

On the Computational Hardness of Quantum One-Wayness

Bruno Cavalar Eli Goldin Matthew Gray Peter Hall
University of Warwick New York University University of Oxford New York University

Yanyi Liu Angelos Pelecanos
Cornell University UC Berkeley

13 December 2023

Abstract

There is a large body of work studying what forms of computational hardness are needed to realize classical cryptography. In particular, one-way functions and pseudorandom generators can be built from each other, and thus require equivalent computational assumptions in order to be realized. Furthermore, the existence of either of these primitives implies that $P \neq NP$, which gives a lower bound on the necessary hardness.

One can also define versions of each of these primitives with quantum output: respectively one-way state generators and pseudorandom state generators. Unlike in the classical setting, it is not known whether either primitive can be built from the other. Although it has been shown that pseudorandom state generators for certain parameter regimes can be used to build one-way state generators, the implication has not been previously known in full generality. Furthermore, to the best of our knowledge the existence of one-way state generators has no known implications in traditional complexity theory.

We show that pseudorandom states compressing n bits to $\log n + 1$ qubits can be used to build one-way state generators and that pseudorandom states compressing n bits to $\omega(\log n)$ qubits are *themselves* one-way state generators. This is a nearly optimal result, since pseudorandom states with fewer than $c \log n$ -qubit output can be shown to exist unconditionally. We also show that any one-way state generator can be broken by a quantum algorithm with classical access to a PP oracle.

An interesting implication of our results is that a $t(n)$ -copy one-way state generator exists unconditionally, for every $t(n) = o(n/\log n)$. This contrasts nicely with the previously known fact that $O(n)$ -copy one-way state generators require computational hardness. We also outline a new route towards a black-box separation between one-way state generators and quantum bit commitments.

Contents

1	Introduction	2
1.1	Results	3
1.2	Technical Overview	6
1.3	Conclusion and Future work	8

2 Preliminaries	9
2.1 Computational complexity	9
2.2 Quantum information theory and cryptography	10
2.3 Probability distributions	11
2.4 Approximate t -designs	12
3 One-way state generators from compressing pseudorandom states	13
3.1 Unconditional OWSGs from efficient approximate t -designs	17
4 Breaking one-way state generators with a PP oracle	18

1 Introduction

The vast majority of useful classical cryptographic primitives share the following property: they can be used to build one-way functions in a black-box manner. A one-way function is a function on bit-strings which can be efficiently evaluated but is hard to invert. In this sense, one-way functions can be thought of as a “minimal” cryptographic primitive. However, any one-way function can be broken by an efficient algorithm with access to an NP oracle. This means that if $P = NP$, then one-way functions do not exist. As it is unknown whether $P = NP$ or not, the existence of one-way functions, and thus all of classical cryptography, must rely on computational assumptions.

This issue led to the natural desire to “map out” the world of classical cryptography. Over many years, cryptographers have done a fairly good job of figuring out which cryptographic primitives can be built from each other. This cartography helps give a sense of the relative strength of assuming the existence of different cryptographic primitives.

As an example, it is known how to construct one-way functions from any key exchange protocol, i.e. a protocol where two parties can agree on a secret using only communication over a public channel [BCG89]. However, there is strong evidence that building a key exchange protocol from a one-way function is difficult [IR89]. The primitives which can be built from one-way function form a crypto-complexity class known as “MiniCrypt” [Imp95]. Two cryptographic primitives in this class of particular note are pseudorandom generators and commitment schemes. A pseudorandom generator is a deterministic function which maps a small amount of randomness to a longer string indistinguishable from random. A commitment scheme is a process by which a party can encode some string into a “commitment”, such that later the party can prove this “commitment” was an encoding of the original string.

In recent years, cryptographers have started to consider what happens if we allow cryptographic primitives to have quantum output. Here, the landscape of relations between primitives looks very different. Of particular note, it was shown that quantum key distribution, a quantum variant of key exchange, exists unconditionally [BB14, Wie83]. On the other hand, it is known that the quantum versions of one-way functions, pseudorandom generators, and commitments cannot be secure against information-theoretic attackers, and thus require some computational hardness in order to exist [LC97, JLS18, KT23]. These variants are known as one-way state generators, pseudorandom state generators, and quantum bit commitments respectively.

However, it is still unclear what this hardness looks like from a complexity perspective. In particular, it is known there exists an oracle relative to which $BQP \supseteq NP$, but all three of these

primitives exist¹ [Kre21]. Furthermore, we are still mapping out the relations between quantum primitives. It was only recently discovered that quantum bit commitments can be built from one-way state generators [KT23], and it is still an open question as to whether pseudorandom state generators can be built from quantum bit commitments.

The main goal of this work is to broaden our understanding of the hardness of quantum primitives, with a particular focus on one-way state generators. In particular, we show two main results:

1. One-way state generators can be built from pseudorandom states for nearly all parameter regimes requiring computational hardness.
2. If one-way state generators exist, then $\text{BQP} \neq \text{PP}$.

These main results bring along a number of interesting implications. The following are of particular note:

1. A fixed-copy version of one-way state generators exists unconditionally.
2. If we can show that quantum bit commitments exist relative to a PP oracle, then there is a black-box separation showing it is unlikely that we will be able to build one-way state generators from quantum bit commitments.

1.1 Results

We now recall a few key concepts from quantum cryptography and give more details about the results we show.

Pseudorandom State Generators (PRS). A pseudorandom state generator, originally defined in [JLS18], is a quantum variant of a pseudorandom generator. Given a classical key k , a PRS maps k to a quantum pure state $|\phi_k\rangle$. The security guarantee is that the output of a PRS on a random input should look like a random state. That is, it is hard for any quantum adversary to distinguish any polynomial number of copies of a random $|\phi_k\rangle$ from polynomial copies of a Haar random state.

The relationship between the length of the input key n and the number of output qubits m determines whether a PRS can exist information-theoretically or requires computational assumptions. In particular, [AGQY22] shows that PRSs with output state length $m \geq \log n$ qubits can be broken by an inefficient adversary, and thus must be a computational object. On the other hand, it is known that PRSs with state length $m \leq c \log n$ exist unconditionally for some $c \in (0, 1)$ [AGQY22, BS20].

One Way State Generators (OWSG). A one-way state generator, originally defined in [MY22a], is a quantum variant of a one-way function. Just like for PRS, a OWSG maps a classical key k to quantum state $|\phi_k\rangle$. The security guarantee of a OWSG is that, given any polynomial number of copies of $|\phi_k\rangle$, it is hard for a quantum algorithm to find keys k' such that $|\phi_k\rangle, |\phi_{k'}\rangle$ have noticeable overlap. OWSGs can also be defined to have mixed state outputs [MY22b], although we will not consider this variant in this work.

¹For one-way state generators and quantum bit commitments, the result follows from [Kre21] and the subsequent works of [MY22b, KT23].

Building OWSGs from PRSs. It is known that any expanding PRS is also a OWSG [MY22a]. Here, an expanding PRS is one which has keys of length n and output states of length $m > (1 + c)n$ for some $c > 0$. We extend this proof to show that any PRS with output length at least $m \geq \log n + 1$ implies OWSGs. Since OWSGs require computational hardness [KT23, LC97], and there exists $d < 1$ such that PRSs with output length $\leq d \log n$ exist unconditionally [BCQ22], this reduction is close to optimal.

Theorem 1.1 (Informal version of Theorem 3.3). *For every $c \geq 1$, if there exists a PRS mapping n -bit strings to $(\log n + c)$ -qubit states, then OWSGs exist.*

Through a closely related argument, we also find that PRSs that map n bits to $\omega(\log n)$ qubits are OWSGs.

Theorem 1.2 (Informal version of Theorem 3.4). *Any PRS that maps n -bit strings to $\omega(\log n)$ -qubit states is also a OWSG.*

Fixed-copy PRSs and OWSGs. Both PRSs and OWSGs are defined to be secure against adversaries that are given any polynomial number of copies of the output state. However, we could instead consider an alternative definition where we fix the number of copies given to the adversary. We will refer to these primitives by the names t -copy PRS and t -copy OWSG. Related primitives have already been considered in a number of works, including [GC01, KT23, LMW23]. A summary of prior work and our results can be found in Table 1.

It is known that for any fixed function t , expanding $t(\lambda)$ -copy PRSs require computational hardness because they can be used to construct quantum bit commitments [LMW23, Yan22, MY22a, BCQ22], where λ is the security parameter. Therefore, any expanding $t(\lambda)$ -copy PRS can be broken by an inefficient attacker. On the other hand, if we do not have an expansion requirement, it can be shown that something called an efficient approximate t -design (defined formally in Section 2.4) is also a t -copy PRS [Kre21]. Since efficient approximate t -designs exist unconditionally [DCEL09, HMMH⁺23, OSP23], so do t -copy (non-expanding) PRSs.

Thus, one may ask the question: for what parameters do t -copy OWSGs require computational assumptions? In a recent work, Khurana and Tomer [KT23] show that Weisner encodings / BB84 states [Wie83, BB14] are 1-copy OWSGs. Additionally, written twenty years before OWSGs were defined, [GC01] shows that t -qubit stabilizer states are $t/2$ -copy OWSGs. The OWSG construction of [GC01] only has a weaker security guarantee, but this can be resolved by amplification.

Note that this means that for any fixed polynomial t , $t(\lambda)$ -copy OWSGs exist unconditionally, where λ is the security parameter. However, [KT23] shows that quantum bit commitments can be built from $\Theta(n)$ -copy OWSGs, where n is the input key length. This is not a contradiction, since the number of copies of security here depends on the input length instead of the security parameter. Thus, we may consider the following refinement of our question:

For what functions $t(\cdot)$ do $t(n)$ -copy OWSGs require computational assumptions?

Our proof of Theorem 1.2 will also imply the following result.

Corollary 1.3 (Informal version of Corollary 3.6). *Every efficient approximate t -design mapping n bits to $\omega(\log n)$ qubits is also a $(t - 1)$ -copy OWSG.*

If we consider state-of-the-art constructions of approximate t -designs [OSP23], we in addition prove the following.

Primitive	Copies	Security	Comments
Expanding PRS	$t \geq 1$ copy	Computational	[MY22a, Yan22, BCQ22]
PRS	$\text{poly}(\lambda)$ -copy	Statistical	Approximate t -designs
OWSG	$\text{poly}(\lambda)$ -copy	Statistical	Theorem 3.4 with approximate t -designs
PRS	$O(n/\log n)$ -copy	Statistical	Approximate t -designs
OWSG	$O(\sqrt{n})$ -copy	Statistical	Stabilizer states [GC01]
OWSG	$\Omega(n)$ -copy	Computational	Implies quantum bit commitments [KT23]
OWSG	$o(n/\log n)$ -copy	Statistical	Corollary 3.7

Table 1: A summary of what is known about the computational and information-theoretic nature of quantum cryptographic primitives, based on the number of copies of the output given to the adversary. We say that security is computational if the existence of the primitive requires computational hardness and we say that the security is statistical if the primitive can be shown to exist unconditionally against statistical adversaries. The rows above the dashed line correspond to constructions where the number of copies is in terms of the security parameter λ . The rows below the dashed line correspond to constructions where the number of copies is in terms of n , the number of input bits.

Corollary 1.4 (Equivalent to Corollary 3.7). *For every $t(n) = o(n/\log n)$, there exists a $t(n)$ -copy OWSG.*

Aside from demonstrating an interesting new property of approximate t -designs, these results give an interesting dichotomy: OWSGs require computational hardness for $\Omega(n)$ -copies, and exist unconditionally for $o(n/\log n)$ copies.

Quantum cryptography, computational complexity, and separations. It is known from [Kre21, AGQY22] that the existence of a PRS which outputs a $(\log n + O(1))$ -qubit state implies that $\text{BQP} \neq \text{PP}$, where BQP refers to the class of problems efficiently solvable by quantum computers, and PP refers to the class of problems such that a probabilistic Turing machine gets the correct answer with probability strictly greater than $\frac{1}{2}$. Thus, like with one-way functions, the existence of PRSs has implications in complexity theory.

However, no similar results are known about OWSGs or quantum bit commitments. In fact, it is conjectured by [LMW23] that quantum bit commitments may exist relative to a random oracle and *any* classical oracle, even ones that depend on the random oracle. Khurana and Tomer [KT23] observe that there exists a classical oracle that breaks OWSGs, which implies that such a conjecture cannot extend to the existence of OWSGs. If the conjecture of [LMW23] is proven, this would provide a black box separation between OWSGs and quantum bit commitments. However, it is not immediately clear that the oracle [KT23] mentions lies inside any interesting complexity class.

We show that the existence of OWSGs indeed does have interesting complexity implications. In particular, we show the following.

Theorem 1.5 (Informal version of Corollary 4.5). *If OWSGs exist, then $\text{BQP} \neq \text{PP}$.*

It then follows that a black box separation between quantum bit commitments and OWSGs can be achieved by proving a weaker version of the conjecture of [LMW23], namely that there exists an oracle \mathcal{O} relative to which quantum bit commitments exist and $\text{PP}^{\mathcal{O}} \subseteq \text{BQP}^{\mathcal{O}}$.

1.2 Technical Overview

We now give an overview of the main technical ingredients of our work.

1.2.1 Building one-way state generators from pseudorandom state generators

Expanding PRS. We will begin this section by detailing the argument from [MY22a] that an expanding PRS is also a OWSG. Recall that an expanding PRS maps n -bit strings to m -qubit quantum states, where $m \geq (1+c)n$ for some $c > 0$. The natural reduction which uses the OWSG adversary to also break the PRS works. Let \mathcal{A} be a OWSG adversary which outputs k' such that $|\phi_{k'}\rangle$ close to $|\phi_k\rangle$ using t copies of $|\phi_k\rangle$. Given $t+1$ copies of a state $|\psi\rangle$, we can test whether it is an output of the PRS or Haar-random as follows: run \mathcal{A} on the first t copies to get a state $|\phi_{k'}\rangle$, and compare $|\phi_{k'}\rangle$ with the last copy of $|\psi\rangle$.

If $|\psi\rangle = |\phi_k\rangle$ for some k , then $|\phi_{k'}\rangle$ will be close to $|\phi_k\rangle$. If $|\psi\rangle$ is Haar-random, then since it is a random state on $(1+c)n$ qubits, with high probability it is far from $|\phi_k\rangle$ for all 2^n values of k . This is because $\frac{2^n}{2^{(1+c)n}}$ is negligible in n .

Improvement to shrinking PRS. To improve this result to PRS with $O(\log n)$ bit output, we simply improve the analysis of the exact same reduction. We make the following simple observation about Haar-random states: for any fixed m -qubit state $|\phi\rangle$, the probability that a Haar-random state $|\psi\rangle$ is “close” to $|\phi\rangle$ is $2^{-\Omega(2^m)}$. Thus, in the reduction above, if $|\psi\rangle$ is a Haar-random state on $2\log n + 1$ qubits, then, by a union bound, with high probability it is far from $|\phi_k\rangle$ for all 2^n values of k . This argument is enough to derive Theorem 1.2. Using amplification [MY22b], we also find that OWSGs can be built even from a PRS which outputs $\log n + O(1)$ qubits, obtaining Theorem 1.1.

Building OWSGs from a fixed-copy PRS. We can instantiate our reduction with a t -copy PRS (i.e. a PRS that is secure against t copies). Our reduction then shows that any t -copy PRS is also a $(t-1)$ -copy OWSG. The fact that approximate t -designs are t -copy PRSs [Kre21] then gives Corollary 1.3.

Recently, O’Donnell, Servedio, and Paredes [OSP23] showed that there exists an efficient $2^{-\lambda}$ -approximate t -design on m -qubit quantum states with seed length $n = O(mt + \lambda)$. Setting $m = \omega(\log n)$ and $t = o\left(\frac{n}{\log n}\right)$ shows that approximate $o\left(\frac{n}{\log n}\right)$ -designs with $\omega(\log n)$ -output bits exist, and thus $o\left(\frac{n}{\log n}\right)$ -copy OWSGs also exist. That is, Corollary 1.4 holds.

1.2.2 Breaking one-way state generators with a PP oracle

The power of PP in the quantum setting. PP is typically defined by referring to Turing machines or randomized algorithms. These definitions are not very useful when dealing with quantum computing. However, it turns out that PP has an equivalent formulation with much more obvious quantum applications, known as PostBQP [Aar05]. PostBQP refers to the class of problems efficiently solvable by uniform quantum circuits with the additional ability to *postselect*. Postselection is another word for performing conditional sampling, and in the quantum setting refers to the ability to choose the result of a measurement and acquire the corresponding residual state. Thus, to break any OWSG with a PP oracle, it suffices to define an algorithm which breaks the OWSG given oracle access to some language in PostBQP.

One-way puzzles. Instead of breaking OWSGs with a PostBQP oracle directly, we rely on a recent result which constructs an interesting classical output primitive, a one-way puzzle, from any OWSG [KT23]. Formally, a one-way puzzle is a pair of algorithms $(\text{Samp}, \text{Ver})$ where Samp samples a key-puzzle pair (k, s) such that $\text{Ver}(k, s)$ outputs 1 with overwhelming probability. Samp is required to be an efficient quantum algorithm, and Ver is allowed to be any arbitrary function. The security requirement is that given s , it is hard for an adversary to find a k' such that $\text{Ver}(k', s) = 1$. Morally, a one-way puzzle is a one-way function where the input and output are sampled together.

It is clear that any one-way puzzle $(\text{Samp}, \text{Ver})$ can be broken by an adversary with the ability to postselect. Given a puzzle s , the adversary can simply run Samp up until it would measure the output state, and then postselect on the output puzzle being s . Measuring the output key will then give a k' which with high probability will satisfy $\text{Ver}(s, k') = 1$. This attack can be viewed as a search version of the attack of [Kre21] against any λ -output PRS.

Attack with decision oracle. But note that this attack makes use of postselecting directly. It is unclear how to translate this into an attack with a decision oracle for PP. To solve this, we show that it is possible in general to perform conditional sampling given access to a PP oracle:

Lemma 1.6 (Informal version of Lemma 4.2). *Let Samp be a (uniform) quantum polynomial time algorithm such that $\text{Samp}(1^n)$ outputs a pair of classical strings (k, s) . There exists a poly-time quantum algorithm \mathcal{A} and a PP language \mathcal{L} such that $\mathcal{A}^{\mathcal{L}}$ takes as input s' and outputs k' , and whose distribution has total variation distance at most $1/n$ from the distribution $(\text{Samp} | s')^{\text{key}}$ defined by*

$$\Pr[(\text{Samp} | s')^{\text{key}} \rightarrow k'] = \Pr_{\text{Samp}(1^n) \rightarrow (k, s)} [k = k' | s = s'].$$

(In other words, we denote by $(\text{Samp} | s')^{\text{key}}$ the distribution of keys generated by $\text{Samp}(1^n)$ conditioned on the puzzle being equal to s' .)

This lemma is in some sense a “search-to-decision” style argument for PP. The argument goes along the same lines as the search-to-decision reduction for SAT. Note that with a PostBQP oracle, we can test whether or not it is possible for an algorithm Samp to produce any given output x by simply postselecting on x being produced by Samp .

A naive approach to sampling k' from $(\text{Samp} | s')^{\text{key}}$ is to sample k' bit by bit. It is possible with a PP oracle to check whether any given output is in the range of Samp . Thus, we can begin by checking whether $(1, s')$ is in the range of Samp with all but the first bit of the key discarded. If so, we can set the first bit of k' to be 1, otherwise 0. In the next step, we can check whether $(k'_1 \circ 1, s')$ is in the range of Samp with all but the second bit of the key discarded. If so, we can set the second bit of k' to be 1, otherwise 0. Repeating this process for each bit of the key will uniformly select an output k' from the range of $(\text{Samp} | s')^{\text{key}}$.

However, $(\text{Samp} | s')^{\text{key}}$ is not necessarily a flat distribution, and so the resulting distribution on k' may be very different from $(\text{Samp} | s')^{\text{key}}$. However, this can be resolved by noticing a key fact. Using a PP oracle, it is possible to *estimate the probability* that the first bit of the output of Samp is 1 conditioned on the output puzzle being s' . Thus, we can use the same technique as before, but instead of just setting each bit of k' , we can sample each bit according to our approximation of the correct conditional distribution. Although the error will add up, we can set our initial error to be small enough that the distribution over k' will be sufficiently close to $(\text{Samp} | s')^{\text{key}}$.

Applying Lemma 1.6 to our postselecting attack against one-way puzzles gives us an efficient quantum attack against one-way puzzles using a PP oracle. As one-way puzzles can be built from one-way state generators, this immediately implies Theorem 1.5.

More on search-to-decision reductions using a PP oracle. The ability of a postselection oracle to aid in search-to-decision reductions was first noted by [INN⁺22], where it was shown that a quantum poly-time algorithm can find a QMA witness by making one quantum query to a PP oracle. They use very different techniques, and in fact our algorithm requires many classical queries instead of one quantum query.

1.3 Conclusion and Future work

We showed two new results about one-way state generators. We proved that their existence is implied by almost all computational PRSs, and that $O(n)$ -copy OWSGs can be broken by QPT algorithms with access to a PP oracle. These results bring the cartography of OWSGs much more in line with that of PRSs.

Further research directions. We outline two major lines of research which are left open by this paper.

Separating OWSG and quantum bit commitments. The two recent papers [KT23] and [LMW23] make clear that proving that quantum bit commitments exist relative to a random oracle and any classical oracle (with access to the random oracle) would suffice to show a black box separation between OWSGs and quantum bit commitments. It follows from our results that a black box separation between quantum bit commitments and OWSGs can be achieved by proving a weaker statement, namely that there exists an oracle \mathcal{O} relative to which quantum bit commitments exist and $\text{PP}^{\mathcal{O}} \subseteq \text{BQP}^{\mathcal{O}}$.

Since it is now known from [KT23] that OWSGs imply quantum bit commitments, proving this conjecture and the accompanying separation would provide strong evidence that quantum bit commitments are (at least among the quantum cryptographic primitives so far proposed) the minimal assumption for quantum cryptography.

OWSG \rightarrow PRS. We have now shown that PRSs imply OWSGs in nearly the greatest possible generality, and that the existence of either have equivalent known complexity ramifications. This gives additional motivation to the question of whether these primitives are equivalent (as their classical equivalents are known to be). Proving either an equivalence or separation (even for a limited copy setting) would be an exciting accomplishment that would greatly clarify the cartography of quantum cryptography.

Open questions. There are two open questions which we would like to see an answer for.

Optimal PRS \rightarrow OWSG Reduction. OWSGs are inherently computational objects, meaning they cannot be implied by statistical primitives. Currently we know that

- PRSs with $m \geq \log n$ output qubits require computational assumptions.
- There exists a $c \in (0, 1)$ s.t. statistical PRSs with $m \leq c \log n$ output qubits exist. So the existence of PRSs in this regime cannot imply OWSGs.
- PRS with $m \geq \log n + 1$ output qubits imply OWSGs.

This leaves the question open of whether a PRS with $m \in (c \log n, \log n]$ output qubits imply the existence of OWSGs. The smallest improvement to this space would be to show that a PRS with $m = \log n$ output qubits imply OWSGs. Other possible tightening would be to show that, for any $c < 1$, a statistical PRS with $m \leq c \log n$ exist, though this would show the surprising result that just over 1 bit of randomness per amplitude is enough to achieve statistical closeness to Haar-random states.

Parameter regimes for which OWSGs don't exist. It would be very interesting to know which parameter regimes it can be shown OWSGs do not exist in. We conjecture that much like how $O(\log n)$ output classical OWFs cannot exist, $m = O(\log \log n)$ output qubit OWSGs cannot exist unconditionally. It would be good to have a complete understanding of how OWSGs behave in different parameter regimes, and a proof of its non-existence in that setting would help to clarify this picture.

2 Preliminaries

In this section, we introduce basic notation and recall definitions from the literature that will be used throughout the rest of this work.

2.1 Computational complexity

We refer the reader to [AB09] for the definition of standard complexity classes such as BQP and PP. We will often abbreviate quantum polynomial-time algorithms by the acronym ‘QPT’.

We recall the definition of the PostBQP complexity class (see, e.g., [Kre21, Definition 12]). We will only be concerned with its promise version in this work.

Definition 2.1 (PromisePostBQP). *A promise problem $\Pi : \{0, 1\}^* \rightarrow \{0, 1, \perp\}$ is in PromisePostBQP (Postselected Bounded-error Quantum Polynomial time) if there exists a QPT algorithm \mathcal{A} whose output is in $\{0, 1, *\}$, and which is such that:*

- (i) *For all $x \in \{0, 1\}^*$, we have $\Pr[\mathcal{A}(x) \in \{0, 1\}] > 0$. When $\mathcal{A}(x) \in \{0, 1\}$, we say that postselection succeeds.*
- (ii) *If $\Pi(x) = 1$, then $\Pr[\mathcal{A}(x) = 1 \mid \mathcal{A}(x) \in \{0, 1\}] \geq \frac{2}{3}$. In other words, conditioned on postselection succeeding, \mathcal{A} outputs 1 with at least $\frac{2}{3}$ probability.*
- (iii) *If $\Pi(x) = 0$, then $\Pr[\mathcal{A}(x) = 0 \mid \mathcal{A}(x) \in \{0, 1\}] \geq \frac{2}{3}$. In other words, conditioned on postselection succeeding, \mathcal{A} outputs 1 with at most $\frac{1}{3}$ probability.*

We remark that the definition of PostBQP is sensitive to the choice of the gate-set, as noticed by [Kup15]. We remark that [Kup15] also proves that every gate-set that satisfies two “reasonable” conditions gives rise to an equivalent “canonical” PostBQP class. Throughout the paper we assume that we are dealing with one such gate-set, such as the $\{\text{CNOT}, H, T\}$ -gate-set. We refer the reader to [Kup15, Section 2.5] for a detailed technical discussion of this matter.

We also recall the following result of Aaronson [Aar05], which states that $\text{PP} = \text{PostBQP}$, remarking that his result also holds for the corresponding promise classes.

Lemma 2.2 (Aaronson [Aar05]). $\text{PromisePostBQP} = \text{PromisePP}$.

Since PP is a syntactic class, we can extend any promise problem in PromisePP into a language in PP. This remark will be fundamental in Section 4, as it was also in [Kre21], when we build a PP oracle with which we can break OWSGs.

2.2 Quantum information theory and cryptography

We denote by $\mathbb{S}(N)$ the set of N -dimensional pure quantum states. We identify the set of n -qubit pure states with $\mathbb{S}(2^n)$. We denote by $\text{Haar}(n)$ the Haar measure on $\mathbb{S}(2^n)$.

Quantum cryptography. We start by introducing the notion of pseudorandom state generators (PRS) [JLS18]. Roughly speaking, given a classical key k , a PRS maps k to a quantum pure state $|\phi_k\rangle$. The security guarantee is that the output of a PRS on a random input should look like a random state. That is, it is hard for any quantum adversary to distinguish a random $|\phi_k\rangle$ from a Haar random state.

Definition 2.3 (Pseudorandom States, Definition 2 of [JLS18]). *Let λ be the security parameter and let \mathcal{K} be a set of binary strings referred to as the key space. Let G be a QPT algorithm that on input $k \in \mathcal{K}$ outputs a pure quantum state $|\phi_k\rangle$ over $n = n(\lambda)$ qubits. We say G is (t, ε) -pseudorandom if the distribution over outputs is ε -indistinguishable from Haar random given t copies. In other words, for any QPT adversary \mathcal{A} , we have*

$$\left| \Pr_{k \leftarrow \mathcal{K}} [\mathcal{A}(|\phi_k\rangle^{\otimes t}) = 1] - \Pr_{|\psi\rangle \leftarrow \text{Haar}(n)} [\mathcal{A}(|\psi\rangle^{\otimes t}) = 1] \right| \leq \varepsilon.$$

We say that G is pseudorandom if it is $(\lambda^c, \frac{1}{\lambda^c})$ -pseudorandom for all $c > 0$, and that it is t -pseudorandom or a t -copy PRS if it is $(t, \frac{1}{\lambda^c})$ -pseudorandom for all $c > 0$.

We turn to introducing the notion of one-way state generators (OWSGs) [MY22a]. A OWSG maps a classical key k to quantum state $|\phi_k\rangle$. The security guarantee of a OWSG is that, given any polynomial number of copies of $|\phi_k\rangle$, it is hard for a quantum algorithm to find keys k' such $|\phi_k\rangle, |\phi_{k'}\rangle$ have noticeable overlap. OWSGs can also be defined to have mixed state outputs [MY22b], although we will not consider this variant in this work.

Definition 2.4 (One-Way State Generators, Definition 4.1 of [MY22b]). *Let λ be the security parameter and let \mathcal{K} be a set of binary strings referred to as the key space. Let G be a QPT algorithm that on input $k \in \mathcal{K}$ outputs a pure quantum state $|\phi_k\rangle$. We say G is (t, ε) -one-way if the outputs are hard to invert with accuracy at least ε . In other words, for any QPT adversary \mathcal{A} , we have*

$$\mathbb{E}_{\substack{k \leftarrow \mathcal{K} \\ k' \leftarrow \mathcal{A}(|\phi_k\rangle^{\otimes t})}} \left[|\langle \phi_k | \phi_{k'} \rangle|^2 \right] \leq \varepsilon.$$

We say that G is one-way or strongly one way if it is $(\lambda^c, \frac{1}{\lambda^c})$ -one-way for all $c > 0$, and that it is t -one-way, a t -copy OWSG, or a t -copy strong OWSG if it is $(t, \frac{1}{\lambda^c})$ -one-way for all $c > 0$.

We remark that all of the results in this paper will also hold for the more general definition of (pure-state) one-way state generators that was introduced in the later work of Morimae and

Yamakawa [MY22a], where the one-way state generator is allowed to have a separate quantum key generation procedure.²

We will rely on the notion of one-way puzzles, defined in [KT23]. A one-way puzzle is a pair of algorithms $(\text{Samp}, \text{Ver})$ where Samp samples a key-puzzle pair (k, s) such that $\text{Ver}(k, s)$ outputs 1 with overwhelming probability. Samp is required to be an efficient quantum algorithm, and Ver is allowed to be any arbitrary function. The security requirement is that given s , it is hard for an adversary to find a k' such that $\text{Ver}(k', s) = 1$.

Definition 2.5 (One-Way Puzzles [KT23]). *Let λ be the security parameter. A one-way puzzle is a pair of sampling and verification algorithms $(\text{Samp}, \text{Ver})$ with the following syntax.*

- $\text{Samp}(1^\lambda) \rightarrow (k, s)$ is a (uniform) quantum polynomial time algorithm that outputs a pair of classical strings (k, s) . We refer to s as the puzzle and k as its key. Without loss of generality we may assume that $k \in \{0, 1\}^\lambda$.
- $\text{Ver}(k, s) \rightarrow \top$ or \perp , is an unbounded algorithm that on input any pair of classical strings (k, s) halts and outputs either \top or \perp .

These satisfy the following properties.

- **Correctness.** *Outputs of the sampler pass verification with overwhelming probability, i.e.,*

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)}[\text{Ver}(k, s) = \top] = 1 - \text{negl}(\lambda).$$

- **Security.** *Given s , it is (quantumly) computationally infeasible to find k satisfying $\text{Ver}(k, s) = \top$, i.e., for every polynomial-sized quantum circuit \mathcal{A} ,*

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^\lambda)}[\text{Ver}(\mathcal{A}(s), s) = \top] = \text{negl}(\lambda).$$

An important recent result due to Khurana and Tomer [KT23] states that one-way state generators imply one-way puzzles. This will be crucial for our algorithm in Section 4.

Theorem 2.6 ([KT23, Theorem 4.2]). *If there exists a $(O(n), \text{negl}(n))$ -OWSG, then there exists a one-way puzzle.*

2.3 Probability distributions

We recall a few standard probability distributions, the chi-squared distribution and the Fisher-Snedecor distribution, the latter of which will be useful when trying to show that compressing PRSs are one-way state generators Lemma 3.1.

Definition 2.7 (Chi-squared distribution). *Let X_1, \dots, X_k be independent normal random variables with mean 0 and variance 1. The chi-square distribution with k degrees of freedom, denoted by $\chi^2(k)$, is the probability distribution of the sum $\sum_{i=1}^k X_i^2$.*

²The results we show in Section 3 imply the existence of one-way state generators in the sense of Definition 2.4 in various settings. Since these are more restricted objects than the ones considered in the more general definition of [MY22a], our results will also imply their existence in the same settings and with an analogous parameter regime. Furthermore, since our oracle QPT algorithm in Section 4 comes from an algorithm breaking one-way puzzles, and one-way puzzles were proved in [KT23] to follow from the “general” pure state OWGs of [MY22a], our algorithm will therefore also break OWGs with a quantum key generation procedure.

Definition 2.8 (Fisher-Snedecor distribution). *Let $a, b \in \mathbb{N}$. The Fisher-Snedecor distribution with a and b degrees of freedom, denoted by $F(a, b)$, is the distribution given by the following ratio:*

$$F(a, b) \sim \frac{\chi^2(a)/a}{\chi^2(b)/b} = \frac{b}{a} \cdot \frac{\chi^2(a)}{\chi^2(b)}.$$

We now recall the definition of the beta function and its variations in order to express the cumulative distribution function of the F -distribution. The *beta function* is defined as

$$B(a, b) := \int_0^1 x^{a-1}(1-x)^{b-1} dx.$$

The *incomplete beta function* is defined as

$$B(t; a, b) := \int_0^t x^{a-1}(1-x)^{b-1} dx.$$

The *regularised incomplete beta function* is defined as

$$I_x(a, b) := \frac{B(x; a, b)}{B(a, b)}.$$

Lemma 2.9 (Cumulative distribution function of the F -distribution). *The cumulative distribution function $p_{a,b}(t)$ of the $F(a, b)$ distribution satisfies*

$$p_{a,b}(t) = I_{at/(at+b)}(a/2, b/2).$$

In particular, we have

$$\Pr_{\mathbf{Y} \leftarrow F(a,b)}[\mathbf{Y} \leq \theta] = \frac{\int_0^{(a\theta)/(a\theta+b)} x^{a-1}(1-x)^{b-1} dx}{\int_0^1 x^{a-1}(1-x)^{b-1} dx}.$$

2.4 Approximate t -designs

In our construction of unconditional t -copy OWSGs, we will use approximate t -designs. Informally, an approximate t -design is a distribution over states such that t copies of an output state is statistically close to t copies of a Haar random state. One can think of t -designs as a quantum version of a t -wise independent distribution, or as a t -copy PRS with a statistical security guarantee.

Definition 2.10 (Approximate t -Design, Definition 2.2 of [OSP23], rephrased). *A probability distribution S over $\mathbb{S}(2^n)$ is an ε -approximate t -design if*

$$\left\| \mathbb{E}_{|\psi\rangle \leftarrow \text{Haar}(n)} [|\psi\rangle\langle\psi|^{\otimes t}] - \mathbb{E}_{|\psi\rangle \leftarrow S} [|\psi\rangle\langle\psi|^{\otimes t}] \right\|_1 \leq \varepsilon,$$

where $\|\cdot\|_1$ is the Schatten 1-norm (or trace norm). We call G an efficient ε -approximate t -design if G is a quantum algorithm running in time $\text{poly}(n, m, t, \log(1/\varepsilon))$ which maps classical strings in $\{0, 1\}^n$ to quantum states over $\mathbb{S}(2^m)$, and such that the output distribution of $G(\cdot)$ on a random n -bit string forms an ε -approximate t -design.

Recently, it has been shown that approximate t -designs with almost optimal seed exist unconditionally.

Theorem 2.11 (Theorem 1.1 of [OSP23], rephrased). *For all $m, t, \varepsilon > 0$, there exists an efficient ε -approximate t -design with input size $n = O(mt + \log(1/\varepsilon))$.*

3 One-way state generators from compressing pseudorandom states

We construct both weak and strong t -copy one-way state generators from $(t + 1)$ -copy compressing pseudorandom states. Before this the only known reduction worked for expanding PRS mapping n -bit strings to cn -qubit states [MY22a] for some $c > 1$. To generalize this reduction, we rely on the following concentration inequality, which informally states that, with high probability, any fixed state is unlikely to be close to Haar-random states. We will then use this in Theorem 3.2 to bound how well a OWSG inverter can distinguish outputs of a PRS from Haar-random states.

Lemma 3.1 (Concentration of Haar States³). *Let $|\phi_0\rangle$ be any state of dimension $N = 2^n$. Then, for any $s > 0$, we have*

$$\Pr_{|\psi\rangle \sim \text{Haar}(n)} \left[|\langle \phi_0 | \psi \rangle|^2 \geq \frac{1}{s} \right] \leq \left(\frac{s}{s+1} \right)^{N-1}.$$

Proof. Since the Haar distribution is invariant under unitary transformations, without loss of generality we assume $|\phi_0\rangle = |0\rangle$.

Muller [Mul59] showed that sampling from the Haar distribution is equivalent to sampling $|\psi\rangle$ as

$$|\psi\rangle \propto \sum_x (\alpha_x + \beta_x i) |x\rangle,$$

where each α_x, β_x is sampled according to the standard Gaussian with expectation 0 and standard deviation 1. Then, we define the random variable Y as

$$Y := \frac{\alpha_0^2 + \beta_0^2}{\sum_{x \neq 0} \alpha_x^2 + \beta_x^2}.$$

Expanding out the inner product gives us

$$|\langle 0 | \psi \rangle|^2 = \frac{\alpha_0^2 + \beta_0^2}{\sum_x (\alpha_x^2 + \beta_x^2)} \leq Y.$$

But observe that each α_x, β_x is sampled independently, and so Y is distributed as the ratio of two chi-squared random variables. So, we see that Y is sampled as a (scaled) F -distribution as follows:

$$Y \sim \frac{\chi^2(2)}{\chi^2(2N-2)} \sim \frac{2}{2N-2} \cdot F(2, 2N-2).$$

³In [Kre21] the author refers to a closely related inequality as following from “standard concentration inequalities, or even an explicit computation”. However we were unable to find an actual proof of the inequality in either the citation listed [BHH16], or in the paper which it cites [HLW06]. Consequently we decided to include a proof here for completeness.

We conclude by Lemma 2.9 that

$$\begin{aligned}
\Pr \left[Y \geq \frac{1}{s} \right] &= \Pr \left[F(2, 2N - 2) \geq \frac{N - 1}{s} \right] \\
&= 1 - \frac{\int_0^{\frac{1}{s+1}} (1-x)^{N-2} dx}{\int_0^1 (1-x)^{N-2} dx} \\
&= 1 - \frac{\left[-\frac{1}{N-1} \cdot (1-x)^{N-1} \right]_0^{1/(s+1)}}{\left[-\frac{1}{N-1} \cdot (1-x)^{N-1} \right]_0^1} \\
&= 1 - \frac{\left[(1-x)^{N-1} \right]_0^{1/(s+1)}}{\left[(1-x)^{N-1} \right]_0^1} \\
&= 1 + \left(\left(1 - \frac{1}{s+1} \right)^{N-1} - 1 \right) \\
&= \left(\frac{s}{s+1} \right)^{N-1},
\end{aligned}$$

where $[f(x)]_0^1 = f(1) - f(0)$. Since

$$\Pr_{|\psi\rangle \sim \text{Haar}(N)} \left[|\langle \phi_0 | \psi \rangle|^2 \geq \frac{1}{s} \right] \leq \Pr \left[Y \geq \frac{1}{s} \right],$$

we are done. \square

Using this lemma we first show a general result stating that state generators that are pseudo-random must also be one-way. We then apply the result in two different parameter regimes. While [Kre21] only claims that PRSs with $m = n$ output bits can be broken by PP, [AGQY22] calls out that the proof can be extended to the case when $m \geq \log n + c$. For this regime, we show that PRSs are weak OWSGs. We then show that slightly less compressing PRSs ($m = 2 \log n + 1$) are strong OWSGs. In all these three results, the number of copies falls by 1 moving from a PRS to a OWSG. While we end up focusing on sublinear-copy PRSs in the next subsection, these reductions work in the default many-copy setting considered in most papers on quantum cryptographic primitives.

Lemma 3.2. *For all $f(n)$ and for*

$$\delta = 2^n \cdot \left(\frac{f(n)}{f(n) + 1} \right)^{(2^m - 1)} + \frac{1}{f(n)},$$

if $G : k \mapsto |\phi_k\rangle$ is a state generator taking n -bit strings to m -qubit pure states which is $(t + 1, \varepsilon)$ -pseudorandom, then it is also $(t, \varepsilon + \delta)$ -one way.

Proof. For the sake of contradiction, assume that there exists an adversary \mathcal{A} that can succeed with probability larger than $\varepsilon + \delta$ in the t -copy one-wayness game (Definition 2.4). We will construct a new adversary \mathcal{A}' for the pseudorandomness game as follows:

Algorithm 1 Adversary \mathcal{A}' in the pseudorandomness game

Input: $|\psi\rangle^{\otimes(t+1)}$.

Output: 0 if $|\psi\rangle$ is pseudorandom, 1 if $|\psi\rangle$ is Haar random.

- 1: Run \mathcal{A} on the first t copies and obtain $k' \leftarrow \mathcal{A}(|\psi\rangle^{\otimes t})$
 - 2: Measure the last copy of $|\psi\rangle$ in the basis $\{|\phi_{k'}\rangle\langle\phi_{k'}|, I - |\phi_{k'}\rangle\langle\phi_{k'}|\}$
 - 3: **return** 1 if the result is $|\phi_{k'}\rangle\langle\phi_{k'}|$, else **return** 0.
-

Observe that when \mathcal{A}' is given a pseudorandom input, it outputs 1 with probability at least

$$\Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}' \left(|\phi_k\rangle^{\otimes(t+1)} \right) = 1 \right] = \mathbb{E}_{k \leftarrow \mathcal{K}} \left[|\langle\phi_k|\phi_{k'}\rangle|^2 \mid k' \leftarrow \mathcal{A} \left(|\phi_k\rangle^{\otimes t} \right) \right] > \varepsilon + \delta,$$

where the key space \mathcal{K} is equal to $\{0, 1\}^n$ here. Thus, it suffices remains to bound the probability that \mathcal{A}' detects a Haar random state:

$$\Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\mathcal{A}' \left(|\psi\rangle^{\otimes(t+1)} \right) = 1 \right] \leq \delta.$$

By construction, we have

$$\begin{aligned} \Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\mathcal{A}' \left(|\psi\rangle^{\otimes(t+1)} \right) = 1 \right] &= \mathbb{E}_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[|\langle\psi|\phi_k\rangle|^2 \mid k \leftarrow \mathcal{A} \left(|\psi\rangle^{\otimes t} \right) \right] \\ &= \int_{\mathbb{S}(2^m)} d\mu(\psi) \cdot \sum_{k \in \{0,1\}^n} \Pr \left[k \leftarrow \mathcal{A} \left(|\psi\rangle^{\otimes t} \right) \right] \cdot |\langle\psi|\phi_k\rangle|^2 \\ &\leq \int_{\mathbb{S}(2^m)} d\mu(\psi) \cdot \max_{k \in \{0,1\}^n} |\langle\psi|\phi_k\rangle|^2. \end{aligned}$$

Where $\mu(\psi)$ is the Haar measure on the space of m -qubit pure states. We will now partition the set of m -qubit pure states into the states that are ‘close’ to a pseudorandom state $|\phi_k\rangle$, and the states that are ‘far’ from all pseudorandom states. Formally, we define

$$A_f := \left\{ |\psi\rangle \in \mathbb{S}(2^m) \mid \max_k |\langle\psi|\phi_k\rangle|^2 \geq \frac{1}{f(n)} \right\}.$$

The set of states that are ‘far’ from all pseudorandom states is its complement,

$$B_f := \left\{ |\psi\rangle \in \mathbb{S}(2^m) \mid \max_k |\langle\psi|\phi_k\rangle|^2 < \frac{1}{f(n)} \right\}.$$

We proceed with computing the integral separately for the two sets:

$$\begin{aligned}
\int_{\mathbb{S}(2^m)} d\mu(\psi) \cdot \max_{k \in \{0,1\}^n} |\langle \psi | \phi_k \rangle|^2 &= \int_{A_f} d\mu(\psi) \cdot \max_{k \in \{0,1\}^n} |\langle \psi | \phi_k \rangle|^2 \\
&\quad + \int_{B_f} d\mu(\psi) \cdot \max_{k \in \{0,1\}^n} |\langle \psi | \phi_k \rangle|^2 \\
&\leq \int_{A_f} d\mu(\psi) + \max_{\substack{|\psi\rangle \in B_f \\ k \in \{0,1\}^n}} |\langle \psi | \phi_k \rangle|^2 \\
&< \Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\exists k \left| |\langle \psi | \phi_k \rangle|^2 \geq \frac{1}{f(n)} \right. \right] + \frac{1}{f(n)} \\
&\leq \sum_{k \in \mathcal{K}} \Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[|\langle \psi | \phi_k \rangle|^2 \geq \frac{1}{f(n)} \right] + \frac{1}{f(n)}.
\end{aligned}$$

Lemma 3.1 implies that

$$\Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\mathcal{A}' \left(|\psi\rangle^{\otimes(t+1)} \right) = 1 \right] \leq 2^n \cdot \left(\frac{f(n)}{1+f(n)} \right)^{2^m-1} + \frac{1}{f(n)} = \delta.$$

We conclude

$$\left| \Pr_{k \leftarrow \mathcal{K}} \left[\mathcal{A}' \left(|\psi\rangle^{\otimes(t+1)} \right) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\mathcal{A}' \left(|\psi\rangle^{\otimes(t+1)} \right) = 1 \right] \right| > \varepsilon. \quad \square$$

Using this general result we can specify to the two theorems below.

Theorem 3.3 (PRs imply OWsGs). *If G is a state generator taking n -bit strings to $m > \log n + c$ qubit states (with $c \geq 1$) which is $(t+1, \varepsilon)$ -pseudorandom, then G is also $(t, \varepsilon + 3/4)$ -one-way.*

Proof. Take Lemma 3.2 with $f(n) = 2$, and $m > \log n + c$. We get that

$$\begin{aligned}
\delta &\leq 2^n \cdot \left(\frac{2}{3} \right)^{n \cdot 2^c - 1} + \frac{1}{2} \\
&= \frac{3}{2} \cdot \left(\frac{2^{2^c} + 1}{3^{2^c}} \right)^n + \frac{1}{2}.
\end{aligned}$$

Since $c > 1$, the left term becomes less than $1/4$ for sufficiently large n , and thus $\delta < 3/4$. \square

Using amplification arguments from [MY22b], these $(t, \varepsilon + 3/4)$ -OWsGs can be used to construct strong OWsGs with negligible probability of inversion.

This reduction is nearly optimal. [BCQ22] showed that there exist statistical many-copy PRs with $c \log n$ output bits for some $c < 1$. As shown later strong OWsGs must imply computational assumptions (e.g. $\text{BQP} \neq \text{PP}$) so they cannot be implied by statistical or information theoretic primitives. Consequently there can be no proof that PRs with output shorter than $c \log n$ qubits imply OWsGs, meaning our proof is optimal up to multiplicative constant factors.

We can also show that PRs with superlogarithmic output are themselves strong OWsGs.

Theorem 3.4 (PRs are strong OWsGs). *If G is a state generator taking n -bit strings to $m = \omega(\log n)$ -qubit states which is $(t+1, \varepsilon)$ -pseudorandom, then it is also $(t, \varepsilon + \text{negl}(n))$ -one way.*

Proof. Take Lemma 3.2 and set $f(n) = \frac{2^m - 1}{2^{m/2}} - 1$. Note that $f(n) = 2^{\Omega(m)} = n^{\omega(1)}$. We get

$$\begin{aligned}
\delta &= 2^n \cdot \left(\frac{f(n)}{f(n) + 1} \right)^{2^m - 1} + \frac{1}{f(n)} \\
&\leq 2^n \cdot \left(1 - \frac{1}{f(n) + 1} \right)^{2^m - 1} + \text{negl}(n) \\
&\leq 2^n \cdot \left(\frac{1}{e} \right)^{2^{m/2}} + \text{negl}(n) \\
&= 2^{n - \Omega(n^{\omega(1)})} + \text{negl}(n) \\
&= \text{negl}(n). \quad \square
\end{aligned}$$

3.1 Unconditional OWSGs from efficient approximate t -designs

We first observe that any efficient $2^{-\lambda}$ -approximate t -design is also a t -copy PRS. This follows definitionally, since t copies from a t -design are statistically close to t copies from a Haar random state.

Proposition 3.5. *Let G be an efficient $2^{-\lambda}$ -approximate t -design that maps n -bit strings to m -qubit pure states. Denote $|\psi_k\rangle := G(k)$. Then for any QPT adversary \mathcal{A} ,*

$$\left| \Pr_{k \leftarrow \{0,1\}^n} \left[\mathcal{A} \left(|\psi_k\rangle^{\otimes t} \right) = 1 \right] - \Pr_{|\psi\rangle \leftarrow \text{Haar}(m)} \left[\mathcal{A} \left(|\psi\rangle^{\otimes t} \right) = 1 \right] \right| \leq \text{negl}(\lambda).$$

In particular, the proposition holds definitionally even for *statistical* adversaries, and so therefore must also hold for QPT adversaries. A simple corollary of Theorem 3.4 then shows that efficient approximate t -designs are also OWSGs.

Corollary 3.6. *Let G be an efficient ε -approximate t -design mapping n bits to $m = \omega(\log n)$ qubits. Then G is $(t, \varepsilon + \text{negl}(n))$ -one-way.*

Theorem 2.11 with $\varepsilon = 2^{-\lambda}$ gives that, for any polynomial t , $t(\lambda)$ -copy PRSs exist unconditionally, and thus Theorem 3.4 concludes that $t(\lambda)$ -copy OWSGs also exist unconditionally.

Contrast this with the recent result of Khurana and Tomer [KT23], where they show that $\Theta(n)$ -copy OWSGs can be used to build quantum bit commitments (and thus require computational hardness [LC97]). This raises the question of what is the largest number of copies t (relative to the number of classical input bits n) for which OWSGs exist unconditionally. We show that the efficient approximate t -designs of [OSP23] approach this computational threshold up to a logarithmic factor.

Corollary 3.7. *For every function $\alpha = \alpha(n) = \omega(1)$, there exists a $\Theta\left(\frac{n}{\alpha \log n}\right)$ -copy OWSG.*

Proof. From Theorem 2.11, we know that there exists some positive constant c such that, for $n = c \cdot mt + c \log(1/\varepsilon)$, there is an efficient ε -approximate t -design mapping n bits to m qubits.

Setting $\varepsilon = 2^{-\lambda}$, $n = 2c \cdot \lambda$, $t = \lambda/(\alpha \log n)$ and $m = \alpha \cdot \log n$ gives us

$$n = 2c\lambda = cat \log n + c\lambda = cmt + c \log(1/\varepsilon),$$

and thus we can build efficient ε -approximate $\Theta(n/(\alpha \log n))$ -designs. But since $m = \omega(\log n)$, using such a design also gives a $\Theta\left(\frac{n}{\alpha \log n}\right)$ -copy strong OWSG from Corollary 3.6. \square

4 Breaking one-way state generators with a PP oracle

In this section, we show how to break one-way state generators (OWSGs) with a PP oracle. Since one-way state generators imply one-way puzzles (Theorem 2.6, due to [KT23]), it suffices to show how to break one-way puzzles using a PP oracle. Recall that a one-way puzzle is a pair of sampling and verification quantum algorithms (Samp, Ver) (see Definition 2.5). The algorithm $\text{Samp}(1^n)$ outputs a pair (k, s) , where k is referred to as the key and s as the puzzle. To break the one-way puzzle, it suffices to create a quantum algorithm \mathcal{A} that, given a puzzle s , returns a key k' such that $\text{Ver}(k', s)$ is accepted with non-negligible probability.

Our strategy to construct \mathcal{A} is as follows. Given a puzzle s sampled according to Samp , we will sample a key k according to the conditional distribution of keys that are sampled together with s . This will suffice to break the one-way puzzle, as Ver accepts pairs from Samp with $(1 - \text{negl})$ -probability. To sample from this distribution, we first show how to use the PP oracle to estimate the conditional probability, and finally how to sample according to a distribution that is close to the true conditional distribution.

We view our strategy as inspired by Kretschmer's [Kre21] idea for breaking pseudorandom state generators with a PP oracle, together with a search-to-decision reduction for PP.

Lemma 4.1. *Let S be a (uniform) quantum polynomial time algorithm that outputs n bits, denoted as a pair of a string x , and a bit b . There exists a poly-time quantum algorithm \mathcal{A} and a PP language \mathcal{L} such that $\mathcal{A}^{\mathcal{L}}$ can estimate the distribution of bit b , conditioned on the output string x . Formally,*

$$p_{x,b} - \frac{1}{n^2} \leq \mathcal{A}^{\mathcal{L}}(S, 1^n, x, b) \leq p_{x,b} + \frac{1}{n^2}$$

where $p_{x,b} = \Pr[S(1^n) = (x, b) | S(1^n) \in \{(x, 0), (x, 1)\}]$.

Proof. Our quantum algorithm \mathcal{A} will take as input the algorithm S , a unary description of the length 1^n , the outputs (x, b) , and will estimate the conditional probability of b , conditioned on the first output of S being x .

Definition of the PP language. We will describe the PP language \mathcal{L} in terms of a PromisePostBQP algorithm $\mathcal{B}(S, x, t)$. This will define a promise problem which is computed by $\mathcal{B}(S, x, t)$. By the equivalence $\text{PromisePostBQP} = \text{PromisePP}$, this gives us a promise problem in PromisePP. Since PromisePP is a syntactic class, this can be extended to a language $\mathcal{L} \in \text{PP}$. However, later on our algorithm \mathcal{A} will only depend on the behaviour of the oracle to \mathcal{L} on inputs that satisfy the promise condition of the promise problem defined by $\mathcal{B}(S, x, t)$.

We first assume without loss of generality that any measurements in S are delayed until the very end of the algorithm. The algorithm \mathcal{B} will simulate S and postselect on the output measurements matching x . Conditioned on postselection succeeding, the output register for b will contain the pure state $\sqrt{p_{x,0}}|0\rangle + \sqrt{p_{x,1}}|1\rangle$. Then \mathcal{B} measures the register of b , and repeats this procedure $r = \Theta(n^4)$ times, outputting 1 if the measurements yields 1 at least a $\frac{t}{2n^2}$ fraction of the time. This ends the description of the PromisePostBQP algorithm $\mathcal{B}(S, x, t)$.

Which inputs (S, x, t) are accepted by \mathcal{B} ? Let $\mathbf{b}_1, \dots, \mathbf{b}_r$ denote the random variables that correspond to the value of the b register at each simulation of S by \mathcal{B} . We define $\mathbf{B} = \sum_i \mathbf{b}_i$ to be their sum. Note that \mathbf{B}/r is the fraction of measurements that yield a 1 in the algorithm \mathcal{B} ,

and recall that we accept if this fraction is at least $t/(2n^2)$. By standard concentration inequalities (e.g., Chebyshev's), when the number of repetitions r is at least $\Theta(n^4)$, then

$$\Pr_{\mathbf{b}_1, \dots, \mathbf{b}_r} \left[\left| \frac{\mathbf{B}}{r} - p_{x,1} \right| > \frac{1}{4n^2} \right] < \frac{1}{3}.$$

This implies that, if $p_{x,1} \geq \frac{t}{2n^2} + \frac{1}{4n^2}$, then $\frac{\mathbf{B}}{r} > \frac{t}{2n^2}$ with probability $2/3$. Moreover, if $p_{x,1} \leq \frac{t}{2n^2} - \frac{1}{4n^2}$, then $\frac{\mathbf{B}}{r} < \frac{t}{2n^2}$ with probability $2/3$. We conclude

$$\begin{cases} (S, x, t) \in \mathcal{L}, & \text{if } p_{x,1} \geq \frac{t}{2n^2} + \frac{1}{4n^2}, \\ (S, x, t) \notin \mathcal{L}, & \text{if } p_{x,1} \leq \frac{t}{2n^2} - \frac{1}{4n^2}. \end{cases} \quad (1)$$

Note that this only gives us information about the output of \mathcal{B} on inputs (S, x, t) such that $|p_{x,1} - \frac{t}{2n^2}| \geq \frac{1}{4n^2}$. Henceforth we will call this inequality the *promise condition*. We remark that, if (S, x, t) does not satisfy the promise condition, then $(S, x, t+1)$ satisfies it. Indeed, if $|p_{x,1} - \frac{t}{2n^2}| < \frac{1}{4n^2}$, we have

$$p_{x,1} \leq \frac{t}{2n^2} + \frac{1}{4n^2} = \frac{t+1}{2n^2} - \frac{1}{2n^2} + \frac{1}{4n^2} = \frac{t+1}{2n^2} - \frac{1}{4n^2},$$

and therefore $|p_{x,1} - \frac{t+1}{2n^2}| \geq \frac{1}{4n^2}$. In particular, we obtain that $(S, x, t+1) \notin \mathcal{L}$.

Description of $\mathcal{A}^{\mathcal{L}}$. Given access to this PP language, \mathcal{A} will query the oracle on the inputs $\{(S, x, 0), \dots, (S, x, 2n^2)\}$. Then \mathcal{A} will output $\frac{t}{2n^2}$, for the smallest t such that (S, x, t) is rejected by the oracle (if (S, x, t) is accepted for all t , then \mathcal{A} outputs 1). Note that, if $(S, x, t-1)$ is accepted by the oracle, and (S, x, t) is rejected, then one of the following three things must have happened:

1. $(S, x, t-1) \in \mathcal{L}$ and $(S, x, t) \notin \mathcal{L}$, or
2. $(S, x, t-1)$ does not satisfy the promise condition, and $(S, x, t) \notin \mathcal{L}$, or
3. $(S, x, t-1) \in \mathcal{L}$ and (S, x, t) does not satisfy the promise condition.

As we observed above, it never occurs that $(S, x, t-1)$ does not satisfy the promise condition, and $(S, x, t) \in \mathcal{L}$. By inspection, using (1), in all of those three cases we get that the value $\frac{t}{2n^2}$ is an additive approximation to $p_{x,1}$ with error at most $\frac{1}{n^2}$. Additionally, since $p_{x,0} = 1 - p_{x,1}$, the value $\frac{2n^2-t}{2n^2}$ is an additive approximation for $p_{x,0}$. \square

When $\text{Samp}(1^n) \rightarrow (k, s)$ is a uniform quantum polynomial-time algorithm outputting a pair of classical strings (k, s) , we will denote by $(\text{Samp} | s')^{\text{key}}$ the distribution of keys k output by $\text{Samp}(1^n)$, conditioned on the puzzle being s' .

Lemma 4.2. *Let Samp be a (uniform) quantum polynomial time algorithm such that $\text{Samp}(1^n)$ outputs a pair of classical strings (k, s) , denoted as the key and the puzzle respectively, where $k \in \{0, 1\}^n$. There exists a poly-time quantum algorithm \mathcal{A} and a PP language \mathcal{L} such that $\mathcal{A}^{\mathcal{L}}$ takes as input a puzzle s' and outputs a key k' , and whose distribution has total variation distance at most $1/n$ from $(\text{Samp} | s')^{\text{key}}$. In other words, we have*

$$(k' | k' \leftarrow \mathcal{A}^{\mathcal{L}}(1^n, s')) \approx_{\frac{1}{n}} (\text{Samp} | s')^{\text{key}}.$$

Proof. On a high level, the algorithm \mathcal{A} will output a key k' by sampling its bits one by one. At the i^{th} iteration, it will use the first $i - 1$ bits of k' to estimate the distribution of the i^{th} bit using a PP oracle, as in Lemma 4.1. Then it will sample the i^{th} bit of k' according to this estimated distribution. The PP language \mathcal{L} that we use is the same as the one shown to exist in that lemma.

In more detail, let us define the sequence $\{\mathcal{S}^{\leq i}\}_{i \in [n]}$ of algorithms based on **Samp**, where $\mathcal{S}^{\leq i}$ is the (uniform) quantum polynomial time algorithm that simulates **Samp**, but only outputs the puzzle s and the first i bits of the key k .

On input $(1^n, s')$, the algorithm \mathcal{A} will proceed in n iterations. In the first iteration, it will use Lemma 4.1 on $\mathcal{S}^{\leq 1}$ with output s' and estimate (up to $\frac{1}{n^2}$ additive error) the probability that the first bit of the key is 1, conditioned on the puzzle s' . Call this estimate \tilde{p}_1 . The algorithm will then sample the first bit of k' according to the Bernoulli distribution defined by \tilde{p}_1 .

Now \mathcal{A} proceeds by sampling the remaining bits. In the i^{th} iteration, it uses Lemma 4.1 with the sampler $\mathcal{S}^{\leq i}$ and outputs an estimate \tilde{p}_i of the probability that the i^{th} bit of k' is 1, conditioned on the puzzle s' and the first $i - 1$ bits of k' . Then it samples the i^{th} bit according to the estimated distribution. Finally, \mathcal{A} outputs k' after the end of the n^{th} iteration.

It remains to show that the output distribution of \mathcal{A} is close to $(\mathbf{Samp} | s')^{\text{key}}$, the output distribution of **Samp**(1^n) conditioned on the puzzle s' . We will show this via a hybrid argument.

Let $\mathcal{D}_0 := (\mathbf{Samp} | s')^{\text{key}}$ be the true distribution of the key k conditioned on the puzzle s' . We define n hybrid distributions \mathcal{D}_i for $i \in \{1, \dots, n\}$ on keys. The hybrid \mathcal{D}_i runs in n iterations and in each iteration samples the next bit of the key. In the first i iterations (that correspond to the first i bits), the distribution \mathcal{D}_i uses the estimated probabilities $\tilde{p}_1, \dots, \tilde{p}_i$. The final $n - i$ bits are sampled according to the true conditional probabilities (i.e., according to $(\mathbf{Samp} | s')^{\text{key}}$, conditioned on the outcome of the previous bits). Note that \mathcal{D}_n now corresponds to the output distribution of our algorithm \mathcal{A} . We show that every two consecutive distributions are at most $\frac{1}{n^2}$ far in total variation distance, and thus the total distance between the true distribution of the key and the output distribution of \mathcal{A} is at most $\frac{1}{n}$ from the triangle inequality.

Claim 4.3. *For every $i \in \{0, \dots, n - 1\}$, we have*

$$d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}_{i+1}) \leq \frac{1}{n^2}.$$

Proof. Before diving into the proof, we introduce some useful notation. We use $k_{[x,y]}$ to denote the length- $(y - x + 1)$ substring that includes bits $\{x, \dots, y\}$ of k . Additionally, even though \mathcal{D}_i is a distribution over n -bit strings, we will abuse notation and consider the probability assigned to substrings of the form $k_{[x,y]}$. In that case, we write

$$\mathcal{D}_i(k_{[x,y]}) := \Pr_{k' \leftarrow \mathcal{D}_i} \left[k'_{[x,y]} = k_{[x,y]} \right].$$

We will also consider the probability assigned to substrings $k[x, y]$, conditioned on a prefix $k_{[1,x]}$, in which case we write

$$\mathcal{D}_i(k_{[x,y]} | k_{[1,x-1]}) := \Pr_{k' \leftarrow \mathcal{D}_i} \left[k'_{[x,y]} = k_{[x,y]} \mid k'_{[1,x-1]} = k_{[1,x-1]} \right].$$

With this notation in place, a direct calculation suffices for the $i = 0$ case. In particular, observe that

$$\begin{aligned} \mathcal{D}_1(k) &= \mathcal{D}_1(k_{[1,1]}) \cdot \mathcal{D}_1(k_{[2,n]} | k_{[1,1]}) \\ &= \mathcal{D}_1(k_{[1,1]}) \cdot \mathcal{D}_0(k_{[2,n]} | k_{[1,1]}), \end{aligned}$$

and similarly $\mathcal{D}_0(k) = \mathcal{D}_0(k_{[1,1]}) \cdot \mathcal{D}_0(k_{[2,n]}|k_{[1,1]})$. Note that, because of Lemma 4.1, we have $|\mathcal{D}_0(k_{[1,1]}) - \mathcal{D}_1(k_{[1,1]})| \leq 1/n^2$. The total variation distance satisfies:

$$\begin{aligned}
d_{\text{TV}}(\mathcal{D}_0, \mathcal{D}_1) &= \frac{1}{2} \sum_{k \in \{0,1\}^n} |\mathcal{D}_0(k) - \mathcal{D}_1(k)| \\
&= \frac{1}{2} \sum_{k \in \{0,1\}^n} |\mathcal{D}_0(k_{[1,1]}) \cdot \mathcal{D}_0(k_{[2,n]}|k_{[1,1]}) - \mathcal{D}_1(k_{[1,1]}) \cdot \mathcal{D}_0(k_{[2,n]}|k_{[1,1]})| \\
&\leq \frac{1}{2} \sum_{k \in \{0,1\}^n} |\mathcal{D}_0(k_{[1,1]}) - \mathcal{D}_1(k_{[1,1]})| \cdot |\mathcal{D}_0(k_{[2,n]}|k_{[1,1]})| \\
&\leq \frac{1}{2n^2} \sum_{k \in \{0,1\}^n} |\mathcal{D}_0(k_{[2,n]}|k_{[1,1]})| \\
&\leq \frac{1}{n^2}.
\end{aligned}$$

Let us now consider the distributions $\mathcal{D}_i, \mathcal{D}_{i+1}$. They both sample bits 1 up to i using the estimated probabilities, and bits $i+2$ up to n with the true conditional probabilities. Write \mathcal{D}_{i+1} as

$$\begin{aligned}
\mathcal{D}_{i+1}(k) &= \mathcal{D}_{i+1}(k_{[1,i]}) \cdot \mathcal{D}_{i+1}(k_{[i+1,i+1]}|k_{[1,i]}) \cdot \mathcal{D}_{i+1}(k_{[i+2,n]}|k_{[1,i+1]}) \\
&= \mathcal{D}_{i+1}(k_{[1,i]}) \cdot \mathcal{D}_{i+1}(k_{[i+1,i+1]}|k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]})
\end{aligned}$$

and similarly \mathcal{D}_i as

$$\begin{aligned}
\mathcal{D}_i(k) &= \mathcal{D}_i(k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+1,i+1]}|k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]}) \\
&= \mathcal{D}_{i+1}(k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+1,i+1]}|k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]}) .
\end{aligned}$$

Note that, because of Lemma 4.1, we have

$$|\mathcal{D}_i(k_{[i+1,i+1]}|k_{[1,i]}) - \mathcal{D}_{i+1}(k_{[i+1,i+1]}|k_{[1,i]})| < 1/n^2.$$

The total variation distance satisfies:

$$\begin{aligned}
& d_{\text{TV}}(\mathcal{D}_i, \mathcal{D}_{i+1}) \\
&= \frac{1}{2} \sum_{k \in \{0,1\}^n} |\mathcal{D}_i(k) - \mathcal{D}_{i+1}(k)| \\
&\leq \frac{1}{2} \sum_{k \in \{0,1\}^n} \mathcal{D}_{i+1}(k_{[1,i]}) \cdot |\mathcal{D}_i(k_{[i+1,i+1]}|k_{[1,i]}) - \mathcal{D}_{i+1}(k_{[i+1,i+1]}|k_{[1,i]})| \cdot \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]}) \\
&\leq \frac{1}{2n^2} \sum_{k \in \{0,1\}^n} \mathcal{D}_{i+1}(k_{[1,i]}) \cdot \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]}) \\
&= \frac{1}{2n^2} \sum_{k_{[1,i]} \in \{0,1\}^i} \mathcal{D}_{i+1}(k_{[1,i]}) \sum_{k_{[i+1,i+1]} \in \{0,1\}} \sum_{k_{[i+2,n]} \in \{0,1\}^{n-i-1}} \mathcal{D}_i(k_{[i+2,n]}|k_{[1,i+1]}) \\
&= \frac{1}{n^2} \sum_{k_{[1,i]} \in \{0,1\}^i} \mathcal{D}_{i+1}(k_{[1,i]}) \\
&= \frac{1}{n^2}. \quad \square
\end{aligned}$$

As observed above, by the triangle inequality we obtain from the Claim that \mathcal{D}_n , which is the output distribution of our algorithm \mathcal{A} , has total variation distance at most $1/n$ from $(\text{Samp} | s')^{\text{key}}$. \square

Theorem 4.4. *For any one-way puzzle $(\text{Ver}, \text{Samp})$, there exists a PP language \mathcal{L} , and a poly-time quantum algorithm $\mathcal{A}^{\mathcal{L}}$ such that*

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [\text{Ver}(\mathcal{A}^{\mathcal{L}}(s), s) = \top] \geq \frac{1}{2}.$$

Proof. From the correctness property of the OWPuzzle, it holds that

$$\Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [\text{Ver}(k, s) = \top] \geq 1 - \text{negl}(n).$$

Recall that $(\text{Samp} | s')^{\text{key}}$ is the distribution over keys output by Samp , conditioned on the puzzle being equal to s' . It is clear that sampling the puzzle first, and then the key does not change their joint distribution, and thus

$$\Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow (\text{Samp} | s)^{\text{key}}}} [\text{Ver}(k', s) = \top] \geq 1 - \text{negl}(n).$$

Given a puzzle s , Lemma 4.2 implies that there exists a quantum polynomial-time algorithm \mathcal{A} and a PP language \mathcal{L} , such that $\mathcal{A}^{\mathcal{L}}(s')$ outputs a key k' according to a distribution $\tilde{D}_{s'}$ such that

$$\tilde{D}_{s'} \approx_{1/n} (\text{Samp} | s')^{\text{key}}. \quad (2)$$

We have

$$\Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow \tilde{D}_s}} [\text{Ver}(k', s) = \top] = \sum_{s'} \Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow \tilde{D}_s}} [\text{Ver}(k', s) = \top | s = s'] \cdot \Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [s = s'].$$

Now note that, conditioned on $s = s'$, the event $\{\text{Ver}(k', s) = \top\}$ depends only on $k' \leftarrow \tilde{D}_{s'}$, and thus we may use (2) to get

$$\begin{aligned}
& \Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow \tilde{D}_s}} [\text{Ver}(k', s) = \top] \\
& \geq \sum_{s'} \left(\Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow (\text{Samp} | s)^{\text{key}}}} [\text{Ver}(k', s) = \top | s = s'] - 1/n \right) \cdot \Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [s = s'] \\
& = \sum_{s'} \left(\Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow (\text{Samp} | s)^{\text{key}}}} [\text{Ver}(k', s) = \top | s = s'] \Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [s = s'] \right) - (1/n) \cdot \Pr_{(k,s) \leftarrow \text{Samp}(1^n)} [s = s'] \\
& = \Pr_{\substack{(k,s) \leftarrow \text{Samp}(1^n) \\ k' \leftarrow (\text{Samp} | s)^{\text{key}}}} [\text{Ver}(k', s) = \top] - 1/n \geq 1 - 1/n - \text{negl}(n) > 1 - 2/n.
\end{aligned}$$

This completes our argument. \square

Now the desired result follows by combining Theorem 4.4 with Theorem 2.6 (Theorem 4.2 of [KT23]).

Corollary 4.5. *For any OWSG G of n -qubit states with key space \mathcal{K} and with security parameter λ , there exists a PP language \mathcal{L} , a $\text{poly}(\lambda)$ -time quantum algorithm $\mathcal{A}^{\mathcal{L}}$, and $t = \text{poly}(\lambda)$ such that*

$$\mathbb{E}_{\substack{k \leftarrow \mathcal{K} \\ k' \leftarrow \mathcal{A}^{\mathcal{L}}(|\phi_k\rangle^{\otimes t})}} \left[|\langle \phi_k | \phi_{k'} \rangle|^2 \right] \geq \frac{1}{2}.$$

Proof. The proof of Theorem 4.2 of [KT23] shows that, for every one-way state generator G , there exists a one-way puzzle $P = (\text{Samp}, \text{Ver})$ such that, if there exists a quantum polynomial-time algorithm \mathcal{A} that breaks P , then there exists a quantum polynomial-time algorithm \mathcal{A}^* that breaks G . The corollary then follows by plugging the algorithm of Theorem 4.4 for P – a polynomial-time quantum algorithm with a PP oracle that breaks P – into their reduction. \square

Acknowledgements

B. P. Cavalari was supported by the Chancellor’s International Scholarship of the University of Warwick. A. Pelecanos was supported by DARPA under Agreement No. HR00112020023. E. Goldin was supported by a National Science Foundation Graduate Research Fellowship. This work was done in part while the authors were visiting the Simons Institute for the Theory of Computing.

References

- [Aar05] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461:3473 – 3482, 2005.

- [AB09] Sanjeev Arora and Boaz Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009.
- [AGQY22] Prabhanjan Ananth, Aditya Gulati, Luowen Qian, and Henry Yuen. Pseudorandom (function-like) quantum state generators: New definitions and applications. In *Theory of Cryptography Conference*, pages 237–265. Springer, 2022.
- [BB14] Charles H. Bennett and Gilles Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, dec 2014.
- [BCG89] Mihir Bellare, Lenore Cowen, and Shafi Goldwasser. On the structure of secret key exchange protocols. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 604–605. Springer, 1989.
- [BCQ22] Zvika Brakerski, Ran Canetti, and Luowen Qian. On the computational hardness needed for quantum cryptography. *arXiv preprint arXiv:2209.04101*, 2022.
- [BHH16] Fernando GSL Brandao, Aram W Harrow, and Michał Horodecki. Local random quantum circuits are approximate polynomial-designs. *Communications in Mathematical Physics*, 346:397–434, 2016.
- [BS20] Zvika Brakerski and Omri Shmueli. Scalable pseudorandom quantum states, 2020.
- [DCEL09] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009.
- [GC01] Daniel Gottesman and Isaac Chuang. Quantum digital signatures, 2001.
- [HLW06] Patrick Hayden, Debbie W Leung, and Andreas Winter. Aspects of generic entanglement. *Communications in mathematical physics*, 265:95–117, 2006.
- [HMMH⁺23] Jonas Haferkamp, Felipe Montealegre-Mora, Markus Heinrich, Jens Eisert, David Gross, and Ingo Roth. Efficient unitary designs with a system-size independent number of non-clifford gates. *Communications in Mathematical Physics*, 397(3):995–1041, 2023.
- [Imp95] R. Impagliazzo. A personal view of average-case complexity. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 134–147, 1995.
- [INN⁺22] Sandy Irani, Anand Natarajan, Chinmay Nirkhe, Sujit Rao, and Henry Yuen. Quantum search-to-decision reductions and the state synthesis problem. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [IR89] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, STOC '89, page 44–61, New York, NY, USA, 1989. Association for Computing Machinery.

- [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 38*, pages 126–152. Springer, 2018.
- [Kre21] William Kretschmer. Quantum pseudorandomness and classical complexity. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [KT23] Dakshita Khurana and Kabir Tomer. Commitments from quantum one-wayness. *CoRR*, abs/2310.11526, 2023.
- [Kup15] Greg Kuperberg. How hard is it to approximate the jones polynomial? *Theory of Computing*, 11(6):183–219, 2015.
- [LC97] Hoi-Kwong Lo and H. F. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78(17):3410–3413, apr 1997.
- [LMW23] Alex Lombardi, Fermi Ma, and John Wright. A one-query lower bound for unitary synthesis and breaking quantum cryptography. Cryptology ePrint Archive, Paper 2023/1602, 2023. <https://eprint.iacr.org/2023/1602>.
- [Mul59] Mervin E. Muller. A note on a method for generating points uniformly on n-dimensional spheres. *Commun. ACM*, 2(4):19–20, apr 1959.
- [MY22a] Tomoyuki Morimae and Takashi Yamakawa. One-wayness in quantum cryptography. *arXiv preprint arXiv:2210.03394*, 2022.
- [MY22b] Tomoyuki Morimae and Takashi Yamakawa. Quantum commitments and signatures without one-way functions. In *Annual International Cryptology Conference*, pages 269–295. Springer, 2022.
- [OSP23] Ryan O’Donnell, Rocco A. Servedio, and Pedro Paredes. Explicit orthogonal and unitary designs, 2023.
- [Wie83] Stephen Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, jan 1983.
- [Yan22] Jun Yan. General properties of quantum bit commitments. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 628–657. Springer, 2022.