Cryptography from Pseudorandom Quantum States

Prabhanjan Ananth* *UC Santa Barbara*
Luowen Qian† *Boston University*
Henry Yuen‡ *Columbia University*

Abstract

Pseudorandom states, introduced by Ji, Liu and Song (Crypto'18), are efficiently-computable quantum states that are computationally indistinguishable from Haar-random states. One-way functions imply the existence of pseudorandom states, but Kretschmer (TQC'20) recently constructed an oracle relative to which there are no one-way functions but pseudorandom states still exist. Motivated by this, we study the intriguing possibility of basing interesting cryptographic tasks on pseudorandom states.

We construct, assuming the existence of pseudorandom state generators that map a λ-bit seed to a $\omega(\log \lambda)$-qubit state, (a) statistically binding and computationally hiding commitments and (b) pseudo one-time encryption schemes. A consequence of (a) is that pseudorandom states are sufficient to construct maliciously secure multiparty computation protocols in the dishonest majority setting.

Our constructions are derived via a new notion called *pseudorandom function-like states* (PRFS), a generalization of pseudorandom states that parallels the classical notion of pseudorandom functions. Beyond the above two applications, we believe our notion can effectively replace pseudorandom functions in many other cryptographic applications.

*prabhanjan@cs.ucsb.edu
†luowenq@bu.edu
‡hyuen@cs.columbia.edu

Contents

1 Introduction

Assumptions are the bedrock of designing provably secure cryptographic constructions. Over the years, theoretical cryptographers have pondered over the precise assumptions needed to achieve cryptographic tasks, often losing sleep over this [Kil88]. The celebrated work of Goldreich [Gol90] shows that most interesting cryptographic tasks (encryption, commitments, pseudorandom generators, etc.) imply the existence of one-way functions – i.e., functions that can be efficiently computed in the forward direction but cannot be efficiently inverted. Thus it appears that the existence of one-way functions is a *minimal* assumption in cryptography.

Quantum information processing presents new opportunities for cryptography. Specifically, in many contexts, the assumptions necessary for cryptographic tasks can be weakened with the help of quantum resources. To illustrate, the seminal work of Bennett and Brassard [BB84] showed that key-exchange can be achieved unconditionally — i.e., without any computational assumptions — using quantum communication. In contrast, key-exchange is known to require computational assumptions if one can only use classical communication. More recently, the works of Grilo, Lin, Song and Vaikuntanathan [GLSV21] and Bartusek, Coladangelo, Khurana, and Ma [BCKM21b] demonstrate that quantum protocols for secure multiparty computation can be constructed from post-quantum one-way functions. On the other hand classical protocols for secure computation cannot be based (in a black-box way) on one-way functions alone [IR89].

These examples suggest that we revisit our belief about the necessity of cryptographic assumptions for quantum cryptographic tasks; that is, tasks that make use of quantum resources (computing ability and communication channels). Specifically, it is not even clear whether one-way functions is a necessary assumption for quantum cryptographic tasks — Goldreich's result [Gol90] only applies to cryptographic primitives and protocols with classical communication.

Our work continues the research agenda carried out by our predecessors [Wie83, BB84, BBCS91, GLSV21, BCKM21b]: *can we achieve cryptographic tasks using quantum communication in a world without one-way functions*[1]*?*

Pseudorandom States, Revisited. Towards understanding the above question, we revisit the notion of pseudorandom quantum states (abbreviated PRS) introduced by Ji, Liu and Song [JLS18]. A *PRS generator* G is a quantum polynomial-time (QPT) algorithm that, given input a key $k \in \{0,1\}^\lambda$, outputs a n-qubit quantum state $|\psi_k\rangle$ satisfying the following guarantee: any polynomial number (in λ) of copies of $|\psi_k\rangle$ is computationally indistinguishable from a polynomial number of copies of a state $|\vartheta\rangle$ that is sampled from the n-qubit *Haar distribution* (i.e. the uniform distribution over n-qubit pure states). Ji, Liu and Song [JLS18] (and subsequently improved by Brakerski and Shmueli [BS19, BS20]) show the existence of PRS assuming post-quantum one-way functions.

This notion is analogous to the pseudorandom generators (PRGs) from classical cryptography which take as input a short seed, say of length λ, and deterministically outputs a larger string of length $n > \lambda$ that is computationally indistinguishable from a string sampled from the uniform distribution. Despite the analogy, it has not been obvious whether pseudorandom quantum states have much cryptographic utility (unlike PRGs, which are ubiquitous in cryptography). Understanding the consequences of pseudorandom quantum states is particularly important in light of a recent result by Kretschmer [Kre21], who showed that there is a relativized world where $BQP = QMA$ (and

[1] Both the works [GLSV21, BCKM21b] explicitly raised the question of basing secure computation on assumptions weaker than one-way functions.

thus post-quantum one-way functions do not exist) while pseudorandom states exist. Kretschmer's result motivates us to focus the aforementioned research agenda on the following question: *what cryptographic tasks can be based solely on pseudorandom quantum states?*

1.1 Our Results

Our contributions in a nutshell are as follows:

- We propose a new notion called *pseudorandom function-like quantum states (PRFS)*.
- Using PRFS, we show how to build (a) statistically binding commitments and (b) pseudo one-time encryption schemes. As a consequence of (a), we obtain maliciously secure computation in the dishonest majority setting.
- Finally, we show that for a certain range of parameters – same as what is needed for the above applications – we can construct PRFS from a PRS.

Before we present the definition of PRFS, we first highlight the need for defining a new notion by describing the challenges for constructing primitives directly from PRS.

1.1.1 Challenges For Basing Primitives On PRS

Although the closest classical analogue of a PRS generator is a PRG, the analogy breaks down in several critical ways. This makes it challenging to use PRS generators in the same way that PRGs are used throughout cryptography.

Specifically, PRS generators appear very *rigid*, meaning that it seems challenging to take an existing PRS generator and generically increase or decrease its output length. Moreover, it is difficult to use output qubits of a PRS generator independently.

Inability to Stretch the Output. A fundamental result about PRGs is that their *stretch* (the output length as a function of the key length) can be amplified arbitrarily. In other words, given a PRG G that maps λ random bits to at least $\lambda+1$ pseudorandom bits, one can construct a PRG G' with λ^c output bits for arbitrarily large c. This fact is implicitly used everywhere in cryptography; specifically, it gives us the flexibility to choose the appropriate stretch of PRG relevant for the application without having to worry about the underlying hardness assumptions.

If PRS generators are analogous to PRGs, then one would expect that a similar amplification result to hold: the existence of PRS with nontrivial output length would (hopefully) imply the existence of PRS with arbitrarily large output length. The natural approach to amplify the stretch of a PRG by iteratively composing it with itself does not immediately work with PRS for a number of reasons; for one, a PRS generator takes as input a classical key while its output is a quantum state!

Inability to Shrink the Output. To add insult to injury, it is not even obvious how to *shrink* the output length of a PRS generator; this was also observed by Brakerski and Shmueli [BB21]. Classically, one can always discard bits from the output of a PRG, and the result is still obviously a PRG. However, discarding a single qubit of an n-qubit pseudorandom state $|\psi_k\rangle$ will leave a mixed state that is efficiently distinguishable from a $(n-1)$-qubit Haar-random state.

Inability to Separate the Output. Since the PRS output is highly entangled, it seems difficult to use the individual output qubits. As an example, suppose we want to encrypt a message of length ℓ. In the classical setting, an ℓ-bit output PRG can be used to encrypt a message of length ℓ by xor-ing the i^{th} PRG output bit with the i^{th} bit of the message. Implicitly, we are using the fact that the output of a PRG can be viewed a tensor product of bits and this feature of classical PRGs is mirrored by our notion of PRFS (explained next). On the other hand, if we have a single (entangled) PRS state (irrespective of the number of qubits it represents), it is unclear how to use each qubit to encode a bit; any operations performed on a single qubit could affect the other qubits that are entangled with this qubit.

1.1.2 New Notion: Pseudorandom Function-Like States

Pseudorandom function-like states (abbreviated PRFS) is a generalization of PRS, where the same key k is used to generate many pseudorandom states. In more details, a (d, n)-*PRFS generator* G is a QPT algorithm that, given as input a key $k \in \{0,1\}^\lambda$ and an input $x \in \{0,1\}^d$, outputs a n-qubit quantum state $|\psi_{k,x}\rangle$, satisfying the following pseudorandomness property: no QPT adversary can distinguish between multiple copies of the output states $\left(|\psi_{k,x_1}\rangle^{\otimes t}, \ldots, |\psi_{k,x_s}\rangle^{\otimes t}\right)$ from a collection of states $\left(|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t}\right)$ where each $|\vartheta_i\rangle$ is sampled independently from the Haar distribution; furthermore the indistinguishability holds even if the adversary knows the inputs $x_1, \ldots, x_s$. (See Section 3 for a formal definition of PRFS generators).

An Alternate Perspective: Tensor Product PRS generators. If PRS generators are analogous to classical pseudorandom generators, then PRFS generators are analogous to classical pseudorandom *functions* (hence the name pseudorandom *function-like*). A PRS generator outputs a single state per key k. On the other hand, we can think of PRFS as a *relaxed* notion of PRS generator that on input k outputs a *tensor product* of states $|\psi_0\rangle \otimes |\psi_1\rangle \otimes \cdots \otimes |\psi_{2^d-1}\rangle$ where each $|\psi_i\rangle$, is indistinguishable from a Haar-random state.

The tensor product feature is quite useful in applications. Let us revisit our earlier example: suppose we want to encrypt a message of length ℓ. If we have a tensor product of ℓ PRS states then we can use each state to encode one bit of the message: if the i^{th} bit is 0 then send the i^{th} PRFS state, otherwise send a random state. The decryptor, using the PRFS key, can decode the message by distinguishing between a PRFS state and a random state.

Additional Observations. Some additional observations of PRFS are in order:

- Assuming one-way functions, we can generically construct (d, n)-PRFS from any n-qubit PRS for any polynomial d, n. To compute PRFS on key k and input x, first compute a classical PRF on (k, x) and use the resulting output as a key for the n-qubit PRS. Since n-qubit PRS can be based on (post-quantum) one-way functions [JLS18, BS20], this shows that even PRFS can be based on (post-quantum) one-way functions.

- In the other direction, we can construct n-qubit PRS from any (d, n)-qubit PRFS. On input k, the PRS simply outputs the result of PRFS on input $(k, 0)$.

- Another interesting aspect about PRFS is that it too, like PRS, is separated from (post-quantum) one-way functions. This can be obtained by a generalization of Kretschmer's result [AQY21].

1.1.3 Implications

We show that PRFS can effectively replace the usage of pseudorandom generators and pseudorandom functions in many primitives one learns about in "Cryptography 101".

Specifically, we focus on two applications.

Implication 1. One-time Encryption with Short Keys and Long Messages. As a starter illustration of the usefulness of PRFS, we construct from a PRFS generator G a one-time encryption scheme for classical messages. The important feature of this construction is the fact that the message length is much larger than the key length. This is impossible to achieve information-theoretically, even in the quantum setting. This type of one-time encryption schemes, also referred to as *pseudo one-time pad*, is already quite useful, as it implies garbling schemes for P/poly [BMR90] and even garbling for quantum circuits [BY20].

Theorem 1.1 (Informal; Pseudo One-time Pad)**.** *Assuming the existence of (d, n)-PRFS with*[2] *$d = O(\log \lambda)$ and $n = \omega(\log \lambda)$, there exists a one-time encryption scheme for messages of length $\ell = 2^d$.*

We emphasize that in the implication to one-time encryption, we only require PRFS with logarithmic-length inputs.

The construction is simple and a direct adaptation of the construction of one-time encryption from pseudorandom generators. To encrypt a message x of length $\ell \gg \lambda$, output the state $G(k, (1, x_i)) \otimes \cdots \otimes G(k, (\ell, x_\ell))$, where $k \in \{0,1\}^\lambda$ is the symmetric key shared by the encryptor and the decryptor. The decryptor using the secret key k can decode[3] the message x. The security of the encryption scheme follows from the pseudorandomness of PRFS.

Implication 2. Statistically binding commitment schemes. We focus on designing commitment schemes with statistical binding and computational hiding properties. In the classical setting, this notion of commitment schemes can be constructed from any length-tripling PRG [Nao91]. Recently, two independent works [GLSV21, BCKM21b] showed that commitment schemes with aforementioned properties imply maliciously secure multiparty computation protocols with quantum communication in the dishonest majority setting. Of particular interest is the work of [BCKM21b] who show that the transformation is robust even if the underlying commitment scheme has quantum communication. They instantiate the underlying commitment scheme from one-way functions.

We design commitment schemes based on PRFS. We present a new definition of statistically binding commitments with quantum communication (see Definition 6.1 for the formal definition). Our definition generalizes all the current known definitions of statistically binding quantum commitments [YWLQ15, Unr16, FUYZ20, BCKM21b, BB21].

Theorem 1.2 (Informal)**.** *Assuming the existence of (d, n)-PRFS*[4] *where $2^{d(\lambda)} \cdot n(\lambda) \geq 7\lambda$, there exists a statistically binding and computationally hiding commitment scheme.*

[2] Recall that λ is the key length.

[3] In the technical sections, we define a QPT algorithm Test that given a state ρ along with k, x, determines if ρ is equal to the output $G(k, x)$. We show the existence of such a test algorithm for any PRFS.

[4] To simplify the analysis, there is an additional technical property of the PRFS not mentioned here that is required by our construction, called *recognizable abort* (Definition 3.5). All known constructions of PRFS and PRS including ours have the recognizable abort property.

By plugging our commitment scheme into the framework of [BCKM21b], we obtain the following corollary.

Corollary 1.3 (Informal)**.** *Assumuing the existence of (d, n)-PRFS with $2^d \cdot n \geq 7\lambda$, there exists a maliciously secure multiparty computation protocol in the dishonest majority setting.*

Our construction is an adaptation of Naor's commitment scheme [Nao91]. We replace the use of the PRG in Naor's construction with a PRFS generator and the first message (which is a random string in Naor's construction) specifies a random Pauli operator (also known as a *quantum one-time pad*).

Other Implications. Besides the above applications, we observe that PRFS can also be used to construct other fundamental primitives such as symmetric-key encryption and message authentication codes. Both primitives guarantee security in the setting when the secret key can be reused multiple times. We sketch both these applications and since they are simple, we omit their formal descriptions in the technical sections.

To design a symmetric-key encryption scheme from a PRFS, denoted by G, we start with the classical construction of symmetric-key encryption from a PRF: to encrypt a single bit m with respect to key k, output $(r, PRF(k, r) \oplus m)$, where r is a λ-bit string chosen uniformly at random and k is the symmetric key. We modify this construction by replacing the PRF with a PRFS. The result is a ciphertext of the form $(r, G(k, (r, m)))$. The decryptor decodes the message to be m by checking if the output is of the form $(r, G(k, (r, 0)))$ or $(r, G(k, (r, 1)))$.

The reusable message authentication code (MAC) from PRFS is even simpler: on input a message m, output the MAC signature $G(k, m)$, where k is the MAC secret key. To check if the MAC signature is valid, it suffices to check whether the signature is of the form $G(k, m)$.

Unlike the earlier implications, both (reusable) encryption and MACs require PRFS with input length to be as long as the message being encrypted/authenticated.

1.1.4 Construction of PRFS

Given the interesting implications of PRFS, the next natural step is to focus on constructing PRFS generators. We show that for some interesting range of parameters, we can achieve PRFS from any PRS.

In particular, we show the following.

Theorem 1.4 (Informal)**.** *For any $d = O(\log \lambda)$ and $n = d + \omega(\log \log \lambda)$, assuming the existence of a $(d + n)$-qubit PRS generator, there exists a (d, n)-PRFS generator.*

A surprising aspect about the above result is that the starting PRS's output length $d + n = \omega(\log \log \lambda)$ could even be much smaller than the key length λ. On the other hand, classical pseudorandom generators with output length less than the input length can be trivially constructed.

We remark that if $d \ll \log \lambda$ then it is easy to build PRFS from PRS; chop up the key k into 2^d blocks; to compute the PRFS generator with key k and input x, compute the PRS generator on the x^{th} block of the key. Unfortunately, PRFS with this range of parameters is not useful for applications. On the other hand, the above theorem allows for 2^d to be (an arbitrarily large) polynomial in the key length; specifically, these are same parameters stated in Theorem 1.1 and Corollary 1.3. We obtain the following corollary.

Corollary 1.5. *Assuming* $(2\log\lambda + \omega(\log\log\lambda))$*-qubit PRS, there exist statistically binding commitment schemes and therefore secure computations. Assuming* $\omega(\log\lambda)$*-PRS, there exist pseudo one-time pad schemes for polynomial (in* λ*) length messages.*

We complement this result by observing that a recent result by Brakerski and Shmueli [BS20] demonstrates the existence of statistically secure $O(\log\lambda)$-qubit PRS. Although it can be shown that statistical $\log(\lambda)$-qubit PRS cannot exist and $(1+\varepsilon)\log(\lambda)$-qubit PRS implies $BQP \neq PP$ [AQY21], when viewed optimistically, this suggests a potential way to construct secure computation unconditionally by further improving the output length of PRS if we only aim for computational security.

Main Insight. The construction of (d, n)-PRFS proceeds as follows: on input key k and $x \in \{0,1\}^d$, first generate a $(d+n)$-PRS state by treating k as the key. Denote the state to be $|\psi\rangle = \sum_{x\in\{0,1\}^d} \alpha_x |x\rangle \otimes |\psi_x\rangle$, where $|\psi_x\rangle$ is a n-qubit state. Suppose we can post-select on the first d qubits being $|x\rangle$ then we can denote $|\psi_x\rangle$ to be the output of PRFS on input (k, x).

There are two main points that need to be mentioned.

- Post-selection, in general, cannot be performed in polynomial time [Aar05]. However, if the event on which we are post-selecting has an inverse polynomial (where the polynomial is known ahead of time) probability of occurring, then we can efficiently perform post-selection.
- Now, we not only need to argue that $|\alpha_x|^2$ is roughly 2^{-d} (an inverse polynomial quantity), but we also need to argue that $|\psi_x\rangle$ is also pseudorandom. Since, a PRS state is indistinguishable from a Haar random state, it suffices to prove these two properties for Haar random states. Luckily, Haar random states have both the two nice properties we desire. Denote $|\vartheta\rangle = \sum_{x\in\{0,1\}^d} \beta_x |x\rangle \otimes |\vartheta_x\rangle$ to be a $(n+d)$-qubit Haar-random state. Then, the following holds: (i) with overwhelming probability, $|\beta_x|^2$ is close to 2^{-d} which in turn is inverse polynomial since d is logarithmic and, (ii) $|\vartheta_x\rangle$ is distributed according to the Haar measure.

Concurrent Work

The recent preprint of Morimae and Yamakawa [MY21] also construct statistically binding and computationally hiding commitment schemes from PRS, adapting Naor's commitment scheme in a manner similar to ours. We note several differences between their work and ours:

1. They show a weaker notion of binding known as *sum-binding*, which roughly says that the *sum* of the probabilities that an adversarial committer can successfully decommit to the bit 0 and the bit 1 is at most a quantity negligibly close to 1. This notion of binding is not known to be sufficient to conclude that PRS implies protocols for secure computation. However our notion of statistical binding (Definition 6.1) is sufficient for leveraging the machinery of [BCKM21b] to obtain quantum protocols for secure computation. Moreover, our definition of statistical binding implies their definition of sum-binding[5].

[5]The sum of probabilities that an adversarial decommitter can decommit to 0 and to 1 in the ideal world of our definition (Definition 6.1) and therefore they sum up to at most negligibly larger than 1 in the real world by our statistical binding guarantee.

2. For the same level of statistical binding security, that is $O(2^{-\lambda})$, they require the existence of a PRS that stretches λ random bits to 3λ qubits of Haar-randomness (i.e., they require the PRS generator to have *stretch*), whereas our result assumes the existence of a PRS that maps λ bits to $2\log\lambda + \omega(\log\log\lambda)$ qubits.

3. The state generation guarantee required from the underlying PRS is much stricter in their setting. In our work, we require the underlying PRS to only satisfy recognizable abort (Definition 3.5) whereas in their work, the underlying PRS needs to satisfy a guarantee that is even stronger than perfect state generation (Definition 3.4). In particular, they assume the existence of a unitary that outputs the state without producing any auxiliary.

Finally, the notion of PRFS, its implications and its construction from PRS is unique to our work.

2 Preliminaries

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 $\mathcal{H}$. 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.

General Measurements. A *general measurement* on a Hilbert space $\mathcal{H}$ is a set $M = \{M_a\}_{a\in A}$ of operators acting on $\mathcal{H}$ indexed by some finite set A of outcomes satisfying the completeness relation

$$\sum_{a\in A} M_a^\dagger M_a = I \ .$$

Applying the measurement M to a density matrix $\rho \in \mathcal{D}(\mathcal{H})$ corresponds to the following operation: outcome a is obtained with probability $\mathrm{Tr}(M_a^\dagger M_a \rho)$, and the post-measurement state is defined to

$$\rho \mapsto \frac{M_a \rho M_a^\dagger}{\mathrm{Tr}(M_a^\dagger M_a \rho)} \ .$$

The 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 2.1. *We have*

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

The following result, known as *Lévy's Lemma*, expresses strong concentration of measure for the Haar measure.

Fact 2.2 (Lévy's Lemma [HLW06])**.** *Let $f : \mathbb{C}^d \to \mathbb{R}$ be a function such that for all unit vectors $|\psi\rangle, |\phi\rangle$ we have*

$$|f(|\psi\rangle) - f(|\phi\rangle)| \leq K \cdot \| |\psi\rangle - |\phi\rangle \|$$

for some number $K > 0$. Then there exists a universal constant $C > 0$ such that

$$\Pr_{|\psi\rangle \leftarrow \mathscr{H}(\mathbb{C}^d)} [|f(|\psi\rangle) - \mathbb{E}\, f| \geq \delta] \leq \exp\Big(-\frac{Cd\delta^2}{K^2}\Big)$$

where $\mathbb{E}\, f$ denotes the average of f over the Haar distribution $\mathscr{H}(\mathbb{C}^d)$.

2.1 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 $p(\lambda)$. We furthermore 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 nonuniform quantum adversaries who may receive quantum states as advice.

The notation we use to describe the inputs/outputs of quantum algorithms will largely mimick 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 Pseudorandom States

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

Definition 3.1 (PRS Generator [JLS18])**.** *We say that a QPT algorithm G is a* pseudorandom state (PRS) generator *if the following holds.*

1. ***State Generation****. There is a negligible function $\varepsilon(\cdot)$ such that 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 $\varepsilon(\lambda)$ 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 3.2 ([JLS18])**.** *If post-quantum one-way functions exist, then there exist PRS generators for all polynomial output lengths.*

3.1 Pseudorandom Function-Like State (PRFS) Generators

In this section, we present our definition of pseudorandom function-like state (PRFS) generators. PRFS generators generalize PRS generators in two ways: first, in addition to the secret key k, the PRFS generator additionally takes in a (classical) input x. The security guarantee of a PRFS implies that, even after revealing x (but not the key k), the output state of the generator is indistinguishable from Haar-random. 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.

Definition 3.3 (PRFS generator)**.** *We say that a QPT algorithm G is a (selectively secure)* pseudorandom function-like state (PRFS) generator *if for all polynomials $s(\cdot), t(\cdot)$ and QPT (nonuniform) distinguishers A there exists a negligible function $\varepsilon(\cdot)$ such that for all λ, for all disjoint $x_1, \ldots, x_{s(\lambda)} \in \{0,1\}^{d(\lambda)}$ we have*

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

We also 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. One could consider stronger notions of security where the indistinguishability property holds even when the adversary is allowed to query the PRFS generator adaptively, or even in superposition. We leave exploring this notion for future work.

State Generation Guarantees. As mentioned above, our definition of PRFS generator does not require that the output of the generator is always a pure state. However, we will see later that a consequence of the PRFS security guarantee is that the output of the generator is *close* to a pure state for an overwhelming fraction of keys k (see Lemma 3.6).

Nevertheless, for applications it is sometimes more useful to also consider a stronger guarantee on the state generation of a PRFS generator.

Definition 3.4 (Perfect state generation)**.** *A $(d(\lambda), n(\lambda))$-PRFS generator G satisfies perfect state generation, if for every $k \in \{0,1\}^\lambda$ and $x \in \{0,1\}^{d(\lambda)}$, there exists an $n(\lambda)$-qubit pure state $|\psi\rangle$ such that $G_\lambda(k,x) = |\psi\rangle\langle\psi|$.*

We observe that an $n(\lambda)$-PRS generator defined in Definition 3.1 is by definition equivalent to an $(0, n(\lambda))$-PRFS generator with perfect state generation.

In general, it may be difficult to construct PR(F)S with perfect state generation as the state generation could occasionally fail; for example, the generator may perform a (quantum) rejection sampling procedure in order to output the state. The scalable PRS generators of Brakerski and Shmueli [BS20] is an example of this. To capture a very natural class of PRFS generators (including the one constructed in this paper), we define the notion of a PRFS generator where $G(k,x)$ outputs a convex combination of a fixed pure state $|\psi_{k,x}\rangle$ or a known abort state $|\perp\rangle$.

Definition 3.5 (Recognizable abort)**.** *A $(d(\lambda), n(\lambda))$-PRFS generator G has the* recognizable abort property *if for every $k \in \{0,1\}^\lambda$ and $x \in \{0,1\}^{d(\lambda)}$ there exists an $n(\lambda)$-qubit pure state $|\psi\rangle$ and $0 \leq \eta \leq 1$ such that $G_\lambda(k,x) = \eta\,|\psi\rangle\langle\psi| + (1-\eta)\,|\perp\rangle\langle\perp|$, where $\perp$ is a special symbol*[6]*.*

Note that this definition alone does not have any constraint on η being close to 1. However, the security guarantee of a PRFS generator implies that η will be negligibly close to 1 with overwhelming probability over the choice of k.[7]

We note that the PRS construction of Brakerski and Shmueli [BS20] also satisfies this recognizable abort guarantee instead of the perfect state generation guarantee. We also note that a PRFS generator with perfect state generation trivially has the recognizable abort property with $\eta = 1$ for all k, x.

3.2 Basic Properties of PRS and PRFS Generators

In this section we present some basic results about PRS and PRFS generators.

The following Lemma establishes some orthogonality and purity properties of the output of PRFS generators, on average over the key.

Lemma 3.6 (Properties of PRFS generator outputs)**.** *Let G be a (d,n)-PRFS generator. Then there exists a negligible function $\varepsilon(\lambda)$ such that for all λ, for all $x, y \in \{0,1\}^{d(\lambda)}$ where $x \neq y$, we have*

1. $\mathbb{E}_{k \leftarrow \{0,1\}^\lambda} \operatorname{Tr}(G_\lambda(k,x)\, G_\lambda(k,y)) \leq 2^{-n(\lambda)} + \varepsilon(\lambda)$;

[6] One can think of $|\perp\rangle$ as the $(n+1)$-qubit state $|100\cdots0\rangle$ with the first qubit indicating whether the generator aborted or not. If the generator doesn't abort, then it outputs $|0\rangle \otimes |\psi\rangle$ for some pure state $|\psi\rangle$ (called the *correct output state* of G on input (k,x)). The distinguisher in the definition of PRFS generator would then ignore the first indicator qubit.

[7] The argument is as follows: if η were on average noticeably far from 1, then a purity test using SWAP tests would distinguish the outputs from Haar random states which are pure. This is formalized in Lemma 3.6.

2. $\left|\mathbb{E}_{k\leftarrow\{0,1\}^\lambda} \operatorname{Tr}(G_\lambda(k,x)^2) - 1\right| \leq \varepsilon(\lambda)$.

Proof. Consider the following QPT algorithm A: on input $(x, y, |\phi_1\rangle, |\phi_2\rangle)$, it performs the SWAP test on $|\phi_1\rangle$ and $|\phi_2\rangle$ and accepts if the SWAP test accepts. If $|\phi_1\rangle, |\phi_2\rangle$ are independently sampled according to the Haar measure on $n(\lambda)$ qubits, the acceptance probability is on average

$$\frac{1}{2} + \frac{1}{2} \mathop{\mathbb{E}}_{|\phi_1\rangle,|\phi_2\rangle\leftarrow\mathscr{H}_n} |\langle\phi_1|\phi_2\rangle|^2 = \frac{1}{2} + \frac{1}{2} 2^{-n(\lambda)} \tag{1}$$

where we used Fact 2.1. On the other hand, if the algorithm A is run on input $(x, y, G_\lambda(k, x), G_\lambda(k, y))$ for randomly chosen k the acceptance probability is on average

$$\frac{1}{2} + \frac{1}{2} \mathop{\mathbb{E}}_{k\leftarrow\{0,1\}^\lambda} \operatorname{Tr}(G_\lambda(k,x)\, G_\lambda(k,y)) \ . \tag{2}$$

Since A is a QPT algorithm, by the pseudorandomness property of the PRFS generator, Equations (1) and (2) are negligibly different. Specifically, their difference is $\varepsilon(\lambda)$, where $\varepsilon(\lambda)$ is the negligible function guaranteed by the pseudorandomness property. This implies the first item of the Lemma.

For the second item of the Lemma, if $|\phi_1\rangle = |\phi_2\rangle$, then the algorithm A accepts $(x, x, |\phi_1\rangle, |\phi_1\rangle)$ with probability 1.

On the other hand, if the algorithm A is run on input $(x, x, G_\lambda(k, x), G_\lambda(k, x))$, then the acceptance probability is on average

$$\frac{1}{2} + \frac{1}{2} \mathop{\mathbb{E}}_{k\leftarrow\{0,1\}^\lambda} \operatorname{Tr}(G_\lambda(k,x)^2) \ .$$

Since the algorithm is efficient and only uses the output of the generator instead of the key, this implies that $\mathbb{E}_{k\leftarrow\{0,1\}^\lambda} \operatorname{Tr}(G_\lambda(k, x)^2)$ is negligibly (specifically, $\varepsilon(\lambda)$) different from 1, as desired. □

3.3 Testing Pseudorandom States

Given a state ρ, it is useful to know whether it is the output of a PRFS generator with key k and input x. One approach would be to invoke the generator to get some number of copies and perform SWAP tests. Unfortunately, this approach would only achieve polynomially small error, which is undesirable for cryptographic applications where we want negligible security. Another approach is to "uncompute" the state generation. The issue with this approach is that it is not clear how to do it when the state generation is not perfect, or if it outputs some additional auxiliary states that we do not know how to uncompute.

In the following, we will show how to use the generator in a semi-black-box way to test any PRFS states. We first state a general Lemma that shows how to convert any circuit that generates a state ρ into a tester (of sorts) for the state ρ.

Lemma 3.7 (Circuit output tester)**.** *Let G denote a (generalized) quantum circuit that takes no input and outputs an n-qubit mixed state ρ. Then there exists a circuit* Test *with boolean output such that:*

1. *For all density matrices σ_{EQ} where Q is an n-qubit register, applying the circuit Test on register Q yields the following state on registers EF where F stores the decision bit:*

$$(I_{\mathsf{E}} \otimes \mathsf{Test}_{\mathsf{Q}})(\sigma_{\mathsf{EQ}}) = \sum_b \mathrm{Tr}_{\mathsf{Q}}\Big((I_{\mathsf{E}} \otimes M_b)\sigma_{\mathsf{EQ}}\Big) \otimes |b\rangle\langle b|_{\mathsf{F}}$$

   where $M_1 = \rho^2$ and $M_0 = I - M_1$.

2. *Furthermore, Test runs the unitary part[8] of G as a black box, and if the complexity of G is T, the complexity of Test is $O(T + n)$.*

Proof. Consider the unitary part $\hat{G}$ of G, which takes as input a register A and outputs registers RB where R has n-qubits and A and B have the appropriate number of qubits.

The circuit Test takes as input an n-qubit register Q and outputs registers FQ where F is a single-qubit accept/reject register. It behaves as follows:

1. Initialize an ancilla register A in the state $|0\cdots0\rangle$, and initialize a single-qubit register F in the state $|0\rangle$.
2. Run the unitary part $\hat{G}$ on register A to obtain registers RB;
3. Swap the registers Q and R;
4. Apply the inverse $\hat{G}^\dagger$ on registers RB to get register A;
5. Measure the register A in the computational basis; if the outcome is $|0\cdots0\rangle$, then flip the qubit in F to $|1\rangle$.
6. Trace out the register Q.

This concludes the description of Test. Item 2 of the Lemma statement follows from inspection.

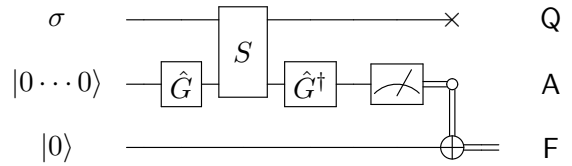

Figure 1: Test circuit. The S gate denotes SWAP between registers Q and R. The F register is set to $|1\rangle$ if and only if the A register measures to all zeroes. The Q and A registers are traced out at the end.

We now prove Item 1 of the Lemma statement. Fix a density matrix σ_{EQ}. Without loss of generality we can assume that σ is a pure state $|\theta\rangle_{\mathsf{EQ}}$ (because we can let E contain the purification). We analyze running the circuit Test on Q of $|\theta\rangle$.

Let $|\theta\rangle_{\mathsf{EQ}} = \sum_i \alpha_i |u_i\rangle_{\mathsf{E}} \otimes |v_i\rangle_{\mathsf{Q}}$ denote the Schmidt decomposition of $|\theta\rangle$ for some orthonormal bases $\{|u_i\rangle\}, \{|v_i\rangle\}$. After Step 2 of the Test circuit, the global state is

$$|\theta\rangle_{\mathsf{EQ}} \otimes \hat{G}\,|0\cdots0\rangle_{\mathsf{A}} \ .$$

[8] See Section 2.1 for a definition of the unitary part of a generalized quantum circuit.

Let $\hat{G}\left|0\right\rangle_{\mathsf{A}} = \sum_j \beta_j \left|\psi_j\right\rangle_{\mathsf{R}} \otimes \left|\phi_j\right\rangle_{\mathsf{B}}$ denote the Schmidt decomposition of $\hat{G}\left|0\cdots0\right\rangle$ for some orthonormal bases $\{\left|\psi_j\right\rangle\}, \{\left|\phi_j\right\rangle\}$. After Step 3 the global state can be written as

$$\sum_{ij} \alpha_i \beta_j \left|u_i\right\rangle_{\mathsf{E}} \otimes \left|\psi_j\right\rangle_{\mathsf{Q}} \otimes \left|v_i\right\rangle_{\mathsf{R}} \otimes \left|\phi_j\right\rangle_{\mathsf{B}} \ .$$

After Step 4 the global state can be written as

$$\sum_{ij} \alpha_i \beta_j \left|u_i\right\rangle_{\mathsf{E}} \otimes \left|\psi_j\right\rangle_{\mathsf{Q}} \otimes \hat{G}^\dagger\Big(\left|v_i\right\rangle_{\mathsf{R}} \otimes \left|\phi_j\right\rangle_{\mathsf{B}}\Big) \ .$$

In Step 5, the A register is measured. If the outcome is all zeroes, then the post-measurement state can be written as (up to normalization)

$$\begin{aligned}
&\sum_{ij} \alpha_i \beta_j \left|u_i\right\rangle_{\mathsf{E}} \otimes \left|\psi_j\right\rangle_{\mathsf{Q}} \otimes \left|0\right\rangle\!\left\langle 0\right|_{\mathsf{A}} \hat{G}^\dagger\Big(\left|v_i\right\rangle_{\mathsf{R}} \otimes \left|\phi_j\right\rangle_{\mathsf{B}}\Big) \\
&= \sum_{ij} \alpha_i \beta_j \left|u_i\right\rangle_{\mathsf{E}} \otimes \left|\psi_j\right\rangle_{\mathsf{Q}} \otimes \Big(\sum_k \beta_k \left\langle\psi_k\right|_{\mathsf{R}} \otimes \left\langle\phi_k\right|_{\mathsf{B}}\Big)\Big(\left|v_i\right\rangle_{\mathsf{R}} \otimes \left|\phi_j\right\rangle_{\mathsf{B}}\Big) \ \left|0\right\rangle_{\mathsf{A}} \\
&= \sum_{ij} \alpha_i \ \beta_j^2 \ \left\langle\psi_j|v_i\right\rangle \ \left|u_i\right\rangle_{\mathsf{E}} \otimes \left|\psi_j\right\rangle_{\mathsf{Q}} \otimes \left|0\right\rangle_{\mathsf{A}} \\
&= (I_{\mathsf{E}} \otimes \rho_{\mathsf{Q}}) \left|\theta\right\rangle_{\mathsf{EQ}} \otimes \left|0\right\rangle_{\mathsf{A}}
\end{aligned}$$

where in the second line we used our Schmidt decomposition for $\hat{G}\left|0\cdots0\right\rangle$ and in the third line we used the orthonormality of the basis $\{\left|\phi_j\right\rangle\}$. The fourth line follows since by definition of G we have $\rho = \sum_j \beta_j^2 \left|\psi_j\right\rangle\!\left\langle\psi_j\right|$.

If the measurement outcome is all zeroes, the register F is set to $\left|1\right\rangle$. Otherwise it remains $\left|0\right\rangle$.

In Step 6, the registers Q and A are traced out. Thus conditioned on getting the all zeroes outcome, the state on register E is

$$\mathrm{Tr}_{\mathsf{Q}}\Big(\rho_{\mathsf{Q}} \ \left|\theta\right\rangle\!\left\langle\theta\right|_{\mathsf{EQ}} \ \rho_{\mathsf{Q}}\Big) = \mathrm{Tr}_{\mathsf{Q}}\Big(\rho_{\mathsf{Q}}^2 \ \left|\theta\right\rangle\!\left\langle\theta\right|_{\mathsf{EQ}}\Big)$$

where we used the cyclicity of the partial trace with respect to operators acting on register Q only. Conditioned on *not* getting the all zeroes outcome, it must be that the state on register E is

$$\mathrm{Tr}_{\mathsf{Q}}\Big((I - \rho^2) \ \left|\theta\right\rangle\!\left\langle\theta\right|_{\mathsf{EQ}}\Big) \ .$$

This establishes that the output of the Test algorithm is as described in the Lemma statement. □

We note that if a PRFS satisfies perfect state generation, then the Test algorithm corresponding to the circuit $G_\lambda(k,x)$ implements a projection onto the state $\left|\psi_{k,x}\right\rangle = G_\lambda(k,x)$ in the case that the Test accepts (i.e. outputs 1). If the PRFS satisfies the weaker recognizable abort property, we get that the Test algorithm implements a *scaled* projection onto the correct state $\left|\psi_{k,x}\right\rangle$.

Corollary 3.8 (PRFS tester with recognizable abort)**.** *Let G be a (d,n)-PRFS generator with the recognizable abort property. Then there exists a QPT algorithm* Test *such that for all λ, $k \in$*

$\{0,1\}^\lambda$ *and* $x \in \{0,1\}^{d(\lambda)}$*, for all density matrices* σ_{EQ} *where* Q *is an* $n(\lambda)$*-qubit register, applying* $\mathsf{Test}(k,x,\cdot)$ *to register* Q *yields the following state on registers* EF *where* F *stores the decision bit:*

$$(I_{\mathsf{E}} \otimes \mathsf{Test}_{\mathsf{Q}})(k,x,\sigma_{\mathsf{EQ}}) = \sum_b \mathrm{Tr}_{\mathsf{Q}}\Big((I_{\mathsf{E}} \otimes M_b)\sigma_{\mathsf{EQ}}\Big) \otimes |b\rangle\langle b|_{\mathsf{F}}$$

where $M_1 = \eta^2 |\psi\rangle\langle\psi|$ *and* $M_0 = I - M_1$ *with* $\eta, |\psi\rangle$ *(which generally depend on* k,x*) are those guaranteed by the recognizable abort property.*

Proof. Fix λ and $k \in \{0,1\}^\lambda, x \in \{0,1\}^{d(\lambda)}$. By the recognizable abort property, we know that $G_\lambda(k,x) = \eta |\psi\rangle\langle\psi| + (1-\eta)|\bot\rangle\langle\bot|$. We implement the circuit Test by first testing whether the input state is $|\bot\rangle$ (which we can do since it is a fixed known state), rejecting if so, and otherwise applying the test circuit from Lemma 3.7 with the circuit $G_{k,x}$ that takes no input and outputs $\rho = G_\lambda(k,x)$. Since we projected the input state to have no overlap with $|\bot\rangle$, we get that

$$\rho\,\sigma\,\rho = \eta^2\ |\psi\rangle\langle\psi|\ \sigma\ |\psi\rangle\langle\psi|$$

as desired. □

Next we analyze a *product* of Test algorithms run in parallel on different qubits of a (possibly entangled) state.

Corollary 3.9 (Product of PRFS testers with recognizable abort)**.** *Let* G *be a* (d,n)*-PRFS generator with the recognizable abort property and let* Test *denote the corresponding tester algorithm given by Corollary 3.8. Fix* $\lambda, t \in \mathbb{N}$*. For all* $k_1,\ldots,k_t \in \{0,1\}^\lambda$ *and for all* $x_1,\ldots,x_t \in \{0,1\}^{d(\lambda)}$*, define the QPT algorithm* $\mathsf{Test}^{\otimes t}$ *that given an* $t \cdot n(\lambda)$*-qubit density matrix* σ *behaves as follows: for all* $i = 1,\ldots,t$*, on the* i*'th block of* $n(\lambda)$ *qubits of* σ*, run the algorithm* $\mathsf{Test}_\lambda(k_i,x_i,\cdot)$*. Output* 1 *if and only if all* t *invocations of* Test *output* 1*.*

Then $\mathsf{Test}^{\otimes t}$ *satisfies the following. For all density matrices* σ_{EQ} *where* Q *is an* $t \cdot n(\lambda)$*-qubit register, applying* $\mathsf{Test}^{\otimes t}$ *to register* Q *yields the following state on registers* EQF *where* F *stores the decision bit:*

$$(I_{\mathsf{E}} \otimes \mathsf{Test}^{\otimes t})(\sigma_{\mathsf{EQ}}) = \sum_b \mathrm{Tr}_{\mathsf{Q}}\Big((I_{\mathsf{E}} \otimes M_b)\sigma_{\mathsf{EQ}}\Big) \otimes |b\rangle\langle b|_{\mathsf{F}}$$

where $M_1 = \eta^2 |\psi\rangle\langle\psi|$ *and* $M_0 = I - M_1$ *with* $|\psi\rangle = |\psi_{k_1,x_1}\rangle \otimes \cdots \otimes |\psi_{k_t,x_t}\rangle$*, and* $\eta = \eta_{k_1,x_1} \cdots \eta_{k_t,x_t}$ *where* $|\psi_{k_i,x_i}\rangle, \eta_{k_i,x_i}$ *for* $i = 1,...,t$ *are the values guaranteed by the recognizable abort property.*

Proof. This follows from the fact that each invocation of $\mathsf{Test}(k_i,x_i,\cdot)$, conditioned on accepting, implements a (scaled) projection $\eta_{k_i,x_i} |\psi_{k_i,x_i}\rangle\langle\psi_{k_i,x_i}|$ on a disjoint register of σ. □

Recall that for a randomly chosen key, by pseudorandomness we know that η is negligibly close to 1 with overwhelming probability. Therefore, the corollaries show that we can estimate the overlap between any state and the correct state up to a negligible error with overwhelming probability when k is honestly sampled.

The next Lemma shows that Lemma 3.7 implies the ability to test the outputs of *any* PRFS generator (even ones without recognizable abort), *on average* over a uniformly random key k.

Lemma 3.10 (Self-testing PRFS). *Let G be a (d, n)-PRFS generator and $\mathsf{Test}(k, x, \cdot)$ denote the tester algorithm for channel $G(k, x)$ given by Lemma 3.7. There exists a negligible function $\nu(\cdot)$ such that for all λ, for all $x \neq y$,*

$$\Pr_k[\mathsf{Test}(k, x, G(k, x)) = 1] \geq 1 - \nu(\lambda),$$

and

$$\Pr_k[\mathsf{Test}(k, x, G(k, y)) = 1] \leq 2^{-n(\lambda)} + \nu(\lambda).$$

Proof. By Lemma 3.7,

$$\Pr_k[\mathsf{Test}(k, x, G(k, x)) = 1] = \mathop{\mathbb{E}}_k \left[\mathrm{Tr}(G(k, x)^3)\right] \geq \mathop{\mathbb{E}}_k \left[\frac{\mathrm{Tr}(G(k, x)^2)^2}{\mathrm{Tr}(G(k, x))}\right] = \mathop{\mathbb{E}}_k \left[\mathrm{Tr}(G(k, x)^2)^2\right],$$

which is negligibly close to 1 by Item 1 of Lemma 3.6 and Markov's inequality, and the inequality is due to the following fact. For any finite-dimensional vector x with non-negative coefficients, by Cauchy–Schwarz we have

$$\|x\|_1 \|x\|_3^3 = \left(\sum_i x_i\right)\left(\sum_i x_i^3\right) \geq \left(\sum_i \sqrt{x_i}\sqrt{x_i^3}\right)^2 = \|x\|_2^4.$$

Similarly,

$$\Pr_k[\mathsf{Test}(k, x, G(k, y)) = 1] = \mathop{\mathbb{E}}_k \left[\mathrm{Tr}(G(k, x)^2 G(k, y))\right] \leq \mathop{\mathbb{E}}_k \left[\mathrm{Tr}(G(k, x) G(k, y))\right],$$

which is at most negligibly larger than $2^{-n(\lambda)}$ by Item 2 of Lemma 3.6, and the inequality is due to the fact that $0 \preccurlyeq G(k, x) \preccurlyeq I$. □

4 Constructing PRFS from PRS

In this section we present our construction of PRFS generators using PRS generators, which are seemingly weaker objects. As mentioned in the introduction, there is a trivial construction of PRFS from PRS. Let G be a PRS generator. Define the PRFS generator G' with input length $d(\lambda) = O(\log \lambda)$, where $G'_{\lambda'}(k, x) = G_\lambda(k_x)$ with $\lambda' = 2^{d(\lambda)}\lambda$ and k_x denoting the x'block of λ bits in $k \in \{0, 1\}^{\lambda'}$.

However, this simple construction is such that the input length is always at most logarithmic in the seed length. This, as far as we can tell, is not very useful for applications.

The construction we present in this section is less trivial: we can build a PRFS generator with input length $d(\lambda)$ that is *any* constant times $\log \lambda$, as long as the the output length of the starting PRS generator is at least $2d(\lambda) + \omega(\log \log \lambda)$. Although the input length may appear modest, such PRFS generators are sufficient for most of the applications we consider in this paper. We find it an intriguing question of whether it is possible to construct PRFS generators with longer input lengths from PRS generators in a black box way, without increasing the seed length by too much.

Theorem 4.1. *Let $d(\lambda), n(\lambda)$ be functions such that $d(\lambda) = O(\log \lambda)$ and $n(\lambda) = d(\lambda) + \omega(\log \log \lambda)$. Let G denote a $(n(\lambda) + d(\lambda))$-PRS generator. Then there exists a $(d(\lambda), n(\lambda))$-PRFS generator F with the recognizable abort property, such that for all λ the circuit F_λ invokes the G_λ as a black box.*

For notational clarity we use the abbreviations $d = d(\lambda)$ and $n = n(\lambda)$.

Define the following circuit $F_\lambda(k, x)$: On input key $k \in \{0,1\}^\lambda$, input $x \in \{0,1\}^d$, repeat the following $2^d \cdot \lambda$ times:

- Compute the $(d+n)$-qubit state $\rho_k \leftarrow G_\lambda(k)$.
- Measure the first d qubits of ρ_k in the computational basis to obtain a string $y \in \{0,1\}^d$. If $y = x$, then output the remaining n qubits. Otherwise, continue.

If the measurement outcomes was different from x in all the $2^d\lambda$ iterations, set $\sigma_{k,x} = |\bot\rangle\langle\bot|$. Let the output be $\sigma_{k,x}$.

The algorithm $F = \{F_\lambda\}_\lambda$ is uniform QPT because for each λ, the running time of the circuit F_λ is going to be $O(2^d \cdot \lambda)$ times the complexity of running G_λ, which is QPT since $d = O(\log \lambda)$ and G is QPT. It is easy to see that if G satisfies recognizable abort, then F also satisfies recognizable abort by construction.

We now argue that the outputs of F satisfy the pseudorandomness property of a PRFS. Let A be a non-uniform QPT algorithm such that there exists $x_1 < \ldots < x_s \in \{0,1\}^d$ such that the following holds:

$$\Big| \Pr_{k\leftarrow\{0,1\}^\lambda}\left[A_\lambda(x_1,\ldots,x_s,F_\lambda(k,x_1)^{\otimes t},\ldots,F_\lambda(k,x_s)^{\otimes t})=1\right] - \Pr_{|\vartheta_1\rangle,\ldots,|\vartheta_s\rangle\leftarrow\mathscr{H}_n}\left[A_\lambda(x_1,\ldots,x_s,|\vartheta_1\rangle^{\otimes t},\ldots,|\vartheta_s\rangle^{\otimes t})=1\right]\Big| = \varepsilon(\lambda)\ . \tag{3}$$

Assume for contradiction that $\varepsilon(\lambda)$ is not negligible in λ.

Let $M = \lambda s t 2^d = \lambda^{O(1)}$. We construct a QPT algorithm B_λ that takes as input $\rho^{\otimes M}$ where ρ is a $(d+n)$-qubit state and does the following:

- For $j = 1, \ldots, s$, repeat the following t times:
  - Repeat the following $\lambda 2^d$ times: Measure the first d qubits of a new copy of ρ in the computational basis to obtain a string $y \in \{0,1\}^d$. If $y = x_j$, then save the remaining n qubits of ρ (which we denote as the state σ_{x_j}). Otherwise, continue.
  - If the outcome x_j was never measured, B_λ aborts.
- Execute $b \leftarrow A_\lambda\left(x_1,\ldots,x_s,\sigma_{x_1}^{\otimes t},\ldots,\sigma_{x_s}^{\otimes t}\right)$.
- Output b.

We show the following:

$$\Big| \Pr_{k\leftarrow\{0,1\}^\lambda}\left[B_\lambda(G_\lambda(k)^{\otimes M})=1\right] - \Pr_{|\vartheta\rangle\leftarrow\mathscr{H}_{d+n}}\left[B_\lambda(|\vartheta\rangle^{\otimes M})=1\right]\Big| \geq \varepsilon(\lambda) - \nu(\lambda) \tag{4}$$

for some negligible function $\nu(\lambda)$. This in turn shows that the algorithm $B = \{B_\lambda\}_\lambda$ violates the pseudorandomness assumption on the PRS generator G, which is a contradiction. Thus $\varepsilon(\lambda)$ must be negligible.

We prove (4) as follows. Consider an intermediate QPT algorithm $\widetilde{B}_\lambda$ that takes as input $(|\vartheta_1\rangle^{\otimes t},\ldots,|\vartheta_s\rangle^{\otimes t})$, where $|\vartheta_i\rangle \leftarrow \mathscr{H}_n$ and output the outcome of $A_\lambda\left(x_1,\ldots,x_s,|\vartheta_1\rangle^{\otimes t},\ldots,|\vartheta_s\rangle^{\otimes t}\right)$. Consider the following two claims.

Lemma 4.2.

$$\left| \Pr_{k \leftarrow \{0,1\}^\lambda} \left[B_\lambda(G_\lambda(k)^{\otimes M}) = 1 \right] - \Pr_{|\vartheta_i\rangle \leftarrow \mathscr{H}_n} \left[\widetilde{B}_\lambda \left(|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t} \right) = 1 \right] \right| = \varepsilon(\lambda) \ .$$

Proof. This follows from (3): (i) the output distribution, over the randomness of k, of output of $B_\lambda \left(G_\lambda(k)^{\otimes M} \right)$ is precisely the output distribution, over the of $A_\lambda(x_1, \ldots, x_s, F_\lambda(k, x_1)^{\otimes t}, \ldots, F_\lambda(k, x_s)^{\otimes t})$ and, (ii) the output distribution of $\widetilde{B}_\lambda \left(|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t} \right)$ is precisely the output distribution of $A_\lambda(x_1, \ldots, x_s, |\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t})$, where $|\vartheta_i\rangle$ is a n-qubit Haar random state. □

Lemma 4.3. *There exists a negligible function $\nu(\lambda)$ such that*

$$\left| \Pr_{|\vartheta\rangle \leftarrow \mathscr{H}_{d+n}} \left[B_\lambda(|\vartheta\rangle^{\otimes M}) = 1 \right] - \Pr_{|\vartheta_i\rangle \leftarrow \mathscr{H}_n} \left[\widetilde{B}_\lambda \left(|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t} \right) = 1 \right] \right| \leq \nu(\lambda) \ .$$

Proof. Consider the behavior of the algorithm B_λ on input $|\vartheta\rangle^{\otimes M}$ for $|\vartheta\rangle$ sampled from the Haar distribution $\mathscr{H}_{d+n}$. Define the distribution $\mathscr{R}$ over $(d+n)$-qubit unitary operators

$$R = \sum_{x \in \{0,1\}^d} |x\rangle\langle x| \otimes R_x$$

where $(R_x)_{x \in \{0,1\}^d}$ is a sequence of i.i.d. Haar-random n-qubit unitaries.

Observe that, by the unitary invariance of the Haar measure, $R|\vartheta\rangle$ is also distributed according to $\mathscr{H}_{d+n}$. Therefore the algorithm B_λ behaves identically on input $(R|\vartheta\rangle)^{\otimes M}$.

Define the event $\mathsf{MeasureFail}$ to be the event that the algorithm B_λ aborts on input $(R|\vartheta\rangle)^{\otimes M}$; this happens only if there exists a $j \in [s]$ such that, even after measuring the first d qubits of $\lambda t 2^d$ copies of $R|\vartheta\rangle$, the string x_j occured fewer than t times as a measurement outcome.

Notice that the event $\mathsf{MeasureFail}$ (and its negation) is *independent* of the choice of randomizing unitaries $(R_x)_x$; that is because applying R to $|\vartheta\rangle$ does not change the distribution of measurement outcomes on the first d qubits. Thus, for all $|\vartheta\rangle = \sum_x \alpha_x |x\rangle \otimes |\vartheta_x\rangle$, conditioning on the event $\neg\mathsf{MeasureFail}$ (the negation of $\mathsf{MeasureFail}$) still leaves the unitary R distributed according to $\mathscr{R}$.

Therefore for all $|\vartheta\rangle$ we have

$$\Pr \left[B_\lambda((R|\vartheta\rangle)^{\otimes M}) = 1 \middle| \neg\mathsf{MeasureFail} \right] = \Pr \left[\widetilde{B}_\lambda((R_{x_1}|\vartheta_{x_1}\rangle)^{\otimes t}, \ldots, (R_{x_s}|\vartheta_{x_s}\rangle)^{\otimes t}) = 1 \right] \quad (5)$$

where the probabilities are over the randomness of the measurements and also the randomness of sampling $R \leftarrow \mathscr{R}$. Since the R_{x_j}'s are i.i.d. Haar-random unitaries and the x_i's are distinct, this is equal to

$$\Pr \left[\widetilde{B}_\lambda(|\phi_1\rangle^{\otimes t}, \ldots, |\phi_s\rangle^{\otimes t}) = 1 \right]$$

where the probability is over the measurement outcomes and $|\phi_i\rangle$'s are i.i.d. Haar-random n-qubit

states. Thus

$$\begin{aligned}
&\Pr\left[B_\lambda(|\vartheta\rangle^{\otimes M}) = 1\right] \\
&= \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1\right] \\
&= \Pr\left[\neg\mathsf{MeasureFail}\right] \cdot \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \neg\mathsf{MeasureFail}\right] \\
&\qquad + \Pr\left[\mathsf{MeasureFail}\right] \cdot \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \mathsf{MeasureFail}\right] \\
&= \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \neg\mathsf{MeasureFail}\right] \\
&\qquad + \Pr\left[\mathsf{MeasureFail}\right] \cdot \Big(\Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \mathsf{MeasureFail}\right] \\
&\qquad\qquad - \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \neg\mathsf{MeasureFail}\right] \Big)
\end{aligned}$$

where we used $\Pr\left[\neg\mathsf{MeasureFail}\right] = 1 - \Pr\left[\mathsf{MeasureFail}\right]$. Thus

$$\begin{aligned}
&\left|\Pr\left[B_\lambda(|\vartheta\rangle^{\otimes M}) = 1\right] - \Pr\left[\widetilde{B}_\lambda\left(|\vartheta_1\rangle^{\otimes t}, \ldots, |\vartheta_s\rangle^{\otimes t}\right) = 1\right]\right| \\
&\leq \Pr\left[\mathsf{MeasureFail}\right] \cdot \left| \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \mathsf{MeasureFail}\right] - \Pr\left[B_\lambda((R\,|\vartheta\rangle)^{\otimes M}) = 1 \middle| \neg\mathsf{MeasureFail}\right] \right| \\
&\leq \Pr\left[\mathsf{MeasureFail}\right] \ .
\end{aligned}$$

We now estimate the probability of the event $\mathsf{MeasureFail}$. Fix a state $|\vartheta\rangle$ and let p_{x_j} denote the probability of obtaining x_j when measuring the first d qubits of $|\vartheta\rangle$, or equivalently since R commutes with the measurement, $R\,|\vartheta\rangle$. Fix a $j \in [s]$. The probability that measuring $\lambda 2^d$ copies of $R\,|\vartheta\rangle$ fails to yield the outcome x_j is equal to

$$\left(1 - p_{x_j}\right)^{\lambda 2^d}$$

The algorithm aborts if this happens in any of the st iterations of the "main loop" of B_λ; thus the probability of $\mathsf{MeasureFail}$ is, by union bound, at most

$$\sum_{j=1}^{s} t\left(1 - p_{x_j}\right)^{\lambda 2^d}$$

The following Lemma establishes deviation bounds on the probabilities p_x:

Lemma 4.4. *Let $|\psi\rangle$ be sampled from the Haar distribution $\mathscr{H}_{d+n}$. For all $x \in \{0,1\}^d$, let p_x denote the probability of measuring the first d qubits of $|\psi\rangle$ in the computational basis and obtaining outcome x. Then for all $\delta > 0$ with probability at least $1 - 2^{d+1} \cdot \exp(-C2^{n+d}\delta^2)$ over $|\psi\rangle$ for some universal constant $C > 0$, we have that*

$$|p_x - 2^{-d}| \leq \delta$$

for all $x \in \{0,1\}^d$.

By Lemma 4.4 (setting $\delta = 2^{-d}/2$), with all but negligible probability over the choice of $|\vartheta\rangle$, each of the p_{x_j}'s are at least $2^{-d}/2$ and therefore the probability of $\mathsf{MeasureFail}$, when averaged over the choice of $|\vartheta\rangle$, is at most

$$st(1 - 2^{-d}/2)^{\lambda 2^d} + 2\exp(-(C2^{n-d} - d)) \leq st\exp(-\Omega(\lambda)) + 2\exp(-(C2^{n-d} - d))$$

which for our choice of s, t, n, d is negligible in λ. □

Applying triangle inequality to Lemmas 4.2 and 4.3, we establish Equation (4), as desired. We conclude this section by proving Lemma 4.4.

Proof of Lemma 4.4. We first show that, with high probability over $|\psi\rangle$, the probability obtaining any *fixed* prefix $x \in \{0,1\}^d$ is going to be exponentially small in 2^n. We then apply a union bound over all $x \in \{0,1\}^d$ to obtain the Lemma statement.

Let Π_x denote the projector onto the first d qubits being in the state $|x\rangle$. Define $p_x = \mathrm{Tr}(\Pi\,|\psi\rangle\langle\psi|)$. On average over the choice of $|\psi\rangle$, this quantity is equal to

$$\mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\mathscr{H}_{d+n}} p_x = \mathrm{Tr}\left(\Pi_x \mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\mathscr{H}_{d+n}} |\psi\rangle\langle\psi|\right) = 2^{-(d+n)}\,\mathrm{Tr}(\Pi_x) = 2^{-d}$$

where we used the fact that the average of a Haar-random state is the maximally mixed state.

We now appeal to *Lévy's Lemma* (Fact 2.2), which shows that p_x concentrates tightly around its expectation. Define $f(|\psi\rangle) = \mathrm{Tr}(\Pi_x\,|\psi\rangle\langle\psi|)$. One can calculate that the Lipschitz constant of f is at most 2;

$$\frac{|f(|\psi\rangle) - f(|\phi\rangle)|}{\|\,|\psi\rangle - |\phi\rangle\,\|} \leq 2$$

for all $|\psi\rangle, |\phi\rangle$. By Fact 2.2, we have

$$\Pr\left[\left|p_x - 2^{-d}\right| \geq \delta\right] \leq 2\exp\left(-C2^{d+n}\delta^2\right)$$

for some universal constant $C > 0$, where the probability is over $|\psi\rangle \leftarrow \mathscr{H}_{d+n}$. □

5 Quantum Pseudo One-Time Pad from PRFS

The first application of PRFS we present is the Quantum Pseudo One-Time Pad (QP-OTP). In classical cryptography, a pseudo one-time pad is like the one-time pad except the key length is shorter than the length of the plaintext message. This is often presented in introductory cryptography courses as a basic example of using pseudorandomness to achieve a cryptographic task that is impossible in the information-theoretic setting. Here, we use a PRFS in place of a PRG to encrypt (classical) messages.

We point out that without knowing about the notion of PRFS, it appears difficult and challenging to construct secure quantum one-time pad schemes directly from PRS generators alone.

Definition 5.1 (Quantum Pseudo One-Time Pad)**.** *We say that a pair of QPT algorithms* $(\mathsf{Enc}, \mathsf{Dec})$ *is a* quantum pseudo one-time pad (QP-OTP) *if the following properties are satisfied: there exist a polynomial* $\ell(\lambda)$ *such that*

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

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

- ***Security***: *For every polynomial $t(\lambda)$, for every nonuniform QPT adversary A, there exists a negligible function $\varepsilon(\cdot)$ where for every λ and $x \in \{0,1\}^\ell$,*

$$\left| \Pr_{\substack{k \leftarrow \{0,1\}^\lambda, \\ \sigma \leftarrow \mathsf{Enc}_\lambda(k,x)}} \left[A_\lambda(\sigma^{\otimes t}) = 1\right] - \Pr_{|\vartheta_1\rangle,\ldots,|\vartheta_\ell\rangle \leftarrow \mathscr{H}_n} \left[A_\lambda((|\vartheta_1\rangle \otimes \cdots \otimes |\vartheta_\ell\rangle)^{\otimes t}) = 1\right] \right| \leq \varepsilon(\lambda),$$

where we have abbreviated $n = n(\lambda)$, $\ell = \ell(\lambda)$, and $t = t(\lambda)$.

Here we are saying the security holds even if the adversary could see multiple copies of the same ciphertexts, which might be useful for certain applications. However, when $t = 1$, we can see that the security implies that the ciphertext is computationally indistinguishable to random bit strings of length ℓn (or a maximally mixed state) by Fact 2.1.

To construct such a quantum pseudo one-time pad, let G be a $(d(\lambda), n(\lambda))$-PRFS generator where $d(\lambda) \geq \lceil \log \ell(\lambda) \rceil + 1$ and $n(\lambda) = \omega(\log \lambda)$. We interpret $G_\lambda(k, \cdot)$ as taking inputs of the form (i, b) where $i \in [\ell(\lambda)]$ and $b \in \{0,1\}$. Let Test denote the test algorithm from Lemma 3.10.

Fix λ and let $\ell = \ell(\lambda)$, $d = d(\lambda)$, and $n = n(\lambda)$.

1. $\mathsf{Enc}_\lambda(k, x)$: on input $k \in \{0,1\}^\lambda$ and a message $x \in \{0,1\}^\ell$, do the following:
   - For every $i \in [\ell]$, compute $\sigma_i \leftarrow G_\lambda(k, (i, x_i))$.
   - Set $\sigma = \sigma_1 \otimes \cdots \otimes \sigma_\ell$.

   Output the ciphertext state σ.

2. $\mathsf{Dec}_\lambda(k, \sigma')$: on input k, ℓn-qubit ciphertext state σ', perform the following operations:
   - Parse σ' as $\sigma'_1 \otimes \cdots \otimes \sigma'_\ell$.
   - For $i \in [\ell]$, execute $\mathsf{Test}(k, (i, 0), \sigma_i)$. If it accepts, set $x_i = 0$. Otherwise, set $x_i = 1$.
   - Output $x = x_1 \cdots x_\ell$.

Lemma 5.2. $(\mathsf{Enc}, \mathsf{Dec})$ *satisfies the correctness property of a quantum pseudo one-time pad according to Definition 5.1.*

Proof. Fix λ and let $\ell = \ell(\lambda)$. Fix a message $x \in \{0,1\}^\ell$. Let $\sigma_{k,i} = G_\lambda(k, (i, x_i))$ and let $\sigma_k = \sigma_{k,1} \otimes \cdots \otimes \sigma_{k,\ell}$.

Consider the decryption process. Fix an index $i \in [\ell]$. By Lemma 3.10, the probability that $\mathsf{Test}\Big(k, (i, 0), \sigma_{k,i}\Big)$ accepts (on average over k) is negligibly close to 1 if $x_i = 0$, and it is negligibly close to 0 if $x_i = 1$, on average over the key k (here we use the fact that the output length of the PRFS generator is $\omega(\log \lambda)$, so that $2^{-n(\lambda)}$ is negligible). Thus the probability that the correct bit x_i gets decoded is negligibly close to 1. Taking a union bound over all indices i, we get that the probability of decoding x is negligibly close to 1, over the randomness of the key k and the decryption algorithm. □

Lemma 5.3. $(\mathsf{Enc}, \mathsf{Dec})$ *satisfies the security property of quantum pseudo one-time pad according to Definition 5.1.*

Proof. We prove the security via a hybrid argument. Let $n(\lambda)$ denote the output length of the PRFS generator G. Fix λ, and let $\ell = \ell(\lambda)$, $n = n(\lambda)$, and $t = t(\lambda)$. Fix a message $x \in \{0,1\}^\ell$. Consider a nonuniform QPT adversary A such that A_λ takes as input t copies of an ℓn-qubit density matrix σ.

Hybrid H_1. Sample $k \leftarrow \{0,1\}^\lambda$. Compute $\sigma \leftarrow \mathsf{Enc}_\lambda(k,x)$. The output of the hybrid is the output of the adversary A_λ on input $\sigma^{\otimes t}$.

Hybrid H_2. Consider the following QPT algorithm B_λ: it takes as input $(i_1,b_1),\ldots,(i_\ell,b_\ell) \in [\ell]\times\{0,1\}$ and a tn-qubit state $\sigma_1^{\otimes t}\otimes\cdots\otimes\sigma_\ell^{\otimes t}$. It runs the adversary A_λ on input $(\sigma_1\otimes\cdots\otimes\sigma_\ell)^{\otimes t}$.

Sample $k \leftarrow \{0,1\}^\lambda$. Compute t copies of $\sigma \leftarrow \mathsf{Enc}_\lambda(k,x)$. The output of this hybrid is the output of B_λ on input $((1,x_1),\ldots,(\ell,x_\ell))$ and $\sigma^{\otimes t} = \sigma_1^{\otimes t}\otimes\cdots\otimes\sigma_\ell^{\otimes t}$.

Hybrid H_3. Sample t copies of Haar-random states $|\vartheta_1\rangle,\ldots,|\vartheta_\ell\rangle \leftarrow \mathscr{H}_n$. The output of this hybrid is the output of B_λ on input $((1,x_1),\ldots,(\ell,x_\ell))$ and $|\vartheta_1\rangle^{\otimes t}\otimes\cdots\otimes|\vartheta_\ell\rangle^{\otimes t}$.

We now argue the indistinguishability of the hybrids. Clearly, hybrids H_1 and H_2 are identical by construction (the adversary B_λ ignores its first input and runs A_λ on input $\sigma^{\otimes t}$). Hybrids H_2 and H_3 are indistinguishable because of the pseudorandomness property of the PRFS generator G. Notice that, by construction, the output of hybrid H_3 is $A_\lambda((|\vartheta_1\rangle\otimes\cdots\otimes|\vartheta_\ell\rangle)^{\otimes t})$. □

6 Quantum Bit Commitments from PRFS

6.1 Definition

We consider the notion of quantum commitment scheme with statistical binding and computational hiding property. This is analogous to a classical commitment scheme where the messages are allowed to be quantum states. We in particular focus on bit commitments where the the committed message is a single bit. We can achieve commitments of long messages generically by composing many instantiations of the bit-commitment scheme in parallel.

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, 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 define the properties satisfied by a commmitment scheme.

Statistical Binding. We start by discussing the statistical binding property. A natural quantum analogue of classical statistical binding property would be to consider the following: for any adversarial (possibly unbounded) committer C_λ^*, we require that at the end of the commit phase, over the randomness of the receiver, there is a unique message that C_λ^* can decommit to in the reveal phase. Unfortunately, this idealistic notion is impossible to achieve: for example, C_λ^* can send a (uniform) superposition of commitments of 0 and 1 and later can open to either 0 or 1 with equal probability. This attack was observed and taken into account in many works, including but not limited to [YWLQ15, Unr16, FUYZ20, BB21].

To account for this issue, we consider a notion where an extractor algorithm needs to be applied on the state of the receiver after the commit phase. The output is the modified receiver's state along with a bit b. We revise the statistical binding property guarantee to be the following: the probability that the extracted bit b is different from the bit decommitted by C^*_λ and the verifier rejects is negligible in λ. We present the definition below.

Definition 6.1 (Statistical Binding)**.** *We say that a quantum commitment scheme (C, R) satisfies* statistical binding *if for any (non-uniform) adversarial committer $C^* = \{C^*_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a (possibly inefficient) extractor algorithm $\mathcal{E}$ such that the following holds:*

$$\mathrm{TD}\left(\mathsf{RealExpt}^{C^*}_\lambda, \mathsf{IdealExpt}^{C^*,\mathcal{E}}_\lambda\right) \leq \nu(\lambda),$$

for some negligible function $\nu(\lambda)$, where the experiments $\mathsf{RealExpt}^{C^*}_\lambda$ and $\mathsf{IdealExpt}^{C^*,\mathcal{E}}_\lambda$ are defined as follows.

- $\mathsf{RealExpt}^{C^*}_\lambda$*: Execute the commit phase to obtain the joint state $\sigma_{C^*R} \leftarrow \mathsf{Commit}\langle C^*_\lambda, R_\lambda\rangle$. Execute the reveal phase to obtain the trit $\mu \leftarrow \mathsf{Reveal}\langle C^*_\lambda, R_\lambda, \sigma_{C^*R}\rangle$. Output the pair (τ_{C^*}, μ) where τ_{C^*} is the final state of the committer.*
- $\mathsf{IdealExpt}^{C^*,\mathcal{E}}_\lambda$*: Execute the commit phase to obtain the joint state $\sigma_{C^*R} \leftarrow \mathsf{Commit}\langle C^*_\lambda, R_\lambda\rangle$. Apply the extractor $\mathcal{E}$ on the receiver's part of σ_{C^*R} to obtain a new joint committer-receiver state σ'_{C^*R} along with $b' \in \{0, 1, \perp\}$. Execute the reveal phase to obtain the trit $\mu \leftarrow \mathsf{Reveal}\langle C^*_\lambda, R_\lambda, \sigma'_{C^*R}\rangle$. Let τ_{C^*} denote the final state of the committer. If $\mu = b'$, then output (τ_{C^*}, μ). Otherwise output $(\tau_{C^*}, \perp)$.*

Remark 6.2. *Many prior works consider statistical binding for quantum commitments. We highlight the main differences between our definition and the prior notions.*

- *Comparison with [YWLQ15, Unr16, FUYZ20]: the statistical binding property is formalized by requiring the states of the (honest) committer when committing to bits 0 and 1 to be far in trace distance. While their definition is cleaner (and probably equivalent to our notion), in our opinion, it is unwieldy to use their definition for applications. Specifically, one has to either implicitly or explicitly come up with an extractor in the security proofs [YWLQ15, FUYZ20] and moreover, show that the indistinguishability of the real and the ideal world holds against dishonest committers. On the other hand, we incorporate these technical difficulties as requirements in our definition making it easier to use in applications.*

  Another downside of the statistical binding property there is that in order for the binding property to hold, the opening phase must follow the "canonical" opening protocol, where the committer sends the purification of the mixed state sent in the committing phase, and the receiver performs a rank-1 projection to check the state. This implies that both *parties must keep their part of the state coherent between the two phases. However, our definition gives the flexibility of the reveal phase having purely classical communication.*

- *Comparison with [BB21]: A related work by [BB21] considers statistical binding of quantum commitments with the added feature that their opening is classical. The main difference is the following. In their notion, the honest receiver applies an operation that collapses the commitment into a quantum state and a classical string in such a way that the classical string*

information theoretically determines the message. They then use this feature to show that in some applications, the opening of the commitment can be classical. Our definition is also more general in the sense that the honest receiver does not apply any such operation and the collapsing part happens implicitly in the execution of extractor.

Remark 6.3. *One has to be careful when using quantum commitments in a larger system if the receiver's state is quantum after the commit phase. As an example, suppose we design a protocol where the quantum commitment sent by the receiver is used inside another cryptographic protocol then we might not be able to invoke binding if the state is destroyed, whereas classically the state could always be copied. Nevertheless, this is a generic caveat of quantum commitments and is not an artifact of any specific definition of binding.*

Computational Hiding. We define the computational hiding property below. This is the natural quantum analogue of the classical computational hiding property. In the literature, this property is also sometimes referred to as quantum concealing.

Definition 6.4 (Computational Hiding)**.** *We say that a quantum commitment scheme (C, R) satisfies computational 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:*

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

for some negligible function $\nu(\cdot)$, where σ_{R^*} is obtained by tracing out the committer's part of the state σ_{CR^*}.

6.2 Construction

We now present the main theorem of this section, which shows that statistically binding quantum commitment schemes can be constructed from PRFS.

Theorem 6.5. *Assuming the existence of $(d(\lambda), n(\lambda))$-PRFS satisfying recognizable abort (Definition 3.5) and $2^d \cdot n \geq 7\lambda$, there exists a commitment scheme satisfying statistical completeness, statistical binding (Definition 6.1) and computational hiding (Definition 6.4).*

We note that, combined with Theorem 4.1 which constructs PRFS generators with $\Omega(\log\lambda)$ input length and recognizable abort property from PRS generators, we can obtain quantum commitment schemes from PRS generators. We present the construction, which is inspired by Naor's commitment scheme [Nao91].

Construction. The main building block is a $(d(\lambda), n(\lambda))$-PRFS, denoted by $G = \{G_\lambda(\cdot,\cdot)\}_{\lambda\in\mathbb{N}}$. Since $n \geq 1$, we assume $d(\lambda) = \lceil \log \frac{7\lambda}{n} \rceil = O(\log\lambda)$ to ensure the efficiency of the algorithm. This is without loss of generality since we can generically shrink the input length for a PRFS by padding zeroes. Let $\mathsf{Test}_\lambda^{\otimes 2^{d(\lambda)}}$ be the product PRFS tester corresponding to G as guaranteed by Corollary 3.9.

We describe the commitment scheme, (C, R) as follows. For notational convenience, we abbreviate $n = n(\lambda)$, $d = d(\lambda)$.

1. *Commit Phase*:

   - The receiver R_λ samples a uniformly random m-qubit Pauli matrix P, where $m = 2^d \cdot n$. We write P as $P_0 \otimes \cdots \otimes P_{2^d-1}$, where P_i is an n-qubit Pauli operator.[9] It sends P to the committer.
   - The committer C_λ on input a bit $b \in \{0,1\}$, does the following:
     - It samples $k \xleftarrow{\$} \{0,1\}^\lambda$.
     - For every $x \in \{0,1\}^d$, computes $\sigma_{k,x} \leftarrow G_\lambda(k,x)$.

     It sends the commitment $\mathbf{c} = \bigotimes_{x\in\{0,1\}^d} \widetilde{\sigma}_{k,x}$, where $\widetilde{\sigma}_{k,x} = P_x^b \sigma_{k,x} P_x^b$, to the receiver.

2. *Reveal Phase*: The committer sends (k,b) as the decommitment to the receiver. If $b \notin \{0,1\}$, receiver outputs $\perp$. The receiever outputs b if and only if $\mathsf{Test}_\lambda^{\otimes 2^d}\left(\{k,x\}_x, P^b \mathbf{c} P^b\right) = 1$ where $P^b = \bigotimes_{x\in\{0,1\}^{2^d}} P_x^b$. Otherwise the receiver outputs $\perp$.

Lemma 6.6. *If G is a PRFS, then there exists a negligible function $\nu(\cdot)$ such that the probability that the honest receiver accepts the honest committer's opening is at least $1-\nu(\lambda)$.*

Proof. This follows immediately from Lemma 3.10 and union bound as 2^d is polynomial in λ. □

Lemma 6.7. *If G is a PRFS, then (C,R) satisfies computational hiding as defined in Definition 6.4.*

Proof. This follows from a standard hybrid argument. Let R^* be a QPT receiver.

Hybrid H_1. This corresponds to C committing to the bit $b=0$. In more detail, let $P = \bigotimes_{x\in\{0,1\}^d} P_x$. be the Pauli sent by R^* to C. Then, C computes $\sigma_{k,x} \leftarrow G_\lambda(k,x)$, for every $x \in \{0,1\}^d$. C sends $\mathbf{c} = \bigotimes_{x\in\{0,1\}^{d(\lambda)}} \widetilde{\sigma}_{k,x}$, where $\widetilde{\sigma}_{k,x} = \sigma_{k,x}$ to R^*.

Hybrid H_2. This hybrid is the same as before, except that $\sigma_{k,x} = |\vartheta_x\rangle\langle\vartheta_x|$, where $|\vartheta_1\rangle, \ldots, |\vartheta_{2^d}\rangle \leftarrow \mathscr{H}_n$.

The output distributions of H_1 and H_2 are computationally indistinguishable from the security of PRFS $\{G_\lambda(\cdot,\cdot)\}_{\lambda\in\mathbb{N}}$.

Hybrid H_3. This hybrid is the same as before, except that $\widetilde{\sigma}_{k,x} = P_x \sigma_{k,x} P_x$. That is, the operator P^b is applied to $\sigma_{k,x}$.

The output distributions of H_2 and H_3 are identical.

Hybrid H_4. This corresponds to C committing to the bit $b=1$. In more detail, let P be the Pauli sent by R^* to C. Then, C computes $\sigma_{k,x} \leftarrow G_\lambda(k,x)$, for every $x \in \{0,1\}^d$. C sends $\mathbf{c} = \bigotimes_{x\in\{0,1\}^{d(\lambda)}} \widetilde{\sigma}_{k,x}$, where $\widetilde{\sigma}_{k,x} = P_x \sigma_{k,x} P_x$ to R^*.

The output distributions of H_3 and H_4 are computationally indistinguishable from the security of PRFS G. □

[9] To sample $P = \bigotimes_i P_i$, the receiver can sample uniformly random bits $\alpha_1, \beta_1, \ldots, \alpha_m, \beta_m$, and let $P_i = X^{\alpha_i} Z^{\beta_i}$ where X and Z are the single-qubit Pauli matrices.

Statistical binding. The rest of the section will be devoted to proving statistical binding of the construction. For this part, we assume recognizable abort to simplify the analysis. However, we believe that with some more work, our construction would still satisfy binding even if a more generic PRFS is used.

Lemma 6.8. *(C, R) satisfies $O(2^{-\lambda})$-statistical binding if the (d, n)-PRFS satisfies recognizable abort property and $2^n \cdot d \geq 7\lambda$.*

Let $C^* = \{C^*_\lambda\}_{\lambda\in\mathbb{N}}$ be an malicious committer. Suppose C^*_λ executes the commit phase with the honest receiver R_λ. Let $\mathbf{c}$ denote the mixed state sent by C^* to R.

We first describe the extractor.

Description of $\mathcal{E}$. On input the commitment $\mathbf{c}$, the extractor $\mathcal{E}$ obtains the description of the Pauli matrix P from the receiver's state, and performs general measurement Λ whose operators are $\{\sqrt{\Lambda_0}, \sqrt{\Lambda_1}, \sqrt{\Lambda_\perp}\}$, where $\Lambda_0, \Lambda_1, \Lambda_\perp$ are positive semi-definite operators defined as follows:

- Define T_0 to be the subspace spanned by $\left\{\bigotimes_{x\in\{0,1\}^{2^d}} |\psi_{k,x}\rangle\langle\psi_{k,x}| \ : \ \forall k \in \{0,1\}^\lambda\right\}$, where the states $|\psi_{k,x}\rangle$ are pure states guaranteed by Definition 3.5. Let Π_0 be a projection that projects onto T_0.
- Define T_1 to be the subspace spanned by $\left\{P\bigotimes_{x\in\{0,1\}^{2^d}} |\psi_{k,x}\rangle\langle\psi_{k,x}| P \ : \ \forall k \in \{0,1\}^\lambda\right\}$. Let Π_1 be a projection that projects onto T_1. Note that $\Pi_1 = P\Pi_0 P$ by definition.
- Let $p = \|\Pi_0 + \Pi_1\|$ (i.e. the maximum eigenvalue of $\Pi_0+\Pi_1$), $\Lambda_0 = p^{-1}\cdot\Pi_0$, and $\Lambda_1 = p^{-1}\cdot\Pi_1$. Define $\Lambda_\perp = I - (\Lambda_0 + \Lambda_1)$. Since Π_0 and Π_1 are projections, $\sqrt{\Lambda_0}$ and $\sqrt{\Lambda_1}$ are well defined. By definition, $\Lambda_0 + \Lambda_1 = p^{-1}(\Pi_0 + \Pi_1) \preccurlyeq I$ and therefore $\Lambda_\perp$ is positive-semidefinite. Thus $\sqrt{\Lambda_\perp}$ is also well-defined.

Let the measurement outcome be $b' \in \{0, 1, \perp\}$ and let the post-measurement state be denoted by $\mathbf{c}'$ after applying the general measurement $\{\sqrt{\Lambda_0}, \sqrt{\Lambda_1}, \sqrt{\Lambda_\perp}\}$. The extractor $\mathcal{E}$ outputs $(\mathbf{c}', b')$. This completes the description of the extractor.

Fact 6.9. *Let $|\phi\rangle$ and $|\psi\rangle$ be two arbitrary m-qubit states. Let $\mathcal{P}_m$ be the m-qubit Pauli group. Then,*

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

Proof. We first observe that $|\langle\psi| P |\phi\rangle|^2 = \mathrm{Tr}\left(P|\psi\rangle\langle\psi| P |\phi\rangle\langle\phi|\right)$; this follows from the fact that the trace of an outer product of two vectors is equivalent to the square of their inner product.

We also use the following fact from [MTW00]: for any m-qubit density matrix ρ,

$$\mathbb{E}_{P\leftarrow\mathcal{P}_m}\left[P\rho P\right] = \frac{I}{2^m}\ . \tag{6}$$

This implies that for all states $|\psi\rangle, |\phi\rangle$,

$$\begin{aligned}
\mathbb{E}_{P\leftarrow\mathcal{P}_m}\left[|\langle\psi|P|\phi\rangle|^2\right] &= \mathbb{E}_{P\leftarrow\mathcal{P}_m}\left[\mathrm{Tr}\left(P|\psi\rangle\langle\psi|P|\phi\rangle\langle\phi|\right)\right] \\
&= \mathrm{Tr}\left(\mathbb{E}_{P\leftarrow\mathcal{P}_m}\left[P|\psi\rangle\langle\psi|P|\phi\rangle\langle\phi|\right]\right) && \text{(from linearity of } \mathbb{E}\text{)} \\
&= \mathrm{Tr}\left(\mathbb{E}_{P\leftarrow\mathcal{P}_m}\left[P|\psi\rangle\langle\psi|P\right]\cdot|\phi\rangle\langle\phi|\right) \\
&= \mathrm{Tr}\left(\frac{I}{2^m}\cdot|\phi\rangle\langle\phi|\right) && \text{(from (6))} \\
&= \frac{1}{2^m}\cdot\mathrm{Tr}\left(|\phi\rangle\langle\phi|\right) \\
&= \frac{1}{2^m}
\end{aligned}$$

as desired. □

Lemma 6.10 (Almost orthogonality of Π_0 and Π_1)**.**

$$\Pr_{P\leftarrow\mathcal{P}_m}\left[p \geq 1 + 3\cdot 2^{-(m-4\lambda)/3}\right] \leq 2^{-(m-4\lambda)/3}.$$

Proof. Let $|\psi\rangle$ be an arbitrary m-qubit pure state. Write $|\psi\rangle = |\alpha\rangle + |\beta\rangle$, where $|\alpha\rangle$ is the projection of $|\psi\rangle$ onto the subspace T_0, and $|\beta\rangle$ is the projection of $|\psi\rangle$ onto the orthogonal complement of T_0. We determine an upper bound for the following quantity:

$$\begin{aligned}
\langle\psi|\left(\Pi_0+\Pi_1\right)|\psi\rangle &= \left(\langle\alpha|+\langle\beta|\right)\left(\Pi_0+\Pi_1\right)\left(|\alpha\rangle+|\beta\rangle\right) \\
&= \left(\langle\alpha|+\langle\beta|\right)\left(|\alpha\rangle+\Pi_1|\alpha\rangle+\Pi_1|\beta\rangle\right) \\
&= \langle\alpha|\alpha\rangle + \langle\alpha|\Pi_1|\beta\rangle + \langle\beta|\Pi_1|\alpha\rangle + \langle\beta|\Pi_1|\beta\rangle + \langle\alpha|\Pi_1|\alpha\rangle \\
&\leq \langle\alpha|\alpha\rangle + \langle\beta|\beta\rangle + 2\left|\langle\alpha|\Pi_1|\beta\rangle\right| + \langle\alpha|\Pi_1|\alpha\rangle \\
&= 1 + 2\left|\langle\alpha|\Pi_1|\beta\rangle\right| + \langle\alpha|\Pi_1|\alpha\rangle \\
&= 1 + 2\sqrt{\langle\alpha|\Pi_1|\beta\rangle\langle\beta|\Pi_1|\alpha\rangle} + \langle\alpha|\Pi_1|\alpha\rangle \\
&\leq 1 + 2\sqrt{\langle\alpha|\Pi_1|\alpha\rangle} + \langle\alpha|\Pi_1|\alpha\rangle \\
&\leq 1 + 3\sqrt{\langle\alpha|\Pi_1|\alpha\rangle} \\
&\leq 1 + 3\sqrt{\mathrm{Tr}(\Pi_0\Pi_1)}
\end{aligned}$$

where we used the fact that since $|\alpha\rangle$ is contained in the support of Π_0, we have $|\alpha\rangle\langle\alpha| \preceq \Pi_0$ and thus $\mathrm{Tr}(|\alpha\rangle\langle\alpha|\,\Pi_1) \leq \mathrm{Tr}(\Pi_0\,\Pi_1)$.

We now estimate the quantity $\mathrm{Tr}(\Pi_0\,\Pi_1)$. Let $\{|u_1\rangle, \ldots, |u_{\dim(T_0)}\rangle\}$ be an orthonormal basis of T_0, so that $\Pi_0 = \sum_{i=1}^{\dim(T_0)} |u_i\rangle\langle u_i|$. Using that $\Pi_1 = P\Pi_0 P$, we have

$$\mathrm{Tr}(\Pi_0\,\Pi_1) = \sum_{i,j=1}^{\dim(T_0)} \langle u_j|P|u_i\rangle\langle u_i|P|u_j\rangle \leq 2^{2\lambda}\cdot\max_{i,j}\left|\langle u_i|P|u_j\rangle\right|^2$$

where we used that $\dim(T_0) \leq 2^\lambda$.

Now, applying Fact 6.9 to $\left|\langle u_i|P|u_j\rangle\right|^2$ and using Markov's inequality we get that for each $i, j \in [\dim(T_0)]$ we have for all $\delta > 0$,

$$\Pr_{P\leftarrow\mathcal{P}_m}\left[\left|\langle u_i|P|u_j\rangle\right|^2 \geq \delta\right] \leq \delta^{-1}2^{-m}.$$

Using a union bound over all i, j,

$$\Pr_{P \leftarrow \mathcal{P}_m}\left[\exists\, i,j : |\langle u_i| P |u_j\rangle|^2 \geq \delta\right] \leq \delta^{-1} 2^{2\lambda - m}$$

which implies

$$\Pr_{P \leftarrow \mathcal{P}_m}\left[\mathrm{Tr}(\Pi_0 \Pi_1) \geq \delta\, 2^{2\lambda}\right] \leq \delta^{-1}\, 2^{2\lambda - m}\ .$$

Putting everything together, since for all $|\psi\rangle$ the quantity $\langle\psi|(\Pi_0 + \Pi_1)|\psi\rangle$ is upper-bounded by a quantity that only depends on $\mathrm{Tr}(\Pi_0\Pi_1)$ which only depends on P, we get

$$\Pr_{P \leftarrow \mathcal{P}_m}\left[\max_{|\psi\rangle}\{\langle\psi|(\Pi_0 + \Pi_1)|\psi\rangle\} \geq 1 + 3\sqrt{\delta}\, 2^{\lambda}\right] \leq \delta^{-1}\, 2^{2\lambda - m}\ .$$

Setting $\delta = 2^{2(\lambda-m)/3}$ we get the desired lemma statement. □

Indistinguishability of Real World and Ideal World. We need to show that the output distributions of $\mathsf{RealExpt}_\lambda^{C^*}$ and $\mathsf{IdealExpt}_\lambda^{C^*,\mathcal{E}}$ as defined in Definition 6.1 are statistically indistinguishable.

To argue this, we set up some notation.

- We assume that after the commit phase, the random Pauli P sent by R in the first message and C^*'s decommitment (k, b) are obtained by measuring some registers of their joint state. Let σ_{XY} denote the joint state of C^* and R after the commit phase, conditioned on the Pauli P and the decommitment (k, b). The register X denotes C^*'s private register and Y denotes R's private register.

- Let ρ_{real} denote the output of $\mathsf{RealExpt}_\lambda^{C^*}$. If $b = \perp$, then $\rho_{\mathsf{real}} = (\sigma_{\mathsf{X}}, |\perp\rangle\langle\perp|)$. Otherwise, since the $\mathsf{Test}_\lambda^{\otimes 2^d}$ is being applied to register Y of $P^b\sigma_{\mathsf{XY}}P^b$ (where P^b is applied to register X), we have

$$\begin{aligned}\rho_{\mathsf{real}} &= \mathop{\mathbb{E}}_{P,k,b} \mathrm{Tr}_{\mathsf{Y}}\left(\mathsf{Test}_\lambda^{\otimes 2^d}(\{k,x\}_x, P^b\sigma_{\mathsf{XY}}P^b)\right)\\ &= \mathop{\mathbb{E}}_{P,k,b} \mathrm{Tr}_{\mathsf{Y}}\left(M_0 P^b\sigma P^b\right) \otimes |b\rangle\langle b| + \mathrm{Tr}_{\mathsf{Y}}\left(M_\perp P^b\sigma P^b\right) \otimes |\perp\rangle\langle\perp|\end{aligned}$$

  where $M_0 = \eta^2\,|\psi\rangle\langle\psi|$ and $M_\perp = I - M_0$ are positive semi-definite operators acting on register Y with $\eta, |\psi\rangle$ given by Corollary 3.9. The expectation is over the choice of random Pauli P and decommitment (k, b).

- Let ρ_{ideal} denote the output of $\mathsf{IdealExpt}_\lambda^{C^*,\mathcal{E}}$. If $b = \perp$, then by definition of the ideal experiment, the output is $(\sigma_{\mathsf{X}}, |\perp\rangle\langle\perp|)$. Otherwise, the general measurement $\{\sqrt{\Lambda_0}, \sqrt{\Lambda_1}, \sqrt{\Lambda_\perp}\}$ is performed first on register Y of the state σ_{XY} to yield outcome $a \in \{0, 1, \perp\}$. Conditioned on outcome a the post-measurement state is $\frac{\sqrt{\Lambda_a}\sigma\sqrt{\Lambda_a}}{\mathrm{Tr}(\Lambda_a\sigma)}$.

  The Pauli operator P^b and then the $\mathsf{Test}_\lambda^{\otimes 2^d}$ circuit is applied to register Y (corresponding to the reveal phase); if the test accepts and the decommitted bit b matches the output a of the

extractor, then the register X and $|\mu\rangle\langle\mu|$ are output. Otherwise the register X and $|\perp\rangle\langle\perp|$ are output. Put together, we get

$$\rho_{\mathsf{ideal}} = \mathop{\mathbb{E}}_{P,k,b} \mathrm{Tr}_{\mathsf{Y}}(N_b\sigma) \otimes |b\rangle\langle b| + \mathrm{Tr}_{\mathsf{Y}}(N_\perp\sigma) \otimes |\perp\rangle\langle\perp|$$

where $N_b = \sqrt{\Lambda_b}P^bM_0P^b\sqrt{\Lambda_b}$ and $N_\perp = I - N_b$. To see that this is correct in the case that the ideal experiment does not output $\perp$, consider that the post-measurement state of the extractor measurement, conditioned on obtaining outcome b, is $\frac{\sqrt{\Lambda_b}\sigma\sqrt{\Lambda_b}}{\mathrm{Tr}(\Lambda_b\sigma)}$. Applying P^b, conditioning on $\mathsf{Test}_\lambda^{\otimes 2^d}$ accepting, and then tracing out the register Y yields the state

$$\mathrm{Tr}_{\mathsf{Y}}\Big(M_0\Big(P^b\sqrt{\Lambda_b}\sigma\sqrt{\Lambda_b}P^b\Big)\Big).$$

Note that all the operators $M_0, P^b, \sqrt{\Lambda_b}$ all act on the register Y, and the partial trace over Y is cyclic with respect to such operators. Thus this is equal to $\mathrm{Tr}_{\mathsf{Y}}(N_b\sigma)$.

We now prove Lemma 6.8.

Proof of Lemma 6.8. Write

$$\rho_{\mathsf{real}} = \mathop{\mathbb{E}}_{P,k,b} \tau^{(b)}_{\mathsf{real}} \otimes |b\rangle\langle b| + \tau^{(\perp)}_{\mathsf{real}} \otimes |\perp\rangle\langle\perp|$$
$$\rho_{\mathsf{ideal}} = \mathop{\mathbb{E}}_{P,k,b} \tau^{(b)}_{\mathsf{ideal}} \otimes |b\rangle\langle b| + \tau^{(\perp)}_{\mathsf{ideal}} \otimes |\perp\rangle\langle\perp|$$

for subnormalized density matrices $\tau^{(\cdot)}_{\mathsf{real}}, \tau^{(\cdot)}_{\mathsf{ideal}}$ which implicitly depend on P, k, b. Since the trace distance is jointly convex we have

$$\begin{aligned}\mathrm{TD}\left(\rho_{\mathsf{real}}, \rho_{\mathsf{ideal}}\right) &\leq \mathop{\mathbb{E}}_{P,k,b} \mathrm{TD}\left(\tau^{(b)}_{\mathsf{real}} \otimes |b\rangle\langle b| + \tau^{(\perp)}_{\mathsf{real}} \otimes |\perp\rangle\langle\perp|, \tau^{(b)}_{\mathsf{ideal}} \otimes |b\rangle\langle b| + \tau^{(\perp)}_{\mathsf{ideal}} \otimes |\perp\rangle\langle\perp|\right) \\ &= \mathop{\mathbb{E}}_{P,k,b} \mathrm{TD}\left(\tau^{(b)}_{\mathsf{real}}, \tau^{(b)}_{\mathsf{ideal}}\right) + \mathrm{TD}\left(\tau^{(\perp)}_{\mathsf{real}}, \tau^{(\perp)}_{\mathsf{ideal}}\right).\end{aligned}$$

Using that $M_\perp = I - M_0$ and $N_\perp = I - N_b$ and that the partial trace is cyclic with respect to operators acting on Y only, we have

$$\begin{aligned}\mathrm{TD}\left(\tau^{(\perp)}_{\mathsf{real}}, \tau^{(\perp)}_{\mathsf{ideal}}\right) &= \mathrm{TD}\left(\mathrm{Tr}_{\mathsf{Y}}\left(M_\perp P^b\sigma P^b\right), \mathrm{Tr}_{\mathsf{Y}}\left(N_\perp\sigma\right)\right) \\ &= \mathrm{TD}\left(\mathrm{Tr}_{\mathsf{Y}}\left(P^bM_0P^b\sigma\right), \mathrm{Tr}_{\mathsf{Y}}\left(N_b\sigma\right)\right) \\ &= \mathrm{TD}\left(\tau^{(b)}_{\mathsf{real}}, \tau^{(b)}_{\mathsf{ideal}}\right).\end{aligned}$$

Thus to prove the Lemma it suffices to prove that the following statement is true: for all $k \in \{0,1\}^\lambda$ and $b \in \{0,1\}$,

$$\mathop{\mathbb{E}}_{P \leftarrow \mathcal{P}_m} \mathrm{TD}\left(\tau^{(b)}_{\mathsf{real}}, \tau^{(b)}_{\mathsf{ideal}}\right) \leq \frac{4}{2^\lambda}. \tag{7}$$

Fix a decommitment (k, b). Recall that $M_0 = \eta^2 |\psi\rangle\langle\psi|$ where $|\psi\rangle = \bigotimes_x |\psi_{k,x}\rangle$, and that $\Lambda_b = p^{-1} \cdot \Pi_b$. Then

$$\begin{aligned} N_b &= \sqrt{\Lambda_b} P^b M_0 P^b \sqrt{\Lambda_b} \\ &= \eta^2 p^{-1} \sqrt{\Pi_b} P^b |\psi\rangle\langle\psi| P^b \sqrt{\Pi_b} \\ &= \eta^2 p^{-1} \Pi_b P^b |\psi\rangle\langle\psi| P^b \Pi_b \end{aligned}$$

where we use the fact that $\sqrt{\Pi_b} = \Pi_b$ since it is a projector. Since Π_b projects onto the span of $\{P^b \bigotimes_x |\psi_{k,x}\rangle : k \in \{0,1\}^\lambda\}$, this means that $\Pi_b P^b |\psi\rangle = P^b |\psi\rangle$, so N_b is equal to

$$\eta^2 p^{-1} P^b |\psi\rangle\langle\psi| P^b = p^{-1} P^b M_0 P^b \ .$$

This means that

$$\begin{aligned} \mathop{\mathbb{E}}_{P \leftarrow \mathcal{P}_m} \mathrm{TD}\left(\tau^{(b)}_{\mathsf{real}}, \tau^{(b)}_{\mathsf{ideal}}\right) &= \mathop{\mathbb{E}}_{P \leftarrow \mathcal{P}_m} \mathrm{TD}\left(\mathrm{Tr}_{\mathsf{Y}}\left(P^b M_0 P^b \sigma\right), \mathrm{Tr}_{\mathsf{Y}}\left(N_b \sigma\right)\right) \\ &= \mathop{\mathbb{E}}_{P \leftarrow \mathcal{P}_m} \mathrm{TD}\left(\mathrm{Tr}_{\mathsf{Y}}\left(P^b M_0 P^b \sigma\right), p^{-1} \mathrm{Tr}_{\mathsf{Y}}\left(P^b M_0 P^b \sigma\right)\right) \end{aligned} \tag{8}$$

If p (which is a function of P) is at most $1 + 3 \cdot 2^{-(4\lambda - m)/3}$ then we say p is *good*, otherwise it is *bad*. By Lemma 6.10 p is bad with probability at most $2^{-(m-4\lambda)/3}$. When p is good, we have

$$\mathrm{TD}\left(\mathrm{Tr}_{\mathsf{Y}}\left(P^b M_0 P^b \sigma\right), p^{-1} \mathrm{Tr}_{\mathsf{Y}}\left(P^b M_0 P^b \sigma\right)\right) \leq 1 - p^{-1} \leq 3 \cdot 2^{-(4\lambda - m)/3}$$

where we used that $\mathrm{TD}\left(\varphi, p^{-1}\varphi\right) \leq 1 - p^{-1} \leq p - 1$ for all subnormalized density matrices φ. Therefore (8) is at most

$$3 \cdot 2^{-(m-4\lambda)/3} + 2^{-(m-4\lambda)/3} = 4 \cdot 2^{-(m-4\lambda)/3}.$$

Averaging over k, b, we get that $\mathrm{TD}\left(\rho_{\mathsf{real}}, \rho_{\mathsf{ideal}}\right) \leq 8 \cdot 2^{-(m-4\lambda)/3}$, which is less than $\frac{8}{2^\lambda}$ when $m \geq 7\lambda$ as desired. □

6.3 Application: Secure Computation

In this section, we show how to base secure computation solely on the existence of a PRS. We start by recalling Bartusek, Coladangelo, Khurana, and Ma's work [BCKM21b] showing the following.

Theorem 6.11 (Implicit from [BCKM21b])**.** *Assuming the existence of quantum statistically binding bit commitments, maliciously secure computation protocols (in the dishonest majority setting) for P/poly exist.*

Remark 6.12. *While both [BCKM21b] and [GLSV21] show that post-quantum one-way functions and quantum communication suffice to obtain protocols for secure computation, the construction of [BCKM21b] has the advantage that it uses the starting commitment scheme as a black box.*

Comparison of the definitions of statistical binding. The application of Theorem 6.11 would be straightforward except for one subtlety, which is that we are using a more general definition of the statistical binding property than required by their paper. Their notion of statistical binding is tailored to commitment schemes with classical messages as it suffices for their purposes. However, their definition is not satisfied by our commitment scheme as explained shortly.

We do not provide the full proof of their theorem with our binding property but we justify below why our notion of statistical binding still suffices for their proof.

We first recall their definition of statistical binding.

Definition 6.13 ([BCKM21a, Definition 3.2])**.** *A bit commitment scheme is statistically binding if for every unbounded-size committer $\mathcal{C}^*$, there exists a negligible function $\nu(\cdot)$ such that with probability at least $1-\nu(\lambda)$ over the measurement randomness in the commitment phase, there exists a bit $b \in \{0,1\}$ such that the probability that the receiver accepts b in the reveal phase is at most $\nu(\lambda)$.*

Our protocol cannot satisfy this property since the honest receiver does not measure the committer's message in any way, and therefore in general it is possible for the committer to commit to an equal superposition of 0 and 1, in which case this binding property is violated. Nonetheless, our statistical binding property (Definition 6.1) is in essence saying the same thing, which is that there is an implicit measurement that could be done to extract the committed bit in a way unnoticeable to the malicious committer. On a high level, we can switch to the ideal world where the bit is extracted, and then complete the proof using the extracted bit.

Compatibility of Theorem 6.11 with Definition 6.1. We claim that our definition of commitments still suffices to recover the proof of Theorem 6.11 using the idea above. Instead of reproducing the proof in full, we instead inform the reader the places where the proof changes if one were to use our definition of commitments. We recommend the reader to first look at [BCKM21b] before reading the details mentioned below.

The proof of Theorem 6.11 uses the statistical binding property in only two places.

1. A special case of [BCKM21b, Theorem 1]: They show how to go from a statistically binding commitment to a computationally-equivocal commitment scheme that preserves the statistical binding property.

2. [BCKM21b, Theorem 2]: They show how to go from an equivocal statistical-binding commitment scheme to an extractable commitment scheme that has statistical hiding property.

For the first step, we note that our statistical binding definition is compatible with their security proof [BCKM21b, Section 4.2]. They essentially use the statistical binding property to argue that the committed bit is information theoretically determined, and use this fact to construct an extractor to show binding. This property is immediately satisfied by our statistical binding property by switching to the ideal world experiment. In the end, we get computationally-equivocal statistically-binding (still according to the relaxed definition) quantum bit commitments from statistically-binding commitments.

For the second step, what they need from the statistical binding property is again some form of inefficient extraction as shown in the following [BCKM21a, Section 5.2], which is also satisfied by our definition: "Since the commitments are statistically binding, the values of $\hat{x}, \hat{\theta}$ that $\mathcal{C}$ would accept are statistically determined after $\mathcal{R}$ commits. Let $|\psi\rangle_{\mathsf{XY}}$ be the joint state of the committer and

receiver where X corresponds to the committer's $2\lambda^3$ registers, and Y corresponds to the receiver's state. Then the committer's check can only pass if the sampling strategy of Lemma 5.3 succeeds for the committed values of $\hat{x}, \hat{\theta}$."

The rest of the proof follows and thus we recover Theorem 6.11 with our statistical binding property. By instantiating the statistically binding bit commitments in Theorem 6.11 with PRS (Theorem 6.5 and Theorem 4.1), we obtain the following corollary.

Corollary 6.14. *Assuming the existence of* $(2\log\lambda + \omega(\log\log\lambda))$*-PRS, there exists maliciously secure computation protocol for P/poly in the dishonest majority setting.*

References

[Aar05] Scott Aaronson. "Quantum Computing, Postselection, and Probabilistic Polynomial-Time". In: *Proceedings: Mathematical, Physical and Engineering Sciences* 461.2063 (2005), pp. 3473–3482. ISSN: 13645021. URL: http://www.jstor.org/stable/30047928 (cit. on p. 8).

[AQY21] Prabhanjan Ananth, Luowen Qian, and Henry Yuen. *Manuscript (in preparation).* 2021 (cit. on pp. 5, 8).

[BB21] Nir Bitansky and Zvika Brakerski. "Classical Binding for Quantum Commitments". In: *Theory of Cryptography - 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8-11, 2021, Proceedings, Part I*. Ed. by Kobbi Nissim and Brent Waters. Vol. 13042. Lecture Notes in Computer Science. Springer, 2021, pp. 273–298. DOI: 10.1007/978-3-030-90459-3_10 (cit. on pp. 4, 6, 23, 24).

[BB84] Charles H. Bennett and Gilles Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Proceedings of International Conference on Computers, Systems & Signal Processing, Dec. 9-12, 1984, Bangalore, India*. 1984, pp. 175–179 (cit. on p. 3).

[BBCS91] Charles H. Bennett, Gilles Brassard, Claude Crépeau, and Marie-Hélène Skubiszewska. "Practical Quantum Oblivious Transfer". In: *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*. Ed. by Joan Feigenbaum. Vol. 576. Lecture Notes in Computer Science. Springer, 1991, pp. 351–366. DOI: 10.1007/3-540-46766-1_29 (cit. on p. 3).

[BCKM21a] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. *One-Way Functions Imply Secure Computation in a Quantum World*. 2021. arXiv: 2011.13486v2 [quant-ph] (cit. on p. 32).

[BCKM21b] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. "One-Way Functions Imply Secure Computation in a Quantum World". In: *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I*. Ed. by Tal Malkin and Chris Peikert. Vol. 12825. Lecture Notes in Computer Science. Springer, 2021, pp. 467–496. DOI: 10.1007/978-3-030-84242-0_17 (cit. on pp. 3, 6–8, 31, 32).

[BMR90] Donald Beaver, Silvio Micali, and Phillip Rogaway. "The Round Complexity of Secure Protocols (Extended Abstract)". In: *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, May 13-17, 1990, Baltimore, Maryland, USA*. Ed. by Harriet Ortiz. ACM, 1990, pp. 503–513. DOI: 10.1145/100216.100287 (cit. on p. 6).

[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*. Ed. by Dennis Hofheinz and Alon Rosen. Vol. 11891. Lecture Notes in Computer Science. Springer, 2019, pp. 229–250. DOI: 10.1007/978-3-030-36030-6_10 (cit. on p. 3).

[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*. Ed. by Daniele Micciancio and Thomas Ristenpart. Vol. 12171. Lecture Notes in Computer Science. Springer, 2020, pp. 417–440. DOI: 10.1007/978-3-030-56880-1_15 (cit. on pp. 3, 5, 8, 12).

[BY20] Zvika Brakerski and Henry Yuen. *Quantum Garbled Circuits*. 2020. arXiv: 2006.01085 [quant-ph] (cit. on p. 6).

[FUYZ20] Junbin Fang, Dominique Unruh, Jun Yan, and Dehua Zhou. *How to Base Security on the Perfect/Statistical Binding Property of Quantum Bit Commitment?* Cryptology ePrint Archive, Report 2020/621. https://ia.cr/2020/621. 2020 (cit. on pp. 6, 23, 24).

[GLSV21] Alex B. Grilo, Huijia Lin, Fang Song, and Vinod Vaikuntanathan. "Oblivious Transfer Is in MiniQCrypt". In: *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part II*. Ed. by Anne Canteaut and François-Xavier Standaert. Vol. 12697. Lecture Notes in Computer Science. Springer, 2021, pp. 531–561. DOI: 10.1007/978-3-030-77886-6_18 (cit. on pp. 3, 6, 31).

[Gol90] Oded Goldreich. "A note on computational indistinguishability". In: *Information Processing Letters* 34.6 (1990), pp. 277–281. ISSN: 0020-0190. DOI: 10.1016/0020-0190(90)90010-U (cit. on p. 3).

[HLW06] Patrick Hayden, Debbie W. Leung, and Andreas Winter. "Aspects of Generic Entanglement". In: *Communications in Mathematical Physics* 265.1 (July 2006), pp. 95–117. ISSN: 1432-0916. DOI: 10.1007/s00220-006-1535-6 (cit. on p. 10).

[IR89] Russell Impagliazzo and Steven Rudich. "Limits on the Provable Consequences of One-Way Permutations". In: *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*. Ed. by David S. Johnson. ACM, 1989, pp. 44–61. DOI: 10.1145/73007.73012 (cit. on p. 3).

[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*. Ed. by Hovav Shacham and Alexandra Boldyreva. Vol. 10993. Lecture Notes in Computer

Science. Springer, 2018, pp. 126–152. DOI: 10.1007/978-3-319-96878-0_5 (cit. on pp. 3, 5, 10, 11).

[Kil88] Joe Kilian. "Founding Cryptography on Oblivious Transfer". In: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*. Ed. by Janos Simon. ACM, 1988, pp. 20–31. DOI: 10.1145/62212.62215 (cit. on p. 3).

[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*. Ed. by Min-Hsiu Hsieh. Vol. 197. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 2:1–2:20. DOI: 10.4230/LIPIcs.TQC.2021.2 (cit. on p. 3).

[MTW00] Michele Mosca, Alain Tapp, and Ronald de Wolf. *Private quantum channels and the cost of randomizing quantum information*. 2000. arXiv: quant-ph/0003101 (cit. on p. 27).

[MY21] Tomoyuki Morimae and Takashi Yamakawa. *Quantum commitments without one-way functions*. 2021. arXiv: 2112.06369 [quant-ph] (cit. on p. 8).

[Nao91] Moni Naor. "Bit commitment using pseudorandomness". In: *Journal of Cryptology* 4.2 (Jan. 1991), pp. 151–158. ISSN: 1432-1378. DOI: 10.1007/BF00196774 (cit. on pp. 6, 7, 25).

[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).

[Unr16] Dominique Unruh. "Computationally Binding Quantum Commitments". In: *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9666. Lecture Notes in Computer Science. Springer, 2016, pp. 497–527. DOI: 10.1007/978-3-662-49896-5_18 (cit. on pp. 6, 23, 24).

[Wie83] Stephen Wiesner. "Conjugate coding". In: *SIGACT News* 15.1 (1983), pp. 78–88. DOI: 10.1145/1008908.1008920 (cit. on p. 3).

[YWLQ15] Jun Yan, Jian Weng, Dongdai Lin, and Yujuan Quan. "Quantum Bit Commitment with Application in Quantum Zero-Knowledge Proof (Extended Abstract)". In: *Algorithms and Computation - 26th International Symposium, ISAAC 2015, Nagoya, Japan, December 9-11, 2015, Proceedings*. Ed. by Khaled M. Elbassioni and Kazuhisa Makino. Vol. 9472. Lecture Notes in Computer Science. Springer, 2015, pp. 555–565. DOI: 10.1007/978-3-662-48971-0_47 (cit. on pp. 6, 23, 24).