0$*, if the following holds:*

$$\mathsf{Pr}\left[\mathsf{Invalid} \leftarrow \mathsf{Verify}\left(\mathbf{x}', \mathsf{Tomography}\left((\Phi_\lambda(\mathbf{x}))^{\otimes L(\lambda)}\right)\right) \ : \ (\mathbf{x}, \mathbf{x}') \leftarrow \mathcal{D}\right] \geq 1 - \mathsf{negl}(\lambda)$$

Sometimes we will use the more general $(\varepsilon, L, \Pi, \mathcal{D})$*-distributional different-input correctness* definition. In this case, the probability of Verify outputting $\mathsf{Invalid}$ is bounded below by $1 - \varepsilon$ instead of $1 - \mathsf{negl}(\lambda)$.

7.2 Verifiable Tomography Procedures

We will consider two different instantiations of $(\mathsf{Tomography}, \mathsf{Verify})$ where the first instantiation will be useful for bit commitments and the second instantiation will be useful for pseudo one-time pad schemes.

In both the instantiations, we use an existing tomography procedure stated in the lemma below.

Lemma 7.5 (Section 1.5.3, [Low21]). *There exists a tomography procedure $\mathcal{T}$ that given sN^2 copies of an N-dimensional density matrix ρ, outputs a matrix M such that $\mathbb{E}\|M-\rho\|_F^2 \le \frac{N}{s}$ where the expectation is over the randomness of the tomography procedure. Moreover, the running time of $\mathcal{T}$ is polynomial in s and N.*

We state and prove a useful corollary of the above lemma.

Corollary 7.6. *There exists a tomography procedure $\mathcal{T}_{\mathsf{imp}}$ that given $4sN^2\lambda$ copies of an N-dimensional density matrix ρ, outputs a matrix M such that the following holds:*

$$\mathsf{Pr}\left[\|M-\rho\|_F^2 \le \frac{9N}{s}\right] \ge 1 - \mathsf{negl}(\lambda)$$

Moreover, the running time of $\mathcal{T}_{\mathsf{imp}}$ is polynomial in s, N and λ.

Proof. Set $\varepsilon = \frac{N}{s}$. We define $\mathcal{T}_{\mathsf{imp}}$ as follows:

$\mathcal{T}_{\mathsf{imp}}\left(\rho^{\otimes 4sN^2\lambda}\right)$: on input $4sN^2\lambda$ copies of an N-dimensional density matrix ρ, do the following:

- For every $i \in [\lambda]$, compute $M_i \leftarrow \mathcal{T}(\rho^{\otimes 4sN^2})$,
- Output M_{i^*}, where $|\{j : \|M_j - M_{i^*}\|_F^2 \le 4\varepsilon\}| > \frac{\lambda}{2}$. If no such $i^* \in [\lambda]$ exists, output $\perp$.

To prove that $\mathcal{T}_{\mathsf{imp}}$ satisfies the condition mentioned in the statement of the corollary, we first consider the following event.

$\mathsf{GoodEvent}$: For every $i \in [\lambda]$, compute $M_i \leftarrow \mathcal{T}(\rho^{\otimes 4sN^2})$. There exists a set $S \subseteq \{M_1, \ldots, M_\lambda\}$ such that for every $M \in S$, $\|M-\rho\|_F^2 \le \varepsilon$ and moreover, $|S| > \frac{\lambda}{2}$.

Consider the following two claims.

Claim 7. $\mathsf{Pr}\left[\mathsf{GoodEvent}\right] \ge 1 - \mathsf{negl}(\lambda)$.

Proof. Applying Markov's inequality to Lemma 7.5, we have $\mathsf{Pr}\left[M \leftarrow \mathcal{T}\left(\rho^{\otimes 4sN^2}\right) \text{ and } \|M-\rho\|_F^2 \ge \varepsilon\right] \le \frac{1}{4}$. Let $\mathbf{X}_i$ be a random variable that is set to 1 if the i^{th} execution of $\mathcal{T}(\rho^{\otimes 4sN^2})$ outputs M_i such that $\|M_i - \rho\|_F^2 \ge \varepsilon$. Observe that $\mathbb{E}[\mathbf{X}_i] \le \frac{1}{4}$. Let $\mu = \mathbb{E}[\sum_{i=1}^{\lambda} \mathbf{X}_i] \le \frac{\lambda}{4}$. By Chernoff bound[13], there exists $\delta \ge 1$ such that $\mathsf{Pr}[\sum_{i=1}^{\lambda} \mathbf{X}_i \ge \frac{\lambda}{2}] = \mathsf{Pr}[\sum_{i=1}^{\lambda} \mathbf{X}_i \ge (1+\delta)\mu] \le \frac{1}{e^{\frac{\lambda}{4}\cdot\frac{1}{2+1}}} = \mathsf{negl}(\lambda)$.

This proves that with overwhelming probability, there exists a subset S of $\{M_1, \ldots, M_\lambda\}$ of size $> \frac{\lambda}{2}$ such that for every $M \in S$, $\|M-\rho\|_F^2 \le \varepsilon$. □

[13] For any set of iid Bernoulli random variables $\mathbf{X}_1, \ldots, \mathbf{X}_N$, for any $\delta > 0$, the following holds: $\mathsf{Pr}\left[\sum_{i=1}^{\lambda} \mathbf{X}_i \ge (1+\delta)\mu\right] \le \frac{1}{e^{\delta\mu\cdot\frac{\delta}{2+\delta}}}$

Claim 8. $\mathsf{Pr}\left[\|M-\rho\|_F^2 \leq 9\varepsilon \;\middle|\; \mathsf{GoodEvent}\right] = 1.$

Proof. Since we are conditioning on $\mathsf{GoodEvent}$, there exists a subset S of $\{M_1,\ldots,M_\lambda\}$ such that for every $M \in S$, $\|M-\rho\|_F^2 \leq \varepsilon$. Moreover, $|S| > \frac{\lambda}{2}$. Suppose $\mathcal{T}_{\mathsf{imp}}$ outputs $M \in S$ then we are done. This is because, for every $M, M' \in S$, it holds that $\|M-M'\|_F^2 \leq 4\varepsilon$.

So we might as well assume that $M \notin S$. By the description of $\mathcal{T}_{\mathsf{imp}}$, it follows that there exists a subset S' of $\{M_1,\ldots,M_\lambda\}$ such that $|S'| > \frac{\lambda}{2}$ and $\|M-M'\|_F^2 \leq 4\varepsilon$ for every $M' \in S'$. Since $S \cap S' \neq \emptyset$, it follows that there exists an $M' \in S$ such that $\|M-M'\|_F^2 \leq 4\varepsilon$. By definition of S and by the fact that Frobenius norm is a matrix norm, it follows that $\|M-\rho\|_F^2 \leq 9\varepsilon$. □

Thus, we have the following:

$$\begin{aligned}
\mathsf{Pr}\left[\|M-\rho\|_F^2 \leq 9\varepsilon\right] &= \mathsf{Pr}\left[\|M-\rho\|_F^2 \leq 9\varepsilon \mid \mathsf{GoodEvent}\right]\mathsf{Pr}[\mathsf{GoodEvent}] \\
&\quad + \mathsf{Pr}\left[\|M-\rho\|_F^2 \leq 9\varepsilon \mid \neg\mathsf{GoodEvent}\right]\mathsf{Pr}[\neg\mathsf{GoodEvent}] \\
&\geq \mathsf{Pr}\left[\|M-\rho\|_F^2 \leq 9\varepsilon \mid \mathsf{GoodEvent}\right]\cdot(1-\mathsf{negl}(\lambda)) && \text{(Claim 7)} \\
&= (1-\mathsf{negl}(\lambda)) && \text{(Claim 8)}
\end{aligned}$$

□

7.2.1 First Instantiation

We will work with a verifiable tomography procedure that will be closely associated with a PRFS. In particular, we will use a $(d(\lambda), n(\lambda))$-PRFS $\{G_\lambda(\cdot,\cdot)\}$ satisfying recognizable abort property (Definition 3.12). Let $\widehat{G}$ be the QPT algorithm associated with G according to Definition 3.12. Note that the output length of $\widehat{G}$ is $n+1$. We set $d(\lambda) = \lceil \frac{\log(\lambda)}{\log(\log(\lambda))} \rceil$ and $n(\lambda) = \lceil 3\log(\lambda) \rceil$.

We will describe the algorithms $(\mathsf{Tomography}, \mathsf{Verify})$ in Figure 2. The set of channels $\mathcal{C} = \{\Phi_\lambda : \lambda \in \mathbb{N}\}$ is associated with $(\mathsf{Tomography}, \mathsf{Verify})$, where Φ_λ is defined as follows:

- Let the input be initialized on register $\mathbf{A}$.
- Controlled on the first register containing the value (P_x, k, x, b), where P_x is an n-qubit Pauli, $k \in \{0,1\}^\lambda, b \in \{0,1\}$, do the following: compute $\left(I \otimes P_x^b\right)\widehat{G}_\lambda(k,x)\left(I \otimes P_x^b\right)$ and store it in the register $\mathbf{B}$.
- Trace out $\mathbf{A}$ and output $\mathbf{B}$.

The channel Φ_λ can be represented as a quantum circuit of size polynomial in λ as the PRFS generator $\hat{G}$ runs in time polynomial in λ.

Distributional Same-Input Correctness. We prove below that $(\mathsf{Tomography}, \mathsf{Verify})$ satisfies distributional same-input correctness. For every $x \in \{0,1\}^{d(\lambda)}$, for every n-qubit Pauli P_x and $b \in \{0,1\}$, define the distribution $\mathcal{D}_{P_x,x,b}$ as follows: sample $k \xleftarrow{\$} \{0,1\}^\lambda$ and output $\mathbf{x} = (P_x, k, x, b)$.

Lemma 7.7. *Let $L = O(2^{3n}\lambda)$. The verifiable tomography scheme* $(\mathsf{Tomography}, \mathsf{Verify})$ *described in Figure 2 satisfies $(L, \mathcal{D}_{P_x,x,b})$-distributional same-input correctness for all P_x, x, b.*

$\underline{\mathsf{Tomography}(\rho^{\otimes L})}$: On input L copies of an $2^{(n+1)}$-dimensional density matrix ρ, compute $\mathcal{T}_{\mathsf{imp}}(\rho^{\otimes L})$ to obtain M, where $\mathcal{T}_{\mathsf{imp}}$ is given in Corollary 7.6. Output M.

$\underline{\mathsf{Verify}(\mathbf{x}, M)}$:

1. Run $\rho^{\otimes L} \leftarrow (\Phi_\lambda(\mathbf{x}))^{\otimes L}$, where $L = 3^8 2^{3(n+1)+2}\lambda$.
2. Compute $\widehat{M} \leftarrow \mathsf{Tomography}\left(\rho^{\otimes L}\right)$.
3. If $\langle \bot | M | \bot \rangle > \frac{1}{9}$ for any $x \in \{0,1\}^d$, output Invalid.
4. If $\|M - \widehat{M}\|_F^2 \le \frac{4}{729}$ output Valid. Output Invalid otherwise.

Figure 2: First instantiation of Tomography

Proof. Define $\mathcal{K}_{\mathsf{good}} \subseteq \{0,1\}^\lambda$ such that for every $k \in \mathcal{K}_{\mathsf{good}}$ and $x \in \{0,1\}^d$, $\widehat{G}_\lambda(k,x)$ can be written as $\eta_{k,x}\left(|0\rangle\langle 0| \otimes |\psi_{k,x}\rangle\langle\psi_{k,x}|\right) + (1-\eta_{k,x})\,|\bot\rangle\langle\bot|$, where $\eta_{k,x} \ge 1 - \mathsf{negl}(\lambda)$ for all $x \in \{0,1\}^d$. From the fact that $\{G_\lambda(\cdot,\cdot)\}$ is a PRFS, it follows that $|\mathcal{K}_{\mathsf{good}}| \ge (1-\mathsf{negl}(\lambda))2^\lambda$.

Fix $k \in \mathcal{K}_{\mathsf{good}}$. Let $x \in \{0,1\}^{d(\lambda)}$, P_x be an n-qubit Pauli and $b \in \{0,1\}$. Set $\mathbf{x} = (P_x, k, x, b)$. Let $M \leftarrow \mathsf{Tomography}\left(\rho^{\otimes L}\right)$, where $\rho = \Phi_\lambda(\mathbf{x})$. We now argue that the probability that $\mathsf{Verify}(\mathbf{x}, M)$ outputs Valid is negligibly close to 1.

Let $\widehat{M} \leftarrow \mathsf{Tomography}\left(\rho^{\otimes L}\right)$, where $\rho = \Phi_\lambda(\mathbf{x})$, be generated during the execution of $\mathsf{Verify}(\mathbf{x}, M)$. Conditioned on the event that both $\|M - \rho\|_F^2 \le \frac{1}{729}$ and $\|\widehat{M} - \rho\|_F^2 \le \frac{1}{729}$ holds, we argue that $\langle \bot | M | \bot \rangle \le \frac{1}{2}$. Consider the following cases.

- Case $b = 0$: In this case, $\rho = \left(\eta_{k,x}\left(|0\rangle\langle 0| \otimes |\psi_{k,x}\rangle\langle\psi_{k,x}|\right) + (1-\eta_{k,x})\,|\bot\rangle\langle\bot|\right)$. Since $(\langle 0| \langle\psi_{k,x}|)\,|\bot\rangle = 0$ (from Definition 3.12), it follows that $\langle \bot | \rho | \bot \rangle = (1-\eta_{k,x}) \le \mathsf{negl}(\lambda)$. Since, we have that $\|M - \rho\|_F^2 \le \frac{1}{729}$, by Fact 3.3, we get $\langle \bot | M | \bot \rangle \le \mathsf{negl}(\lambda) + \frac{1}{27} + \sqrt{\frac{2-2\mathsf{negl}(\lambda)}{729}} \le \frac{1}{9}$.
- Case $b = 1$: In this case, $\rho_x = (I \otimes P_x)\left(\eta_{k,x}\left(|0\rangle\langle 0| \otimes |\psi_{k,x}\rangle\langle\psi_{k,x}|\right) + (1-\eta_{k,x})\,|\bot\rangle\langle\bot|\right)(I \otimes P_x)$. Since $(\langle 0| \langle\psi_{k,x}|)\,|\bot\rangle = 0$ (from Definition 3.12), it follows that $\langle \bot | \rho_x | \bot \rangle = (1-\eta_{k,x})\,\langle \bot | (I \otimes P_x) | \bot \rangle \langle \bot | (I \otimes P_x) | \bot \rangle$. Since, for any unitary A and any state $|\phi\rangle$, we have $\langle \phi | A | \phi \rangle \le \langle \phi | \phi \rangle$, we get $\langle \bot | \rho_x | \bot \rangle \le (1-\eta_{k,x}) \le \mathsf{negl}(\lambda)$. Similar to the above case, we get $\langle \bot | M | \bot \rangle \le \mathsf{negl}(\lambda) + \frac{1}{27} + \sqrt{\frac{2-2\mathsf{negl}(\lambda)}{729}} \le \frac{1}{9}$.

From Corollary 7.6, it follows that (a) $\Pr[\|M - \rho\|_F^2 \le \frac{92^{(n+1)}}{3^8 2^{(n+1)}} \le \frac{1}{729}] \ge 1 - \mathsf{negl}(\lambda)$ and similarly, (b) $\Pr[\|\widehat{M} - \rho\|_F^2 \le \frac{1}{729}] \ge 1 - \mathsf{negl}(\lambda)$, where the probability is over the randomness of Tomography. Thus, it follows that $\Pr[\|M - \widehat{M}\|_F^2 \le \frac{4}{729}] \ge 1 - \mathsf{negl}(\lambda)$.

□

Distributional Different-Input Correctness. We prove below that $(\mathsf{Tomography}, \mathsf{Verify})$ satisfies $(\varepsilon, L, \Pi, \mathcal{D}_x)$-different-input correctness, where Π and $\mathcal{D}_x$ are defined as follows:

$$\Pi\left((P_0, k_0, x_0, b_0), (P_1, k_1, x_1, b_1)\right) = \begin{cases} 0 & P_0 = P_1, x_0 = x_1 \text{ and } b_0 \neq b_1, \\ 1 & \text{otherwise.} \end{cases}$$

The sampler for $\mathcal{D}_x$ is defined as follows: sample $P_x \xleftarrow{\$} \mathcal{P}_n$, $k_0, k_1 \xleftarrow{\$} \{0,1\}^\lambda$ and output $((P_x, k_0, x, 0),$ $((P_x, k_1, x, 1))$. We first prove an intermediate lemma that will be useful for proving distributional different-input correctness. Later on, this lemma will also be useful in the application of bit commitments.

Lemma 7.8. *Let $P_x \in \mathcal{P}_n$ and there exists a density matrix M such that* $\mathsf{Verify}(P_x \| k_0 \| x \| 0, M) = \mathsf{Valid}$ *and* $\mathsf{Verify}(P_x \| k_1 \| x \| 1, M) = \mathsf{Valid}$, *for some $k_0, k_1 \in \{0,1\}^\lambda$. Then*

$$\mathrm{Tr}\left(P_x \left|\psi_{k_1,x}\right\rangle\!\left\langle\psi_{k_1,x}\right| P_x \left|\psi_{k_0,x}\right\rangle\!\left\langle\psi_{k_0,x}\right|\right) \geq \frac{542}{729}.$$

Proof. Since $\mathsf{Verify}(P_x \| k_0 \| x \| 0, M) = \mathsf{Valid}$, therfore

$$\langle\perp| M |\perp\rangle \leq \frac{1}{9}$$

and

$$\|M - M_0\|_F^2 \leq \frac{4}{729},$$

where $M_0 = \mathsf{Tomography}\left(\left(\Phi_\lambda\left(P_x \| k_0 \| x \| 0\right)\right)^{\otimes L}\right)$.

Similarly, since $\mathsf{Verify}(P_x \| k_1 \| x \| 1, M) = \mathsf{Valid}$,

$$\langle\perp| (I \otimes P_x) M (I \otimes P_x) |\perp\rangle \leq \frac{1}{9}$$

and

$$\|M - M_1\|_F^2 \leq \frac{4}{729},$$

where $M_1 = \mathsf{Tomography}\left(\left(\Phi_\lambda\left(P_x \| k_1 \| x \| 1\right)\right)^{\otimes L}\right)$.

Since, $M_0 = \mathsf{Tomography}\left(\left(\Phi_\lambda\left(P_x \| k_0 \| x \| 0\right)\right)^{\otimes L}\right)$ and $M_1 = \mathsf{Tomography}\left(\left(\Phi_\lambda\left(P_x \| k_1 \| x \| 1\right)\right)^{\otimes L}\right)$,

$$\|M_0 - \left(\eta_0 \left(|0\rangle\langle 0| \otimes \left|\psi_{k_0,x}\right\rangle\!\left\langle\psi_{k_0,x}\right|\right) + (1-\eta_0)\,|\perp\rangle\langle\perp|\right)\|_F^2 \leq \frac{1}{729},$$

and

$$\|M_1 - (I \otimes P_x)\left(\eta_1 \left(|0\rangle\langle 0| \otimes \left|\psi_{k_1,x}\right\rangle\!\left\langle\psi_{k_1,x}\right|\right) + (1-\eta_1)\,|\perp\rangle\langle\perp|\right)(I \otimes P_x)\|_F^2 \leq \frac{1}{729}.$$

By triangle inequality,

$$\|M - \left(\eta_0 \left(|0\rangle\langle 0| \otimes \left|\psi_{k_0,x}\right\rangle\!\left\langle\psi_{k_0,x}\right|\right) + (1-\eta_0)\,|\perp\rangle\langle\perp|\right)\|_F^2 \leq \frac{1}{81},$$

and

$$\|M - (I \otimes P_x)\left(\eta_1 \left(|0\rangle\langle 0| \otimes \left|\psi_{k_1,x}\right\rangle\!\left\langle\psi_{k_1,x}\right|\right) + (1-\eta_1)\,|\perp\rangle\langle\perp|\right)(I \otimes P_x)\|_F^2 \leq \frac{1}{81}.$$

By Fact 3.3,

$$\langle\perp| \left(\eta_0 |0\rangle\langle 0| \otimes |\psi_{k_0}\rangle\langle\psi_{k_0}| + (1-\eta_0) |\perp\rangle\langle\perp|\right) |\perp\rangle \leq 10/27,$$

and

$$\langle\perp| \left(I \otimes P_x\right) \left(\eta_1 |0\rangle\langle 0| \otimes |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| + (1-\eta_1) |\perp\rangle\langle\perp|\right) \left(I \otimes P_x\right) |\perp\rangle \leq 10/27.$$

Simplifying, we get

$$\eta_0 \geq \frac{17}{27},$$

and

$$\eta_1 \geq \frac{17}{27}.$$

Also, by triangle inequality,

$$\begin{aligned}\|(I \otimes P_x) \left(\eta_1 |0\rangle\langle 0| \otimes |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| + (1-\eta_1) |\perp\rangle\langle\perp|\right) (I \otimes P_x) \\ - \left(\eta_0 |0\rangle\langle 0| \otimes |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}| + (1-\eta_0) |\perp\rangle\langle\perp|\right) \|_F^2 \leq 4/81.\end{aligned}$$

Or,

$$\begin{aligned}\|\eta_1 (I \otimes P_x) |0\rangle\langle 0| \otimes |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| (I \otimes P_x) - \eta_0 |0\rangle\langle 0| \otimes |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}| \\ - \left((1-\eta_0) |\perp\rangle\langle\perp| - (1-\eta_1)(I \otimes P_x)x |\perp\rangle\langle\perp| (I \otimes P_x)\right) \|_F^2 \leq 4/81.\end{aligned}$$

By Fact 3.2, since $\langle\psi_0| (|0\rangle |\psi_k\rangle) = 0$,

$$\begin{aligned}\|\eta_1 (I \otimes P_x) |0\rangle\langle 0| \otimes |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| (I \otimes P_x) - \eta_0 |0\rangle\langle 0| \otimes |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}| \|_F^2 \\ + \| \left((1-\eta_0) |\perp\rangle\langle\perp| - (1-\eta_1)(I \otimes P_x) |\perp\rangle\langle\perp| (I \otimes P_x)\right) \|_F^2 \leq 4/81.\end{aligned}$$

Or,

$$\|\eta_1 (I \otimes P_x) |0\rangle\langle 0| \otimes |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| (I \otimes P_x) - \eta_0 |0\rangle\langle 0| \otimes |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}| \|_F^2 \leq 4/81.$$

By Fact 3.2,

$$\eta_1^2 + \eta_0^2 - \eta_1\eta_0 \mathrm{Tr}\left(P_x |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| P_x |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}|\right) \leq 4/81.$$

Hence,

$$\mathrm{Tr}\left(P_x |\psi_{k_1,x}\rangle\langle\psi_{k_1,x}| P_x |\psi_{k_0,x}\rangle\langle\psi_{k_0,x}|\right) \geq \frac{542}{729}.$$

□

With the above lemma in mind, we can prove the different-input correctness.

Lemma 7.9. $(\mathsf{Tomography}, \mathsf{Verify})$ *described in Figure 2 satisfies* $(O(2^{-n}), L, \Pi, \mathcal{D}_x)$*-different-input correctness, where* $L = O(2^{3n}\lambda)$*.*

Proof. Define $\mathcal{D}'_x$ as follows: sample $P_x \xleftarrow{\$} \mathcal{P}_n$, $k_0, k_1 \xleftarrow{\$} \mathcal{K}_{\mathsf{good}}$ such that $k_0 \neq k_1$ and output $((P_x, k_0, x, 0), ((P_x, k_1, x, 1))$. Let

$$\begin{aligned}p = \Pr[\mathsf{Valid} \leftarrow \mathsf{Verify}\left(P_x||k_0||x||0, \mathsf{Tomography}\left(\Phi_\lambda\left(P_x||k_1||x||1\right)\right)^{\otimes L(\lambda)}\right) : \\ \left(\left(P_x, k_0, x, 0\right), \left(P_x, k_1, x, 1\right)\right) \leftarrow \mathcal{D}_x],\end{aligned}$$

and

$$p' = \mathsf{Pr}[\mathsf{Valid} \leftarrow \mathsf{Verify}\left(P_x||k_0||x||0, \mathsf{Tomography}\left(\Phi_\lambda\left(P_x||k_1||x||1\right)\right)^{\otimes L(\lambda)}\right) : \left(\left(P_x, k_0, x, 0\right), \left(P_x, k_1, x, 1\right)\right) \leftarrow \mathcal{D}'_x]$$

Then, since $|\mathcal{K}_{\mathsf{good}}| \geq (1 - \mathsf{negl}(\lambda))2^\lambda$, we know that

$$p = p' \cdot \mathsf{Pr}[k_0, k_1 \in \mathcal{K}_{\mathsf{good}} : \left(\left(P_x, k_0, x, 0\right), \left(P_x, k_1, x, 1\right)\right) \leftarrow \mathcal{D}_x] + \mathsf{negl}(\lambda),$$

or

$$p = p'(1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda).$$

By Lemma 7.7, we know that

$$\mathsf{Pr}[\mathsf{Valid} \leftarrow \mathsf{Verify}\left(P_x||k_1||x||1, \mathsf{Tomography}\left(\Phi_\lambda\left(P_x||k_1||x||1\right)\right)^{\otimes L(\lambda)}\right) : k_1 \xleftarrow{\$} \mathcal{K}_{\mathsf{good}}] \geq 1 - \mathsf{negl}(\lambda).$$

Hence,

$$p' = \mathsf{Pr}\left[\begin{smallmatrix}\mathsf{Verify}(P_x||k_0||x||0, M_x) = \mathsf{Valid}, \\ \mathsf{Verify}(P_x||k_1||x||1, M_x) = \mathsf{Valid}\end{smallmatrix} : \begin{smallmatrix}((P_x, k_0, x, 0), (P_x, k_1, x, 1)) \leftarrow \mathcal{D}'_x \\ M_x = \mathsf{Tomography}(\Phi_\lambda(P_x||k_1||x||1))^{\otimes L(\lambda)}\end{smallmatrix}\right] \cdot (1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda).$$

By Lemma 7.8,

$$p' \leq \mathsf{Pr}\left[\mathrm{Tr}\left(P_x \left|\psi_{k_1,x}\right\rangle\left\langle\psi_{k_1,x}\right| P_x \left|\psi_{k_0,x}\right\rangle\left\langle\psi_{k_0,x}\right|\right) \geq \frac{542}{729} : \left(\left(P_x, k_0, x, 0\right), \left(P_x, k_1, x, 1\right)\right) \leftarrow \mathcal{D}'_x\right] \cdot (1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda),$$

Or

$$p' \leq \mathsf{Pr}\left[\left|\left\langle\psi_{k_1,x}\right| P_x \left|\psi_{k_0,x}\right\rangle\right|^2 \geq \frac{542}{729} : \left(\left(P_x, k_0, x, 0\right), \left(P_x, k_1, x, 1\right)\right) \leftarrow \mathcal{D}'_x\right] \cdot (1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda).$$

We use the following fact [AQY21, Fact 6.9]: Let $|\psi\rangle$ and $|\phi\rangle$ be two arbitrary n-qubit states. Then,

$$\mathop{\mathbb{E}}_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[\left|\langle\psi| P_x |\phi\rangle\right|^2\right] = 2^{-n}.$$

For any k_0, k_1, x by the above fact, $\mathbb{E}_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[\left|\left\langle\psi_{k_0,x}\right| P_x \left|\psi_{k_1,x}\right\rangle\right|^2\right] = 2^{-n}$. Using Markov's inequality we get that for all $\delta > 0$,

$$\Pr_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[\left|\left\langle\psi_{k_0,x}\right| P_x \left|\psi_{k_1,x}\right\rangle\right|^2 \geq \delta\right] \leq \delta^{-1}2^{-n}.$$

Hence,

$$p' \leq \frac{729}{542}2^{-n} \cdot (1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda),$$

and

$$p \leq \frac{729}{542}2^{-n} \cdot (1 - \mathsf{negl}(\lambda)) + \mathsf{negl}(\lambda).$$

Hence, the scheme satisfies satisfies $(O(2^{-n}), L, \Pi, \mathcal{D}_x)$-different-input correctness.

□

7.2.2 Second Instantiation

Similar to the first instantiation, we will start with a $(d(\lambda)+1, n(\lambda))$-PRFS $\{G_\lambda(\cdot,\cdot)\}$ satisfying recognizable abort property (Definition 3.12). We set $d(\lambda) = \lceil \log(\lambda) \rceil$ and $n(\lambda) = \lceil \log(\lambda) \rceil$.

We will describe the algorithms $(\mathsf{Tomography}, \mathsf{Verify})$ in Figure 3. The set of channels $\mathcal{C} = \{\Phi_\lambda : \lambda \in \mathbb{N}\}$ is associated with $(\mathsf{Tomography}, \mathsf{Verify})$, where Φ_λ is defined as follows:

- Let the input be initialized on register $\mathbf{A}$.
- Controlled on the first register containing the value $(k, i, b) \in \{0,1\}^{\ell(\lambda)}$, where $k \in \{0,1\}^\lambda, i \in \{0,1\}^d, b \in \{0,1\}$, do the following: compute $G_\lambda(k, i \| b)$ and store the result in the register $\mathbf{B}$.
- Trace out $\mathbf{A}$ and output $\mathbf{B}$.

$\underline{\mathsf{Tomography}(\rho^{\otimes L})}$: On input L copies of an 2^n-dimensional density matrix ρ, compute $\mathcal{T}_{\mathsf{imp}}(\rho^{\otimes L})$ to obtain M, where $\mathcal{T}_{\mathsf{imp}}$ is given in Corollary 7.6. Output M.

$\underline{\mathsf{Verify}(\mathbf{x}, M)}$:

1. Run $\rho^{\otimes L} \leftarrow (\Phi_\lambda(\mathbf{x}))^{\otimes L}$, where $L = 2^{3n+11}\lambda$.
2. Run $\mathsf{Tomography}\left(\rho^{\otimes L}\right)$ to get $\widehat{M}$.
3. If $\|M - \widehat{M}\|_F^2 \leq \frac{9}{128}$, output Valid. Output $\mathsf{Invalid}$ otherwise.

Figure 3: Second instantiation of $\mathsf{Tomography}$

Same-Input Correctness. We prove below that $(\mathsf{Tomography}, \mathsf{Verify})$ satisfies same-input correctness.

Lemma 7.10. $(\mathsf{Tomography}, \mathsf{Verify})$ *described in Figure 3 satisfies L-same-input correctness.*

Proof. Suppose $M \leftarrow \mathsf{Tomography}\left(\rho^{\otimes L}\right)$, where $\rho = \Phi_\lambda(\mathbf{x})$ for some $\mathbf{x} \in \{0,1\}^{\ell(\lambda)}$. We prove that $\mathsf{Verify}(\mathbf{x}, M)$ outputs Valid with overwhelming probability. Let $\widehat{M} \leftarrow \mathsf{Tomography}\left(\rho^{\otimes L}\right)$ be generated during the execution of $\mathsf{Verify}(\mathbf{x}, M)$.

Let us condition on the event that $\|M - \rho\|_F^2 \leq \frac{9}{512}$ and $\|\widehat{M} - \rho\|_F^2 \leq \frac{9}{512}$. Using triangle inequality, we have that $\|M - \widehat{M}\|_F^2 \leq \frac{9}{128}$.

All that is left is to prove that $\|M - \rho\|_F^2 \leq \frac{9}{512}$ and $\|\widehat{M} - \rho\|_F^2 \leq \frac{9}{512}$ holds with overwhelming probability. We can invoke Corollary 7.6 since the dimension of ρ is 2^n and the number of copies of ρ used in tomography is $4 \cdot 2^{n+9} \cdot (2^n)^2 \cdot \lambda$. In more detail, we have the following:

$$\begin{aligned}
\mathsf{Pr}\left[\|M - \rho\|_F^2 \leq \frac{9}{512} \text{ and } \|\widehat{M} - \rho\|_F^2 \leq \frac{9}{512}\right] &= \mathsf{Pr}\left[\|M - \rho\|_F^2 \leq \frac{9}{512}\right] \mathsf{Pr}\left[\|\widehat{M} - \rho\|_F^2 \leq \frac{9}{512}\right] \\
&\geq (1 - \mathsf{negl}(\lambda))(1 - \mathsf{negl}(\lambda)) \quad \text{(Corollary 7.6)} \\
&\geq (1 - \mathsf{negl}(\lambda))
\end{aligned}$$

□

Distributional Different-Input Correctness. We prove below that $(\mathsf{Tomography}, \mathsf{Verify})$ satisfies different-input correctness.

Lemma 7.11. *Assuming the security of $\{G_\lambda(\cdot,\cdot)\}$, $(\mathsf{Tomography}, \mathsf{Verify})$ described in Figure 3 satisfies $(L, \Pi, \mathcal{D}_i)$-different-input correctness, for every $i \in \{0,1\}^d$, where $\Pi : \{0,1\}^{\ell(\lambda)} \times \{0,1\}^{\ell(\lambda)} \to \{0,1\}$ and a distribution $\mathcal{D}_i$ are defined as follows:*

- *$\Pi((k,i,b),(k',i',b')) = 0$ if and only if $k = k'$, $i = i'$ and $b \neq b'$,*
- *$\mathcal{D}_i$ is a distribution that samples $k \xleftarrow{\$} \{0,1\}^\lambda$ and outputs $((k,i,0),(k,i,1))$.*

Proof. Define $\mathcal{K}_{\mathsf{good}} \subseteq \{0,1\}^\lambda$ as follows: for every $k \in \{0,1\}^\lambda$, $k \in \mathcal{K}_{\mathsf{good}}$ if and only if for every $i \in \{0,1\}^d$, $\|G_\lambda(k,i\|0) - G_\lambda(k,i\|1)\|_F^2 > \frac{81}{512}$.

We show the following.

Claim 9. *For every $k \in \mathcal{K}_{\mathsf{good}}$, for every bit b,*

$$\mathsf{Pr}\left[\mathsf{Invalid} \leftarrow \mathsf{Verify}((k,i,b),M) \;\;:\;\; M \leftarrow \mathsf{Tomography}((\Phi_\lambda((k,i,1-b))^{\otimes L}))\right] \geq 1 - \mathsf{negl}(\lambda)$$

Proof. We show this for the case when $b = 0$; the argument for the case when $b = 1$ symmetrically follows. Let $\widehat{M} \leftarrow \mathsf{Tomography}((\Phi_\lambda(k,i,0))^{\otimes L})$ be generated during the execution of $\mathsf{Verify}((k,i,0),M)$.

Let us condition on the event that $\|\widehat{M} - G_\lambda(k,i,0)\|_F^2 \leq \frac{9}{512}$ and $\|\widehat{M} - G_\lambda(k,i,1)\|_F^2 \leq \frac{9}{512}$. We prove that $\|M - \widehat{M}\|_F^2 > \frac{9}{512}$. Suppose not, then the following holds:

$$\begin{aligned}
\|G_\lambda(k,i\|0) - G_\lambda(k,i\|1)\|_F &\leq \|\widehat{M} - G_\lambda(k,i,0)\|_F + \|\widehat{M} - G_\lambda(k,i,1)\|_F + \|M - \widehat{M}\|_F \\
&\leq \sqrt{\frac{9}{512}} + \sqrt{\frac{9}{512}} + \sqrt{\frac{9}{512}} \leq 3\sqrt{\frac{9}{512}}
\end{aligned}$$

Thus, $\|G_\lambda(k,i\|0) - G_\lambda(k,i\|1)\|_F^2 \leq \frac{81}{512}$. This contradicts the fact that $\|G_\lambda(k,i\|0) - G_\lambda(k,i\|1)\|_F^2 > \frac{81}{512}$.

To summarise, conditioned on the event that $\|M - G_\lambda(k,i,0)\|_F^2 \leq \frac{9}{512}$ and $\|\widehat{M} - G_\lambda(k,i,1)\|_F^2 \leq \frac{9}{512}$, it holds that $\|M - \widehat{M}\|_F^2 > \frac{9}{512}$. Thus, Verify outputs $\mathsf{Invalid}$.

As we showed in the proof of Lemma 7.10, it follows that $\|\widehat{M} - G_\lambda(k,i,0)\|_F^2 \leq \frac{9}{512}$ and $\|\widehat{M} - G_\lambda(k,i,1)\|_F^2 \leq \frac{9}{512}$ holds with probability at least $(1 - \mathsf{negl}(\lambda))$.

Combining the above observations, we have the following: for every $k \in \mathcal{K}_{\mathsf{good}}$, for every bit b, we show that $\mathsf{Verify}((k,i,0),M)$ outputs $\mathsf{Invalid}$ with probability at least $(1 - \mathsf{negl}(\lambda))$. This completes the proof. □

All that is left is to show that $|\mathcal{K}_{\mathsf{good}}|$ is large enough.

Claim 10. $|\mathcal{K}_{\mathsf{good}}| \geq (1 - \mathsf{negl}(\lambda))2^\lambda$.

Proof. We invoke the security of $\{G_\lambda(\cdot,\cdot)\}$ to show this. Suppose the statement of the claim is false (that is $|\mathcal{K}_{\mathsf{good}}| \leq (1-\delta)2^\lambda$, where δ is non-negligible), we prove that $\{G_\lambda(\cdot,\cdot)\}$ is insecure.

Consider the following QPT distinguisher D that distinguishes whether it is given (classical) oracle access to $G_\lambda(\cdot,\cdot)$ or whether it is given access to an oracle, call it $\mathcal{O}_{\mathsf{Haar}}$, that on any input $x \in \{0,1\}^d$, outputs an iid Haar random state. D queries the oracle on all the inputs $\{0,1\}^{d+1}$ and for every input $x \in \{0,1\}^d$, it obtains L copies of two states $\rho_{x\|0}$ and $\rho_{x\|1}$. It then computes $M_{x\|0} \leftarrow \mathsf{Tomography}\left(\rho_{x\|0}^{\otimes L}\right)$ and $M_{x\|1} \leftarrow \mathsf{Tomography}\left(\rho_{x\|1}^{\otimes L}\right)$. It outputs 1 if there exists $x \in \{0,1\}^d$ such that $\|M_{x\|0} - M_{x\|1}\|_F^2 \leq \frac{15}{\sqrt{512}}$, otherwise it outputs 0.

We consider two cases below.

- D has oracle access to $\mathcal{O}_{\mathsf{Haar}}$: for any $x \in \{0,1\}^d$, from Fact 3.4, it follows that $\mathsf{Pr}[\|\rho_{x\|0} - \rho_{x\|1}\|_F^2 \leq \frac{15}{\sqrt{512}}] \leq \mathsf{negl}(\lambda)$. By union bound, it follows that the probability that there exists an $x \in \{0,1\}^d$ such that $\|\rho_{x\|0} - \rho_{x\|1}\|_F^2 \leq \frac{15}{\sqrt{512}}$ is negligible in λ. Thus, D outputs 1 with negligible probability.

- D has oracle access to $G_\lambda(k,\cdot)$, where $k \leftarrow \{0,1\}^\lambda$: Let us condition on the event that PRFS key $k \notin \mathcal{K}_{\mathsf{good}}$. This means that there exists an input $x \in \{0,1\}^d$ such that $\|G_\lambda(k,i\|0) - G_\lambda(k,i\|1)\|_F^2 \leq \frac{81}{512}$. Moreover, from Corollary 7.6, it follows that with probability at least $(1-\mathsf{negl}(\lambda))$, $\|M_{x\|0} - G_\lambda(k,x\|0)\|_F^2 \leq \frac{9}{512}$ and $\|M_{x\|1} - G_\lambda(k,i\|1)\|_F^2 \leq \frac{9}{512}$. Thus, with probability at least $1-\mathsf{negl}(\lambda)$, we have the following:

$$\begin{aligned}\|M_{x\|0} - M_{x\|1}\|_F &\leq \|M_{x\|0} - G_\lambda(k,x,0)\|_F + \|M_{x\|1} - G_\lambda(k,x,1)\|_F + \|G_\lambda(k,x,0) - G_\lambda(k,x,1)\|_F \\ &\leq \sqrt{\frac{9}{512}} + \sqrt{\frac{9}{512}} + \sqrt{\frac{81}{512}} = \frac{15}{\sqrt{512}}\end{aligned}$$

  Since the probability that $k \notin \mathcal{K}_{\mathsf{good}}$ is at least δ, we have that the probability that D outputs 1 is at least $\delta(1-\mathsf{negl}(\lambda))$, which is non-negligible in λ.

Since the difference in the probability that D outputs 0 in both the above cases is at least non-negligible, we have that D violates the security of PRFS. □

□

8 Applications

In this section, we show how to use PRFS to constrtuct a variety of applications:

1. Bit commitments with classical communication and,
2. Pseudo one-time pad schemes with classical communication.

To accomplish the above applications, we use verifiable tomography from Section 7.

8.1 Commitment scheme

We construct bit commitments with classical communication from pseudorandom function-like quantum states. We recall the definition by [AQY21].

A (bit) commitment scheme is given by a pair of (uniform) QPT algorithms (C, R), where $C = \{C_\lambda\}_{\lambda\in\mathbb{N}}$ is called the *committer* and $R = \{R_\lambda\}_{\lambda\in\mathbb{N}}$ is called the *receiver*. There are two phases in a commitment scheme: a commit phase and a reveal phase.

- In the (possibly interactive) commit phase between C_λ and R_λ, the committer C_λ commits to a bit, say b. We denote the execution of the commit phase to be $\sigma_{CR} \leftarrow \mathsf{Commit}\langle C_\lambda(b), R_\lambda\rangle$, where σ_{CR} is a joint state of C_λ and R_λ after the commit phase.
- In the reveal phase C_λ interacts with R_λ and the output is a trit $\mu \in \{0, 1, \bot\}$ indicating the receiver's output bit or a rejection flag. We denote an execution of the reveal phase where the committer and receiver start with the joint state σ_{CR} by $\mu \leftarrow \mathsf{Reveal}\langle C_\lambda, R_\lambda, \sigma_{CR}\rangle$.

We require that the above commitment scheme satisfies the correctness, computational hiding, and statistical binding properties below.

Definition 8.1 (Correctness)**.** *We say that a commitment scheme (C, R) satisfies correctness if*

$$\mathsf{Pr}\left[b^* = b \;:\; \substack{\sigma_{CR}\leftarrow\mathsf{Commit}\langle C_\lambda(b),R_\lambda\rangle,\\ b^*\leftarrow\mathsf{Reveal}\langle C_\lambda,R_\lambda,\sigma_{CR}\rangle}\right] \geq 1 - \nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function.

Definition 8.2 (Computational Hiding)**.** *We say that a commitment scheme (C, R) satisfies computationally hiding if for any malicious QPT receiver $\{R^*_\lambda\}_{\lambda\in\mathbb{N}}$, for any QPT distinguisher $\{D_\lambda\}_{\lambda\in\mathbb{N}}$, the following holds:*

$$\left|\Pr_{(\tau,\sigma_{CR^*})\leftarrow\mathsf{Commit}\langle C_\lambda(0),R^*_\lambda\rangle}[D_\lambda(\sigma_{R^*}) = 1] - \Pr_{(\tau,\sigma_{CR^*})\leftarrow\mathsf{Commit}\langle C_\lambda(1),R^*_\lambda\rangle}[D_\lambda(\sigma_{R^*}) = 1]\right| \leq \varepsilon(\lambda),$$

for some negligible $\varepsilon(\cdot)$.

Definition 8.3 (Statistical Binding)**.** *We say that a commitment scheme (C, R) satisfies statistical binding if for every QPT sender $\{C^*_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a (possibly inefficient) extractor $\mathcal{E}$ such that the following holds:*

$$\mathsf{Pr}\left[\mu \neq b^* \wedge \mu \neq \bot \;:\; \substack{(\tau,\sigma_{C^*R})\leftarrow\mathsf{Commit}\langle C^*_\lambda,R_\lambda\rangle,\\ b^*\leftarrow\mathcal{E}(\tau),\\ \mu\leftarrow\mathsf{Reveal}\langle C^*_\lambda,R_\lambda,\sigma_{C^*R}\rangle}\right] \leq \nu(\lambda),$$

where $\nu(\cdot)$ is a negligible function and τ is the transcript of the Commit *phase.*

Remark 8.4 (Comparison with [AQY21])**.** *In the binding definition of [AQY21], given the fact that the sender's and the receiver's state could potentially be entangled with each other, care had to be taken to ensure that after the extractor was applied on the receiver's state, the sender's state along with the decision bit remains (indistinguishable) to the real world. In the above definition, however, since the communication is entirely classical, any operations performed on the receiver's end has no consequence to the sender's state. As a result, our definition is much simpler than [AQY21].*

8.1.1 Construction

Towards constructing a commitment scheme with classical communication, we use a verifiable tomography from Figure 2.

Construction. We present the construction in Figure 4. In the construction, we require $d(\lambda) = \lceil \log \frac{3\lambda}{n} \rceil \geq 1$.

Commit(b):

- The reciever R_λ samples an m-qubit Pauli $P = \bigotimes_{x \in \{0,1\}^d} P_x$ where $m = 2^d n$. It sends P to the commiter.
- The committer C_λ on intput $b \in \{0,1\}$ does the following:
  - Sample $k \xleftarrow{\$} \{0,1\}^\lambda$.
  - For all $x \in \{0,1\}^d$
    - $*$ Generate $\sigma_x^{\otimes L} \leftarrow \left(\Phi_\lambda \left(P_x||k||x||b\right)\right)^{\otimes L}$, where $L = 3^8 2^{3n+5}\lambda$.
    - $*$ $M_x \leftarrow \mathsf{Tomography}\left(\sigma_x^{\otimes L}\right)$.
  - Send $M = (M_x)_{x \in \{0,1\}^d}$ to the reciever.

Reveal:

- The commiter sends (k, b) as the decommitment. If $b \notin \{0,1\}$, the reciever outputs $\perp$. Output b if for each $x \in \{0,1\}^d$, $\mathsf{Verify}(P_x||k||x||b, M) = \mathsf{Valid}$, output $\perp$ otherwise.

Figure 4: Commitment scheme

We prove that the construction in Figure 4 satisfies correctness, computational hiding and statistical binding properties.

Lemma 8.5 (Correctness)**.** *The commitment scheme in Figure 4 satisfies correctness.*

Proof. This follows from Lemma 7.7. □

Lemma 8.6 (Computational Hiding)**.** *The commitment scheme in Figure 4 satisfies computational hiding.*

Proof. We prove the security via a hybrid argument. Fix $\lambda \in \mathbb{N}$. Consider a QPT adversary R_λ^*.

Hybrid $H_{1,b}$, for all $b \in \{0,1\}$. This corresponds to C commiting to the bit b.

Hybrid $H_{2,b}$, for all $b \in \{0,1\}$. This hybrid is the same as before except that for all $x \in \{0,1\}^d$, $\Phi_\lambda\left(P||k||x||b\right)$ replaced with $\left(\left(I \otimes P_x^b\right)\left(|0\rangle\langle 0| \otimes |\vartheta_x\rangle\langle\vartheta_x|\right)\left(I \otimes P_x^b\right)\right)$ where $|\vartheta_1\rangle, ..., |\vartheta_{2^d}\rangle \leftarrow \mathscr{H}_n$.

The hybrids $H_{1,b}$ and $H_{2,b}$ are computationally indistinguishable because of the security of $PRFS$. $H_{2,0}$ and $H_{2,1}$ are identical by the unitary invariance property of Haar distribution. Hence, $H_{1,0}$ and $H_{1,1}$ are computationally indistinguishable.

□

Lemma 8.7 (Statistical Binding)**.** *The commitment scheme in Figure 4 satisfies $O(2^{-0.5\lambda})$-statistical binding.*

Proof of Lemma 8.7. Let $C^* = \{C^*_\lambda\}_{\lambda\in\mathbb{N}}$ be a malicous committer. Execute the commit phase between C^*_λ and R_λ. Let τ be the classical transcript and let σ_{C^*R} be the joint state of C^*R. We first provide the description of an extractor.

Description of $\mathcal{E}$. On the input $\tau = (P, M)$, the extractor does the following:

1. For all $k'||b' \in \{0,1\}^\lambda \times \{0,1\}$, run for all $x \in \{0,1\}^d$, $\mathsf{Verify}(P_x||k'||x||b', M)$.

2. If for all $x \in \{0,1\}^d$, $\mathsf{Verify}(P||k'||x||b', M) = \mathsf{Valid}$, output b'.

3. Else output $\perp$.

Fact 8.8. *Let $\mathcal{P}_m$ be the m-qubit Pauli group. Then,*

$$\Pr_{P \xleftarrow{\$} \mathcal{P}_m}\left[\exists k_0, k_1 : \forall x \in \{0,1\}^d, |\langle\psi_{k_0,x}|\, P_x \,|\psi_{k_1,x}\rangle|^2 \geq \delta\right] \leq \delta^{-2^d} 2^{2\lambda - m}.$$

Proof. We use the following fact [AQY21, Fact 6.9]: Let $|\psi\rangle$ and $|\phi\rangle$ be two arbitrary n-qubit states. Then,

$$\mathop{\mathbb{E}}_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[|\langle\psi|\, P_x \,|\phi\rangle|^2\right] = 2^{-n}.$$

For any k_0, k_1, x by the above fact, $\mathbb{E}_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[|\langle\psi_{k_0,x}|\, P_x \,|\psi_{k_1,x}\rangle|^2\right] = 2^{-n}$. Using Markov's inequality we get that for all $\delta > 0$,

$$\Pr_{P_x \xleftarrow{\$} \mathcal{P}_n}\left[|\langle\psi_{k_0,x}|\, P_x \,|\psi_{k_1,x}\rangle|^2 \geq \delta\right] \leq \delta^{-1} 2^{-n}.$$

Since, all P_x's are independent,

$$\Pr_{P \xleftarrow{\$} \mathcal{P}_m}\left[\forall x \in \{0,1\}^d, |\langle\psi_{k_0,x}|\, P_x \,|\psi_{k_1,x}\rangle|^2 \geq \delta\right] \leq \left(\delta^{-1} 2^{-n}\right)^{2^d}.$$

Using a union bound over all k_0, k_1,

$$\Pr_{P \xleftarrow{\$} \mathcal{P}_m}\left[\exists k_0, k_1 : \forall x \in \{0,1\}^d, |\langle\psi_{k_0,x}|\, P_x \,|\psi_{k_1,x}\rangle|^2 \geq \delta\right] \leq \delta^{-2^d} 2^{2\lambda - m}.$$

□

Let the transcript be (P, M) where P is chosen uniformly at random. Let

$$p = \mathsf{Pr}\left[\mu \neq b^* \wedge \mu \neq \perp \;:\; \begin{array}{c} (\tau, \sigma_{C^*R}) \leftarrow \mathsf{Commit}\langle C^*_\lambda, R_\lambda\rangle, \\ b^* \leftarrow \mathcal{E}(\tau), \\ \mu \leftarrow \mathsf{Reveal}\langle \tau, \sigma_{C^*R}\rangle \end{array}\right]$$

Then

$$p = \Pr_{P \xleftarrow{\$} \mathcal{P}_m} \left[\exists k_0, k_1, b_0, b_1 \ : \forall x \in \{0,1\}^d \ \substack{\mathsf{Verify}(P_x||k_0||x||b_0, M_x) = \mathsf{Valid}, \\ \mathsf{Verify}(P_x||k_1||x||b_1, M_x) = \mathsf{Valid}, \\ b_0 \neq b_1} \right].$$

Without loss of generality we can assume $b_0 = 0$ and $b_1 = 1$,

$$p = \Pr_{P \xleftarrow{\$} \mathcal{P}_m} \left[\exists k_0, k_1 \ : \forall x \in \{0,1\}^d \ \substack{\mathsf{Verify}(P_x||k_0||x||0, M_x) = \mathsf{Valid}, \\ \mathsf{Verify}(P_x||k_1||x||1, M_x) = \mathsf{Valid}} \right].$$

By Lemma 7.8,

$$p \leq \Pr_{P \xleftarrow{\$} \mathcal{P}_m} \left[\exists k_0, k_1 \ : \forall x \in \{0,1\}^d, \ Tr(P_x \left|\psi_{k_1,x}\right\rangle\!\left\langle\psi_{k_1,x}\right| P_x \left|\psi_{k_0,x}\right\rangle\!\left\langle\psi_{k_0,x}\right|) \geq 542/729 \right]$$

By Fact 8.8,

$$p \leq \frac{729}{542}^{2^d} \left(2^{2\lambda - m}\right).$$

For $m \geq 3\lambda$, the protocol satisfies $O(2^{-0.5\lambda})$-statistical binding.

□

8.2 Encryption scheme

We construct a psuedo one-time pad scheme with classical communication from psuedorandom function-like quantum states. We first present the definition below.

Definition 8.9 (Psuedo One-Time Pad)**.** *We say that a pair of QPT algorithms* $(\mathsf{Enc}, \mathsf{Dec})$ *is a psuedo one-time pad if the following properties are satisfied: there exists a polynomial* $M(\lambda)$ *such that*

Correctness: *There exists a negligible function* $\varepsilon(\cdot)$ *such that for every* λ*, every* $x \in \{0,1\}^{M(\lambda)}$*,*

$$\Pr_{\substack{k \leftarrow \{0,1\}^\lambda \\ \mathsf{ct} \leftarrow \mathsf{Enc}_\lambda(k,x)}} [\mathsf{Dec}_\lambda(k, \mathsf{ct}) = x] \geq 1 - \varepsilon(\lambda).$$

Security: *For every* λ*, for every QPT adversary* A_λ*, there exists a negligible function* $\varepsilon(\cdot)$ *such that* $x_1, x_2 \in \{0,1\}^{M(\lambda)}$*,*

$$\left| \Pr_{\substack{k \leftarrow \{0,1\}^\lambda \\ \mathsf{ct} \leftarrow \mathsf{Enc}_\lambda(k,x_1)}} [A_\lambda(\mathsf{ct}) = 1] - \Pr_{\substack{k \leftarrow \{0,1\}^\lambda \\ \mathsf{ct} \leftarrow \mathsf{Enc}_\lambda(k,x_2)}} [A_\lambda(\mathsf{ct}) = 1] \right| \leq \varepsilon(\lambda).$$

8.2.1 Construction

Towards constructing an encryption scheme with classical communication, we use the verifiable tomography $(\mathsf{Tomography}, \mathsf{Verify})$ described in Figure 3 satisfying L-same-input correctness and $(L, \Pi, \mathcal{D})$-distributional different-input correctness. Let the set of channels associated with $(\mathsf{Tomography}, \mathsf{Verify})$ be $\mathcal{C} = \{\Phi_\lambda \ : \ \lambda \in \mathbb{N}\}$. Recall that $\mathcal{C}$ is associated with (d, n)-PRFS. Refer to Figure 3 for the description of $L, \Pi, \mathcal{D}$ and Φ_λ.

$\mathsf{Enc}_\lambda(k, x)$:

- Parse $x = x_1 \cdots x_M$.
- For $i \in [M]$, generate $L(\lambda)$ copies of $\sigma_i \leftarrow (\Phi_\lambda((k, i, x_i)))$.
- Generate $u_i \leftarrow \mathsf{Tomography}\left(\sigma_i^{\otimes L(\lambda)}\right)$.
- Output the ciphertext $\mathsf{ct} = \{u_i\}_{i \in [M]}$.

$\mathsf{Dec}_\lambda(k, \mathsf{ct})$:

- Parse ct as $\{u_i\}_{i \in [M]}$.
- For $i \in [M]$, run $(\mathsf{Verify}(k||i||0, u_i))$. If Valid, set $x_i = 0$, else set $x_i = 1$.
- Output $x = x_1 \cdots x_M$.

Figure 5: Encryption scheme

Construction. We present the construction in Figure 5.

Theorem 8.10. $(\mathsf{Enc}_\lambda, \mathsf{Dec}_\lambda)$ *satisfies the correctness property of a psuedo one-time pad.*

Proof. Let $\mathsf{ct} \leftarrow \mathsf{Enc}_\lambda(k, x)$. Parse ct as $\{u_i\}_{i \in [l]}$.

For every $i \in [M]$, the following holds:

- From Lemma 7.10, $\mathsf{Verify}(k||i||x_i, u_i)$ outputs Valid with probability $1 - \mathsf{negl}(\lambda)$.
- From Lemma 7.11, $\mathsf{Verify}(k||i||(1 - x_i), u_i)$ outputs Valid with probability at most $\mathsf{negl}(\lambda)$.

Thus, for any given i, the decryption algorithm can correctly determine x_i with probability at least $1 - \mathsf{negl}(\lambda)$. By union bound, it then follows that the probability that the decryption algorithm correctly determines all bits of x is negligibly close to 1. □

Theorem 8.11. $(\mathsf{Enc}_\lambda, \mathsf{Dec}_\lambda)$ *satisfies the security property of a pseudo one-time pad.*

Proof. We prove the security via a hybrid argument. Fix λ and the messages $x_0, x_1 \in \{0,1\}^M$. Consider a QPT adversary A_λ.

Hybrid $H_{1,b}$, for $b \in \{0,1\}$. Sample $k \leftarrow \{0,1\}^\lambda$. Compute $\mathsf{ct} \leftarrow \mathsf{Enc}_\lambda(k, x_b)$. Output ct.

Hybrid H_2. Sample $L(\lambda)$ copies of n-qubit Haar-random states $|\vartheta_1\rangle, ..., |\vartheta_l\rangle \leftarrow \mathscr{H}_n$. For every $i \in [M]$, compute $u_i \leftarrow \mathsf{Tomography}\left(|\vartheta_1\rangle^{\otimes L}\right)$. The output of this hybrid is $(u_1, \ldots, u_l)$.

The hybrids $H_{1,b}$, for $b \in \{0,1\}$, and H_2 are computationally indistinguishable from the security of PRS. □

Acknowledgements

The authors would like to thank the anonymous TCC 2022 reviewers for their helpful comments. The authors would also like to thank Fermi Ma for his suggestions that improved the bounds and the analysis in Section 5.1.

References

[AE07] Andris Ambainis and Joseph Emerson. "Quantum t-designs: t-wise Independence in the Quantum World". In: *22nd Annual IEEE Conference on Computational Complexity (CCC 2007), 13-16 June 2007, San Diego, California, USA*. 2007, pp. 129–140. DOI: 10.1109/CCC.2007.26 (cit. on pp. 3, 24).

[AKN98] Dorit Aharonov, Alexei Y. Kitaev, and Noam Nisan. "Quantum Circuits with Mixed States". In: *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*. 1998, pp. 20–30. DOI: 10.1145/276698.276708 (cit. on p. 52).

[AQY21] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. *Cryptography from Pseudorandom Quantum States*. 2021. arXiv: 2112.10020v2 (cit. on pp. 3–5, 7, 8, 14–17, 32, 41, 45, 47).

[BHH16] Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. "Local Random Quantum Circuits are Approximate Polynomial-Designs". In: *Communications in Mathematical Physics* 346.2 (2016), pp. 397–434. DOI: 10.1007/s00220-016-2706-8 (cit. on p. 11).

[BS19] Zvika Brakerski and Omri Shmueli. "(Pseudo) Random Quantum States with Binary Phase". In: *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*. Vol. 11891. 2019, pp. 229–250. DOI: 10.1007/978-3-030-36030-6_10 (cit. on pp. 3, 6, 18, 22, 25).

[BS20] Zvika Brakerski and Omri Shmueli. "Scalable Pseudorandom Quantum States". In: *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part II*. Vol. 12171. 2020, pp. 417–440. DOI: 10.1007/978-3-030-56880-1_15 (cit. on pp. 3–5, 14, 34).

[DCEL09] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. "Exact and approximate unitary 2-designs and their application to fidelity estimation". In: *Phys. Rev. A* 80 (1 2009), p. 012304. DOI: 10.1103/PhysRevA.80.012304 (cit. on pp. 3, 24).

[Gav12] Dmitry Gavinsky. "Quantum Money with Classical Verification". In: *Proceedings of the 27th Conference on Computational Complexity, CCC 2012, Porto, Portugal, June 26-29, 2012*. 2012, pp. 42–52. DOI: 10.1109/CCC.2012.10 (cit. on p. 4).

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. "How to Construct Random Functions". In: *J. ACM* 33.4 (1986), pp. 792–807. DOI: 10.1145/6490.6503 (cit. on pp. 3, 4).

[Har13] Aram W. Harrow. *The Church of the Symmetric Subspace*. 2013. arXiv: 1308.6595 (cit. on p. 18).

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. "A Pseudorandom Generator from any One-way Function". In: *SIAM Journal on Computing* 28.4 (1999), pp. 1364–1396. DOI: 10.1137/S0097539793244708 (cit. on p. 3).

[HKP20] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements". In: *Nature Physics* 16.10 (2020), pp. 1050–1057. DOI: 10.1038/s41567-020-0932-7 (cit. on p. 34).

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. "Pseudorandom Quantum States". In: *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*. Vol. 10993. 2018, pp. 126–152. DOI: 10.1007/978-3-319-96878-0_5 (cit. on pp. 3, 6, 13, 14, 18, 25, 31).

[Kre21] William Kretschmer. "Quantum Pseudorandomness and Classical Complexity". In: *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, July 5-8, 2021, Virtual Conference*. Vol. 197. 2021, 2:1–2:20. DOI: 10.4230/LIPIcs.TQC.2021.2 (cit. on pp. 3–5, 32, 34).

[LC97] Hoi-Kwong Lo and H. F. Chau. "Is Quantum Bit Commitment Really Possible?" In: *Phys. Rev. Lett.* 78 (17 1997), pp. 3410–3413. DOI: 10.1103/PhysRevLett.78.3410 (cit. on p. 5).

[Low21] Lowe, Angus. "Learning Quantum States Without Entangled Measurements". MA thesis. University of Waterloo, 2021. HDL: 10012/17663 (cit. on pp. 8, 36).

[May97] Dominic Mayers. "Unconditionally Secure Quantum Bit Commitment is Impossible". In: *Phys. Rev. Lett.* 78 (17 1997), pp. 3414–3417. DOI: 10.1103/PhysRevLett.78.3414 (cit. on p. 5).

[MY21] Tomoyuki Morimae and Takashi Yamakawa. *Quantum commitments and signatures without one-way functions*. 2021. arXiv: 2112.06369 (cit. on pp. 3, 4).

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010. DOI: 10.1017/CBO9780511976667 (cit. on p. 9).

[RS09] Aidan Roy and A. J. Scott. "Unitary designs and codes". In: *Designs, Codes and Cryptography* 53.1 (2009), pp. 13–31. DOI: 10.1007/s10623-009-9290-2 (cit. on p. 3).

[Zha12a] Mark Zhandry. "How to Construct Quantum Random Functions". In: *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*. 2012, pp. 679–687. DOI: 10.1109/FOCS.2012.37 (cit. on pp. 12, 13, 52).

[Zha12b] Mark Zhandry. "Secure Identity-Based Encryption in the Quantum Random Oracle Model". In: *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*. Vol. 7417. 2012, pp. 758–775. DOI: 10.1007/978-3-642-32009-5_44 (cit. on p. 13).

[Zha15] Mark Zhandry. "A note on the quantum collision and set equality problems". In: *Quantum Inf. Comput.* 15.7&8 (2015), pp. 557–567. DOI: 10.26421/QIC15.7-8-2 (cit. on p. 32).

A Small-Range Distributions over Unitary Operators

Let $\mathcal{U}$ be a distribution over unitary operators over a finite Hilbert space, r be a positive integer and $\mathcal{X}$ be a finite set. (Looking ahead, we are going to sample a unitary from a distribution and potentially invoke this unitary multiple times. This is not a channel as all invocations are going to use the same unitary.) We define $\mathcal{U}^{\mathcal{X}} := \bigoplus_{x\in\mathcal{X}} U_x$ to be the unitary that maps $|x\rangle \otimes |y\rangle \to |x\rangle \otimes U_x |y\rangle$, with each U_x being an independent sample from $\mathcal{U}$. We define small-range distributions $\mathsf{SR}_r^{\mathcal{U}}(\mathcal{X})$ sampled as follows:

- For each $i \in [r]$, sample a U_i from $\mathcal{U}$.
- For each $x \in \mathcal{X}$, sample a random $i_x \in [r]$, so that the unitary maps $|x\rangle \otimes |y\rangle \to |x\rangle \otimes U_{i_x} |y\rangle$ for any state $|y\rangle$.

Theorem A.1 ([Zha12a, Corollary VII.5]). *Let D be a distribution over bit strings , and let $\mathcal{U}$ be the unitary distribution for random oracles , i.e. it corresponds to D where for each bit string m, we associate a unitary that maps $|y\rangle \to |y \oplus m\rangle$ in the computational basis. The output distributions of a quantum algorithm making q quantum queries to an oracle either drawn from $\mathsf{SR}_r^D(\mathcal{X})$ from $D^{\mathcal{X}}$ are $\ell q^3/r$-close, where $\ell = 8\pi^2/3 < 27$.*

The goal of this section is to generalize this theorem to the setting where D is a (possibly infinite-support) distribution of arbitrary quantum unitary oracles instead of a distribution over (finite-length) bit strings.

Generalizing the proof of Zhandry's theorem to this setting seems difficult. For example, the proof uses the fact that the query oracle operates in the computational basis, which a generic unitary might not satisfy. Therefore, we instead are going to use a downsampling trick, which allows us to invoke Zhandry's theorem as a black box. Furthermore, we are going to crucially rely on the fact that the bound given by Zhandry's theorem is independent of the size of D or $\mathcal{X}$.

We invoke some inequalities about the diamond norm. The diamond norm intuitively captures the best distinguishing advantage from a single invocation of one of two channels. We refer the readers to the work of Aharonov et al. [AKN98] for the formal definitions.

Lemma A.2 ([AKN98, lemma 12.6]). *Let V, W be operators. If $\|V\|, \|W\| \leq 1$, then*

$$\left\| V(\cdot)V^\dagger - W(\cdot)W^\dagger \right\|_\diamond \leq 2\|V - W\|.$$

Lemma A.3 ([AKN98, lemma 13]). *Let T_1, T_2, T_1', T_2' be super-operators with norm ≤ 1, such that $\left\| T_j' - T_j \right\|_\diamond \leq \varepsilon_j$ for $j = 1, 2$. Then $\|T_2'T_1' - T_2T_1\|_\diamond \leq \varepsilon_1 + \varepsilon_2$.*

Lemma A.4 (Finite-support downsampling lemma). *For any distribution $\mathcal{U}$ over unitary operators over a finite Hilbert space and any $\varepsilon > 0$, there exists a distribution $\mathcal{F}$ with finite support such that any q-query (otherwise unbounded) quantum algorithm cannot distinguish a unitary from either distribution with advantage more than εq.*

Proof. Let $\mathcal{N}$ be the epsilon-net for the unitary group with $|\mathcal{N}| < \infty$ such that for any unitary U, there exists an approximation $U' \in \mathcal{N}$ such that $\|U - U'\| \leq \frac{\varepsilon}{2}$. We construct $\mathcal{F}$ by mapping every unitary from $\mathcal{U}$ to its approximation in $\mathcal{N}$. It follows from Lemmas A.2 and A.3 that for any fixed U, U', any q-query quantum algorithm's distinguishing advantage is at most εq. Therefore the lemma follows by averaging over sampling from $\mathcal{U}$. □

Lemma A.5 (Uniform downsampling lemma)**.** *For any distribution $\mathcal{U}$ over unitary operators over a finite Hilbert space and any $\varepsilon > 0$, there exists an integer $n > 0$ and a family of unitaries $\mathcal{T} = (U_1, ..., U_{2^n})$ such that any q-query quantum algorithm cannot distinguish a unitary from a uniformly sampled U_i from $\mathcal{T}$, or $\mathcal{U}$ with advantage more than εq.*

Proof. Let $\mathcal{F}$ be the distribution guaranteed by Lemma A.4 with distinguishing advantage $\varepsilon q/2$. Consider a random family $\mathcal{T}$ where each entry is an independent sample from $\mathcal{F}$. By the law of large numbers, the expected total variation distance to $\mathcal{F}$ goes to 0 as n goes to infinity, which implies the existence of a sequence of families whose distance is at most $\varepsilon/2 \leq \varepsilon q/2$. Therefore, by triangular inequality, the overall distinguishing advantage is at most εq. □

This lemma shows that we can downsamples an arbitrary distribution to a uniform distribution (with potentially huge support) followed by postprocessing, which is compatible with Zhandry's theorem. We do not make explicit the support of the downsampled distribution, and this is also not necessary as the loss from Zhandry's theorem is actually independent from the size of the support. Indeed, we combine these to show that we can generalize Zhandry's theorem over arbitrary unitary distribution with (asymptotically) the same loss.

Theorem A.6. *The output distributions of a quantum algorithm making q queries to an oracle either drawn from $\mathsf{SR}_r^{\mathcal{U}}(\mathcal{X})$ from $\mathcal{U}^{\mathcal{X}}$ are $300q^3/r$-close.*

Proof. Let $\varepsilon = \frac{q^2}{r|\mathcal{X}|}$, and let $\mathcal{T}$ be the distribution with length n guaranteed by Lemma A.5. We show the theorem by the following sequence of hybrids:

Hybrid 0 The oracle is drawn from $\mathcal{U}^{\mathcal{X}}$.

Hybrid 1 The oracle is drawn from $\mathcal{T}^{\mathcal{X}}$. Since the oracle is essentially a direct sum of $|\mathcal{X}|$ independent unitary samples, the distinguishing advantage between the output distributions of this and the last hybrid is at most $\varepsilon q \cdot |\mathcal{X}| = q^3/r$.

Hybrid 2 We instead sample a classical random oracle $(\{0,1\}^n)^{\mathcal{X}}$, and simulate the distribution of $\mathcal{T}^{\mathcal{X}}$ using two queries to the random oracle. In particular, we first query the random oracle to get $i \in \{0,1\}^n$, apply U_i, and then uncompute i in superposition. By construction, this is perfectly indistinguishable to the last hybrid.

Hybrid 3 We change the classical random oracle to be instead sampled by $\mathsf{SR}_r^{\{0,1\}^n}(\mathcal{X})$. By Theorem A.1 and the construction, the distance to the last hybrid is at most $27(2q)^3/r$.

Hybrid 4 We switch the oracle to $\mathsf{SR}_r^{\mathcal{T}}(\mathcal{X})$. This is perfectly indistinguishable to the last hybrid.

Hybrid 5 We switch the oracle to $\mathsf{SR}_r^{\mathcal{U}}(\mathcal{X})$. The distance for this is again at most q^3/r.

By triangle inequality, the two distributions are $< 300q^3/r$ indistinguishable. □