Quantum Pseudorandomness Cannot Be Shrunk In a Black-Box Way

Samuel Bouaziz--Ermann[1] and Garazi Muguruza[2,3]

[1] Sorbonne Université, CNRS, LIP6, France
[2] Informatics Institute, University of Amsterdam, Netherlands
[3] QuSoft, Netherlands

Abstract. Pseudorandom Quantum States (PRS) were introduced by Ji, Liu and Song as quantum analogous to Pseudorandom Generators. They are an ensemble of states efficiently computable but computationally indistinguishable from Haar random states. Subsequent works have shown that some cryptographic primitives can be constructed from PRSs. Moreover, recent classical and quantum oracle separations of PRS from One-Way Functions strengthen the interest in a purely quantum alternative building block for quantum cryptography, potentially weaker than OWFs.
However, our lack of knowledge of extending or shrinking the number of qubits of the PRS output still makes it difficult to reproduce some of the classical proof techniques and results. Short-PRSs, that is PRSs with logarithmic size output, have been introduced in the literature along with cryptographic applications, but we still do not know how they relate to PRSs. Here we answer half of the question, by showing that it is not possible to shrink the output of a PRS from polynomial to logarithmic qubit length while still preserving the pseudorandomness property, in a relativized way. More precisely, we show that relative to Kretschmer's quantum oracle (TQC 2021) short-PRSs cannot exist (while PRSs exist, as shown by Kretschmer's work).

1 Introduction

Pseudorandomness is an important concept in cryptography since almost all relevant classical cryptographic primitives require the existence of one-way functions (OWF) or equivalent objects such as pseudorandom generators (PRG) and pseudorandom functions (PRF) [HILL99]. It corresponds to a deterministic function whose output cannot be distinguished from the uniform distribution by a computationally bounded algorithm.

In 2018, Ji, Liu and Song [JLS18] introduced a quantum analog of PRGs, called *Pseudorandom Quantum States* (PRS) that consists of a family of polynomial size keyed-states $\{|\phi_k\rangle\}_{k\in\mathcal{K}}$ such that no quantum polynomial-time algorithm can distinguish between a polynomial number of copies[4] of a randomly sampled element from the PRSs family or a polynomial number of copies of a Haar-random state (see Definition 1 for a formal definition). We already know how to construct several cryptographic primitives from (variants of) PRSs: public key encryption with quantum keys [BGHD+23], quantum digital signatures [MY22b], pseudo one-time pad encryption schemes [AQY22], statistically binding and computationally hiding commitments [AQY22,MY22a,KT23] and quantum computational zero knowledge proofs [BCQ23]. Such rapid interest derives probably from the fact that PRSs can be constructed from OWFs [JLS18] (and thus PRGs), but there are oracle separations found between OWFs and PRSs [Kre21,KQST23], which makes them a potentially weaker building block for quantum cryptography, with a purely quantum description.

Classically, practical applications of PRGs require stretching the output, which can be done arbitrarily by just using the output of the PRG as an input and roughly composing the PRG with itself. However, no analogous intuitive construction is possible with PRSs because the input is a classical bit string and the output is a quantum state. Moreover, although classically shrinking a PRG does not make sense, it is still possible by discarding part of the output of the original PRG. Quantumly, discarding half of a quantum

[4] The reason we give the adversary a polynomial number of copies of the state is that an arbitrary quantum state cannot be cloned, by the no-cloning theorem.

state could lead to a maximally mixed state, making it easy to distinguish from a Haar-random state, which means that pseudorandomness is not a property respected by subsets of the registers.

This motivates the definition of short-PRS, a PRS that on input $k \in \{0,1\}^\lambda$ outputs a qubit of size $n(\lambda) = O(\log \lambda)$. Brakerski and Shmueli [BS20] proved that $c \cdot \log \lambda$-output PRSs exist for any $c \geq 1$. While classically, "short-PRGs" and PRGs are equivalent, in the quantum setting, the following question remains open

Question 1: *What is the relation between short and long PRSs?*

In this work we answer half of the question, by showing that there exists a quantum oracle relative to which PRSs exist but short-PRSs do not, thus PRSs are a weaker assumption than short-PRSs. Note that this question appeared in the literature [BS20,BB21,AQY22].

Our separation relies on another fundamental difference between quantum and classical cryptography; while pseudorandomness has received much attention in classical cryptography, the same cannot be said about *pseudodeterminism*. A pseudodeterministic variant of PRGs (PD-PRG) was first introduced by Ananth, Lin and Yuen [ALY24], defined as a quantum polynomial-time algorithm that outputs a pseudorandom string on a *fraction* of the input keys. In their work, this fraction is polynomial in size and this define polynomial pseudodeterminism, and we define overwhelming pseudodeterminsm when this fraction is overwhelming in size. They showed that short-PRSs can be used to construct PD-PRGs. Later Barhoush, Behera, Ozer, Salvail and Sattath [BBO+24] introduced the analogous pseudodeterministic one-way functions (PD-OWF) and showed that PD-PRGs imply PD-OWFs.

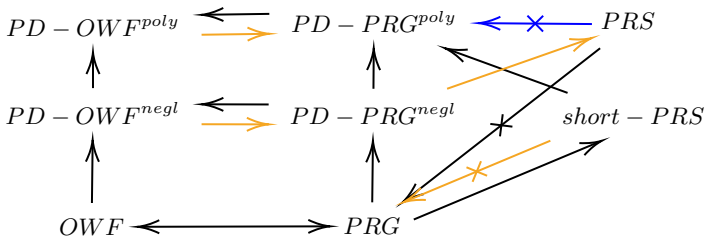

Fig. 1. Relation between different primitives. Proven, expected and our result.

Although a quantum state can encode arbitrary classical information, this information is not necessarily *accessible* for an observer. For example, while OWFs are separated from PRSs, PD-OWFs can be built from short-PRSs [BBO+24]. Another interesting question is thus

Question 2: *What is the role of pseudodeterminism in quantum cryptography?*

Pseudodeterministic variants of OWFs and PRGs seem to be enough to build many interesting tasks in cryptography. For example, [ALY24] state that overwhelming PD-PRGs are enough to build PRSs, as these can be used to generate pseudorandom phases pseudodeterministically, and they build polynomial PD-PRGs from short-PRSs. However, it is unclear how to build overwhelming PD-PRGs from short-PRSs, or if polynomial PD-PRGs are enough to build PRSs[5]. In other words, there seems to be a gap between the cryptography we can build from polynomial or overwhelming pseudodeterminism, with short-PRSs being nearly as strong

[5] The same construction from PRGs does not work as the polynomial error will induce too much error in the resulting state to be a PRS.

as overwhelming pseudodeterminism while PRSs are weaker than polynomial pseudodeterminism, as we show in this work. What is clear is that pseudodeterminism is a non-trivial property in quantum cryptography.

Our contribution. We show that Kretschmer's oracle [Kre21] not only implies that OWFs do not exist, but also none of the pseudoterministic variants do either. Since PD-PRGs and PD-OWFs can be constructed from short-PRSs [ALY24,BBO+24], our work gives a separation between PRSs and short-PRSs[6] and can be stated as follows.

Theorem 1 (Theorem 4, informal). *There exists a quantum oracle $\mathcal{O}$ relative to which PRSs exist but short-PRSs do not.*

This result might sound counterintuitive, as it shows that we cannot shrink a pseudorandom quantum state to a smaller one. An explanation of this result could be that requiring a polynomial number of copies of a logarithmic quantum state to be indistinguishable from a polynomial number of copies of a random state is a strong assumption, that is similar to OWFs, as shown by previous works.

Relative to Kretschmer's oracle we know that poly-size PRSs exist and $\mathsf{PromiseBQP} = \mathsf{PromiseQMA}$. Here we show that not only the existence of OWFs, but also the existence of polynomial error pseudodeterministic OWFs (a possibly weaker assumption, but implied by short-PRSs) also implies $\mathsf{PromiseBQP} \neq \mathsf{PromiseQMA}$. For the proof to work, we rely on the promise version of the complexity classes. Promises problems are such that there are yes instances and no instances, but also other instances where the output of an algorithm does not matter. In our proof, we define a language with the yes instances as the values for which there exists a *high* probability pre-image, and the no instances are values for which there is no *low* probability pre-image. Thus there is a gap between the possible success probabilities, and this gap is needed to distinguish between the yes and no instances in polynomial time.

Open questions. We still leave open the question in the opposite direction, of building PRSs from short-PRSs. There are two partial results regarding this in the literature; the first is that it is possible to construct PRSs with a bigger dimension, but at the cost of the security definition to hold only for single-copy PRSs [GJMZ23], the second gives a positive answer by assuming the existence of pseudorandom isometries [AGKL23]. However, the general question, without assumptions, still remains open.

Note that our separation is relative to a quantum oracle, thus a *classical* oracle separation of short-PRSs from PRSs is unclear. The classical oracle from [KQST23] gives P=NP, but this is not enough for PD-OWFs, as we need promise problems in our proof.

As mentioned above, pseudodeterminism is a non-trivial feature of quantum cryptography. It is unclear where PRSs and short-PRSs lie in the graph of Figure 1: are short-PRSs equivalent to $\mathsf{negl}(n)$ error pseudodeterminism? Which cryptographic tasks require negligible error pseudodeterminism and which ones inverse polynomial? Moreover, although intuitively the answer should be yes, it is still unclear if all the pseudodeterministic variants of PRGs/PRFs/OWFs are equivalent to each other. Constructing PD-OWFs from PD-PRGs is possible by the generalization lemma of [BBO+24], but proving the other direction is still open.

Finally, our results imply that the quantum oracle separation from [Kre21] is not enough to separate short-PRSs from OWFs, thus it is natural to ask if this is possible. Indeed quantum oracles for "short" PRUs will not lead to $\mathsf{PromiseBQP} = \mathsf{PromiseQMA}$, because the concentration inequality on the Haar measure does not hold with small dimensions for the unitaries. As for classical oracles, there is a classical oracle separation between 1-copy short-PRSs and OWFs [KQST23], but a multi-copy separation is still unknown (both for short-PRSs and PRSs).

2 Preliminaries

We use λ to denote the security parameter. We use $\mathsf{negl}(\cdot)$ to denote a negligible function. We use ϵ to denote the empty string. We use $||$ to denote the concatenation operator. We use $x \prec y$ to denote the fact that x is

[6] Note that relative to Kretschmer's oracle, not only do we have PRSs, but we also have pseudorandom unitaries (PRU).

a prefix of y, i.e. there exists x' such that $y = x||x'$. We use $x \not\prec y$ to denote the fact that x is not a prefix of y. We use $\leftarrow \mathcal{A}$ to denote uniform sampling from the set $\mathcal{A}$.

We use $\mathcal{H}_n$ to denote the Haar measure over n-qubit space, i.e. $\mathcal{H}((\mathbb{C})^{\otimes n})$. The Haar measure over $\mathbb{C}^d$ is the uniform measure over all d-dimensional unit vectors over $\mathbb{C}$.

We include here the relevant pseudorandom notions.

Definition 1 (Pseudorandom quantum states [JLS18]). *Let $\lambda \in \mathbb{N}$, and let $n(\lambda)$ be the number of qubits in the quantum system. A keyed family of n-qubit quantum states $\{|\phi_k\rangle\}_{k\in\{0,1\}^\lambda}$ is* pseudorandom *if the following two conditions hold:*

1. ***Efficient generation.*** *There is a QPT algorithm G that on input $k \in \{0,1\}^\lambda$ generates*
$$G_\lambda(k) = |\varphi_k\rangle\langle\varphi_k| .$$
2. ***Pseudorandomness.*** *For any QPT adversary $\mathcal{A}$ and all polynomials $t(\cdot)$, we have*
$$\left| \Pr_{k\leftarrow\{0,1\}^\lambda}\left[\mathcal{A}\left(1^\lambda, |\varphi_k\rangle^{\otimes t(\lambda)}\right) = 1\right] - \Pr_{|\nu\rangle\leftarrow\mathcal{H}_n(\lambda)}\left[\mathcal{A}\left(1^\lambda, |\nu\rangle^{\otimes t(\lambda)}\right) = 1\right]\right| \leq \mathsf{negl}(\lambda).$$

We say that a $n(\lambda)$-PRS is a *short-PRS* if the output is logarithmic in the security parameter, i.e. $n(\lambda) = \Theta(\log\lambda)$. From now on we will use PRSs to refer to long-output PRSs and short-PRSs for logarithmic output.

We also include a pseudodeterministic primitive.

Definition 2 (Quantum Pseudo-deterministic One-Way Functions[7] [BBO+24, Definition 9]). *A QPT algorithm $F : \{0,1\}^{m(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ is a* quantum pseudo-deterministic one-way function *if the following conditions hold:*

- ***Pseudodeterminism.*** *There exists a constant $c > 0$ and function $\mu(\lambda) = O(\lambda^{-c})$ such that for all $\lambda \in \mathbb{N}$, there exists a set $\mathcal{K}_\lambda \subset \{0,1\}^{m(\lambda)}$:*
  1. $\Pr\left[x \in \mathcal{K}_\lambda \,\middle|\, x \leftarrow \{0,1\}^{\ell(\lambda)}\right] \geq 1 - \mu(\lambda)$.
  2. *For any $x \in \mathcal{K}_\lambda$, it holds that*
$$\max_{y\in\{0,1\}^{\ell(\lambda)}} \Pr\left[y = F_\lambda(x)\right] \geq 1 - \mathsf{negl}(\lambda), \tag{1}$$
  where the probability is over the randomness of F_λ.
- ***Security.*** *For every QPT inverter $\mathcal{A}$:*
$$\Pr_{x\leftarrow\{0,1\}^{m(\lambda)}}\left[F\left(\mathcal{A}(F(x))\right) = F(x)\right] \leq \mathsf{negl}(\lambda), \tag{2}$$
  where the probability is over the randomness of F and $\mathcal{A}$.

Note that the pseudodeterminism factor in the above definition comes from the size of the *good* key space $\mu(\lambda)$, which is an inverse-polynomial in the security parameter λ. This means that for a non-negligible number of elements in the key space, the OWF could behave arbitrarily. We could also define a negligible variant by requesting $\mu(\lambda)$ to be a negligible function in λ.

We can build PD-OWFs from short-PRSs.

Theorem 2 (Adapted from [BBO+24, Theorem 6][8]). *Assuming the existence of $(c\log\lambda)$-PRSs with $c > 12$, there exists a $O(\lambda^{-c/12+1})$-PD-OWF $F : \{0,1\}^{\ell(\lambda)} \to \{0,1\}^{\ell(\lambda)}$ with input/output length $\ell(\lambda) = \lambda^{c/6}$.*

[7] In [BBO+24] they actually define Quantum Pseudo-deterministic One-Way *Hash* Functions (PD-OWHF). We omit the *hash* property here for simplicity, but since the security properties of both functionalities are equivalent our proof also trivially works for PD-OWHF.

[8] Here we also use the PD-OWF variant of their theorem originally for PD-OWHF. This choice affects the parameters of the domain and range in the theorem statement because constructing a PD-OWHF requires more steps than constructing a PD-OWF (we only need the first step of their proof). However, note that changing the domain/range of the function to some different polynomials in λ would still make the proof go through by changing some parameters in the proof.

Finally, we will need Kretschmer's (quantum) oracle $\mathcal{O}$ relative to which OWFs do not exist, but PRSs do. The former is because PromiseBQP and PromiseQMA are equal relative to this oracle, we only include here the definitions of these languages for the sake of self-containment, but we refer the reader to the original paper for a lengthier explanation of the complexity classes.

Definition 3. *A promise problem $\mathcal{L} = \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}} \cup \mathcal{L}_{\perp}$ with $\mathcal{L} \subseteq \{0,1\}^*$ is in* PromiseQMA *(Quantum Merlin–Arthur) if there exists a polynomial-time quantum algorithm $\mathsf{V}(x, |\psi\rangle)$ called a* verifier *and a polynomial p such that:*

1. *(Completeness) If $x \in \mathcal{L}_{\text{yes}}$, then there exists a quantum state $|\psi\rangle$ on $p(|x|)$ qubits (called a* witness *or* proof*) such that $\Pr[\mathsf{V}(x, |\psi\rangle) = 1] \geq \frac{2}{3}$.*
2. *(Soundness) If $x \in \mathcal{L}_{\text{no}}$, then for every quantum state $|\psi\rangle$ on $p(|x|)$ qubits, $\Pr[\mathsf{V}(x, |\psi\rangle) = 1] \leq \frac{1}{3}$.*

Definition 4. *A promise problem $\mathcal{L} = \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}} \cup \mathcal{L}_{\perp}$ with $\mathcal{L} \subseteq \{0,1\}^*$ is in* PromiseBQP *(Bounded-error Quantum Polynomial time) if there exists a randomized polynomial-time quantum algorithm $\mathcal{A}(x)$ such that:*

1. *If $x \in \mathcal{L}_{\text{yes}}$, then $\Pr[\mathcal{A}(x) = 1] \geq \frac{2}{3}$.*
2. *If $x \in \mathcal{L}_{\text{no}}$, then $\Pr[\mathcal{A}(x) = 1] \leq \frac{1}{3}$.*

Theorem 3 ([Kre21]). *There exists a quantum oracle $\mathcal{O}$, such that:*

1. $\mathsf{PromiseBQP}^{\mathcal{O}} = \mathsf{PromiseQMA}^{\mathcal{O}}$.
2. *λ-PRSs exist relative to $\mathcal{O}$.*[9]

3 Separating PRSs from short-PRSs

In this section, we prove our main theorem, that we can separate poly-size PRSs and log-size PRSs. Formally, we will be proving the following theorem.

Theorem 4. *There exists a quantum oracle $\mathcal{O}$ such that relative to $\mathcal{O}$, λ-PRSs exist, but $(c \log \lambda)$-PRSs with $c > 12$ do not exist.*

The oracle necessary for the separation is actually the same oracle that Kretschmer used to separate PRSs and OWFs (Theorem 3). It turns out that this oracle is also separating PD-OWFs from PRSs, we prove here that if PromiseBQP = PromiseQMA, then we do not have PD-OWFs. This in addition with Barhoush and Salvail's result that short-PRSs are enough to build PD-OWFs (Theorem 2) will give us the theorem.

According to the above considerations, the main theorem follows directly from the following proposition.

Proposition 1. *If PD-OWFs exist, then* PromiseBQP $\neq$ PromiseQMA.

Proof. Let $\lambda \in \mathbb{N}$ and $F : \{0,1\}^* \to \{0,1\}^{\ell(\lambda)}$ be a PD-OWF. Let us define a promise language $\mathcal{L} = \mathcal{L}_{\text{yes}} \cup \mathcal{L}_{\text{no}} \cup \mathcal{L}_{\perp}$ with $\mathcal{L} \subseteq \{0,1\}^*$ where *yes* instances have a pre-image with respect to F_λ but *no* instances do not. Formally,

$$\mathcal{L}_{\text{yes}} = \Big\{ (1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)} \Big| \exists x \in \{0,1\}^*, x' \prec x \text{ and } \Pr[y = F_\lambda(x)] \geq 1 - \mathsf{negl}(\lambda) \Big\}, \quad (3)$$

$$\mathcal{L}_{\text{no}} = \Big\{ (1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)} \Big| \forall x \in \{0,1\}^*, x' \not\prec x \text{ or } \Pr[y = F_\lambda(x)] \leq 1 - \frac{1}{\lambda} \Big\}. \quad (4)$$

We claim that $\mathcal{L} \in$ PromiseQMA but $\mathcal{L} \notin$ PromiseBQP, thus there must be a separation between both complexity classes. These claims are proven in Lemma 1 and Lemma 2 respectively. □

[9] The random oracle is a random unitary operation, which gives the existence of PRUs and thus PRSs [JLS18].

We now prove the two claims from the proposition. We start by showing that the language defined in proposition 1 is in PromiseQMA, this is, we will construct an algorithm (verifier) that given an element of the domain and a (quantum) proof can distinguish if the element is a yes or no instance of the language.

Lemma 1. *Let F be a PD-OWF and let $\mathcal{L}$ be the language defined in Proposition 1, then $\mathcal{L} \in$ PromiseQMA.*

Proof. We define a quantum polynomial-time algorithm $\mathcal{A}$ that given an element of the domain $(1^{\ell(\lambda)}, x', y) \in \{1^{\ell(\lambda)}\} \times \{0,1\}^* \times \{0,1\}^{\ell(\lambda)}$ and a classical proof $x \in \{0,1\}^{\ell(\lambda)}$, will check if the proof x is indeed a pre-image of the PD-OWF by checking if it coincides with the input y.

Algorithm 1: $\mathcal{A}((1^{\ell(\lambda)}, x', y), x)$

```
if x' ⊀ x :
    return 0
for 1 ≤ i ≤ 2λ :
    if F_λ(x) ≠ y :
        return 0
return 1
```

Note that the algorithm runs in polynomial-time trivially because computing F_λ is done efficiently by definition and we make 2λ calls to it. We now prove that the algorithm distinguishes between the yes/no instances.
(i) Let $(1^{\ell(\lambda)}, x', y) \in \mathcal{L}_{\text{yes}}$. Then by definition there exists a proof $x \in \{0,1\}^*$ such that

$$x' \prec x \text{ and } \Pr[y = F_\lambda(x)] \geq 1 - \mathsf{negl}(\lambda).$$

Then the proof x will be an element of the input $((1^{\ell(\lambda)}, x', y), x)$ for which the algorithm $\mathcal{A}$ will output 1 with high probability because

$$\begin{aligned}\Pr\left[\mathcal{A}((1^{\ell(\lambda)}, x', y), x) = 1\right] &= \Pr\left[\forall 1 \leq i \leq 2\lambda, \text{ the execution of } F_\lambda(x) \text{ outputs y}\right] \\ &\geq (1 - \mathsf{negl}(\lambda))^{2\lambda} \geq 2/3,\end{aligned}$$

which holds whenever $\lambda \geq 6$. Indeed, recall that $\mathsf{negl}(\lambda) \leq 1/\lambda^c$ for all $c > 1$, thus in particular $\mathsf{negl}(\lambda) \leq 1/\lambda^2$, hence

$$(1 - \mathsf{negl}(\lambda))^{2\lambda} \geq \left(1 - \frac{1}{\lambda^2}\right)^{2\lambda} \geq \frac{2}{3},$$

whenever $\lambda \geq 6$, where we used that we have an increasing function in λ.
(ii) Let $(1^{\ell(\lambda)}, x', y) \in \mathcal{L}_{\text{no}}$. Then by definition for every potential proof $x \in \{0,1\}^*$ we have that either

$$x' \not\prec x \quad \text{or} \quad \Pr[y = F_\lambda(x)] \leq 1 - \frac{1}{\lambda}.$$

Then for every possible input $((1^{\ell(\lambda)}, x', y), x)$ the algorithm will output 0 with high probability because

$$\begin{aligned}\Pr\left[\mathcal{A}((1^{\ell(\lambda)}, x', y), x) = 1\right] &= \Pr\left[x' \not\prec x \wedge \forall 1 \leq i \leq 2\lambda, \text{ the execution of } F_\lambda(x) \text{ outputs y}\right] \\ &\leq \Pr\left[\forall 1 \leq i \leq 2\lambda, \text{ the execution of } F_\lambda(x) \text{ outputs y}\right] \\ &\leq \left(1 - \frac{1}{\lambda}\right)^{2\lambda} \leq e^{-2} \leq 1/3.\end{aligned}$$

□

Lemma 2. *Let F be a PD-OWF and let $\mathcal{L}$ be the language defined in Proposition 1, then $\mathcal{L} \notin \mathsf{PromiseBQP}$.*

Proof. We will prove this by contradiction. Let us assume that instead $\mathcal{L} \in \mathsf{PromiseBQP}$, this is, there exists a BQP algorithm $\mathcal{A}$ such that:

1. If $(1^{\ell(\lambda)}, x', y) \in \mathcal{L}_{\text{yes}}$, then $\Pr\left[\mathcal{A}(1^{\ell(\lambda)}, x', y) = 1\right] \geq 2/3$.
2. If $(1^{\ell(\lambda)}, x', y) \in \mathcal{L}_{\text{no}}$, then $\Pr\left[\mathcal{A}(1^{\ell(\lambda)}, x', y) = 1\right] \leq 1/3$.

Without loss of generality, we can assume that the algorithm $\mathcal{A}$ has completeness $1 - \frac{1}{\ell(\lambda)}$ and soundness $\frac{1}{\ell(\lambda)}$. We will now show how we can construct a QPT algorithm $\mathcal{A}'$ that finds a pre-image of every F_λ with high probability when it exists, by querying the original BQP algorithm $\mathcal{A}$ at most $\ell(\lambda) + 1$ times.

Algorithm 2: $\mathcal{A}'(1^{\ell(\lambda)}, y)$

```
b ← A(1^{ℓ(λ)}, ϵ, y)
if b = 0 :
    return ⊥
x_0 ← ϵ
for 1 ≤ i ≤ ℓ(λ) :
    b ← A(1^{ℓ(λ)}, x_0||0, y)
    if b = 1 :
        x_0 = x_0||0
    else :
        x_0 = x_0||1
return x_0
```

Indeed if $(1^{\ell(\lambda)}, \epsilon, y) \in \mathcal{L}_{\text{yes}}$, then the probability that the algorithm $\mathcal{A}'$ outputs a correct pre-image is very high

$$\Pr\left[y = \underset{y \in \{0,1\}^{\ell(\lambda)}}{\arg\max} \Pr\left[y = F_\lambda(x)\right] \,\middle|\, x \leftarrow \mathcal{A}'\left(1^{\ell(\lambda)}, y\right)\right] \geq \left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}. \tag{5}$$

However, this raises a contradiction with the security of the PD-OWF from the assumption Definition 2,

$$\begin{aligned}
&\Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x)\right] \\
&\quad= \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\middle|\, x \in \mathcal{K}_\lambda\right] \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[x \in \mathcal{K}_\lambda\right] \\
&\qquad+ \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\middle|\, x \notin \mathcal{K}_\lambda\right] \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[x \notin \mathcal{K}_\lambda\right] \\
&\quad\geq \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\middle|\, x \in \mathcal{K}_\lambda\right] \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[x \in \mathcal{K}_\lambda\right] \\
&\quad\geq \Pr_{x \leftarrow \{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x) \,\middle|\, x \in \mathcal{K}_\lambda\right] (1 - \mu(\lambda)),
\end{aligned}$$

where the first equality comes from the law of total probability and the second inequality comes from the property of $\mathcal{K}_\lambda$. We can rewrite the last element as:

$$\begin{aligned}
&\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left[y_3 = y_1 \;\middle|\; y_1, y_2, y_3 \leftarrow F_\lambda(x), x_1 \leftarrow \mathcal{A}'(1^{\ell(\lambda)}, y_2), x \in \mathcal{K}_\lambda\right] \\
&\quad\geq \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left[\left(y_3 = y_2 = y_1 = \operatorname*{arg\,max}_{y\in\{0,1\}^{\ell(\lambda)}} \Pr[y = F_\lambda(x)]\right) \wedge (x_1 = x) \;\middle|\; \begin{array}{c} y_1, y_2, y_3 \leftarrow F_\lambda(x) \\ x_1 \leftarrow \mathcal{A}'(1^{\ell(\lambda)}, y_2), x \in \mathcal{K}_\lambda \end{array}\right] \\
&\quad= \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left[F_\lambda(x) = \operatorname*{arg\,max}_{y\in\{0,1\}^{\ell(\lambda)}} \Pr[y = F_\lambda(x)] \;\middle|\; x \in \mathcal{K}_\lambda\right]^3 \\
&\qquad\cdot \Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left[\mathcal{A}'\left(1^{\ell(\lambda)}, \operatorname*{arg\,max}_{y\in\{0,1\}^{\ell(\lambda)}} \Pr[y = F_\lambda(x)]\right) = x \;\middle|\; x \in \mathcal{K}_\lambda\right] \\
&\quad\geq (1 - \mathsf{negl}(\lambda))^3 \left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1},
\end{aligned}$$

where the first inequality comes from the law of total probability, and the last inequality from the definition of a PD-OWF and Equation (5). This gives that:

$$\begin{aligned}
\Pr_{x\leftarrow\{0,1\}^{\ell(\lambda)}}\left[F_\lambda(\mathcal{A}'(1^{\ell(\lambda)}, F_\lambda(x))) = F_\lambda(x)\right] &\geq (1 - \mathsf{negl}(\lambda))^3 \left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1} (1 - \mu(\lambda)) \\
&\geq (1 - O(\mu(\lambda))) \left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1}.
\end{aligned}$$

Note that this bound is not negligible since

$$\left(1 - \frac{1}{\ell(\lambda)}\right)^{\ell(\lambda)+1} \geq \frac{1}{10},$$

whenever $\lambda \geq 2$, which contradicts Equation (2). □

Acknowledgements

We deeply thank Alex Bredariol Grilo for many precious discussions. GM was supported by the Quantum Delta NL visitor's programme travel grant, which enabled the collaboration.

References

AGKL23. Prabhanjan Ananth, Aditya Gulati, Fatih Kaleoglu, and Yao-Ting Lin. Pseudorandom isometries, 2023.

ALY24. Prabhanjan Ananth, Yao-Ting Lin, and Henry Yuen. Pseudorandom Strings from Pseudorandom Quantum States. In Venkatesan Guruswami, editor, *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, volume 287 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:22, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

AQY22. Prabhanjan Ananth, Luowen Qian, and Henry Yuen. Cryptography from pseudorandom quantum states. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 208–236. Springer, Heidelberg, August 2022.

BB21. Nir Bitansky and Zvika Brakerski. Classical binding for quantum commitments. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 273–298. Springer, Heidelberg, November 2021.

BBO+24. Mohammed Barhoush, Amit Behera, Lior Ozer, Louis Salvail, and Or Sattath. Signatures from pseudorandom states via ⊥-prfs, 2024.

BCQ23. Zvika Brakerski, Ran Canetti, and Luowen Qian. On the Computational Hardness Needed for Quantum Cryptography. In Yael Tauman Kalai, editor, *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, volume 251 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:21, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

BGHD+23. Khashayar Barooti, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Quoc-Huy Vu, and Michael Walter. Public-key encryption with quantum keys. In Guy Rothblum and Hoeteck Wee, editors, *Theory of Cryptography*, pages 198–227, Cham, 2023. Springer Nature Switzerland.

BS20. Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 417–440. Springer, Heidelberg, August 2020.

GJMZ23. Sam Gunn, Nathan Ju, Fermi Ma, and Mark Zhandry. Commitments to quantum states. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1579–1588, New York, NY, USA, 2023. Association for Computing Machinery.

HILL99. Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

JLS18. Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 126–152. Springer, Heidelberg, August 2018.

KQST23. William Kretschmer, Luowen Qian, Makrand Sinha, and Avishay Tal. Quantum cryptography in algorithmica. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1589–1602, New York, NY, USA, 2023. Association for Computing Machinery.

Kre21. William Kretschmer. Quantum Pseudorandomness and Classical Complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021)*, volume 197 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 2:1–2:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

KT23. Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. Cryptology ePrint Archive, Paper 2023/1620, 2023. https://eprint.iacr.org/2023/1620.

MY22a. Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. Cryptology ePrint Archive, Paper 2022/1336, 2022. https://eprint.iacr.org/2022/1336.

MY22b. Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 269–295. Springer, Heidelberg, August 2022.