Cryptography in the Common Haar State Model: Feasibility Results and Separations

Prabhanjan Ananth$^{(\boxtimes)}$, Aditya Gulati, and Yao-Ting Lin

UCSB, Santa Barbara, USA
prabhanjan@cs.ucsb.edu, {adityagulati,yao-ting_lin}@ucsb.edu

Abstract. Common random string model is a popular model in classical cryptography. We study a quantum analogue of this model called the common Haar state (CHS) model. In this model, every party participating in the cryptographic system receives many copies of one or more i.i.d Haar random states.

We study feasibility and limitations of cryptographic primitives in this model and its variants:

- We present a construction of pseudorandom function-like states with security against computationally unbounded adversaries, as long as the adversaries only receive (a priori) bounded number of copies. By suitably instantiating the CHS model, we obtain a new approach to construct pseudorandom function-like states in the plain model.
- We present separations between pseudorandom function-like states (with super-logarithmic length) and quantum cryptographic primitives, such as interactive key agreement and bit commitment, with classical communication. To show these separations, we prove new results on the indistinguishability of identical versus independent Haar states against LOCC (local operations, classical communication) adversaries.

1 Introduction

In classical cryptography, the common random string and the common reference string models were primarily introduced to tackle cryptographic tasks that were impossible to achieve in the plain model. In the common reference string model, there is a trusted setup who produces a string that every party has access to. In the common random string model, the common string available to all the parties is sampled uniformly at random. Due to the lack of structure required from the common random string model, it is in general the more desirable model of the two. There have been many constructions proposed over the years in these two models, including non-interactive zero-knowledge [BFM19], secure computation with universal composition [CF01,CLOS02] and two-round secure computation [GS22,BL18].

It is a worthy pursuit to study similar models for quantum cryptographic protocols. In the quantum world, there is an option to define models that are intrinsically quantum in nature. For instance, we could define a model wherein

E. Boyle and M. Mahmoody (Eds.): TCC 2024, LNCS 15365, pp. 94–125, 2025.
https://doi.org/10.1007/978-3-031-78017-2_4

a trusted setup produces a quantum state and every party participating in the cryptographic system receives one or more copies of this quantum state. Indeed, two works by Morimae, Nehoran and Yamakawa [MNY23] and Qian [Qia23] consider this model, termed as the *common reference quantum state model* (CRQS). They proposed a construction of unconditionally secure commitments in this model. Quantum commitments is a foundational notion in quantum cryptography. In recent years, quantum commitments have been extensively studied [AQY22, MY21, AGQY22, MY23, BCQ23, Bra23] due to its implication to secure computation [BCKM21, GLSV21]. The fact that information-theoretically secure commitments are impossible in the plain model [LC97, May97, CLM23] renders the contributions of [MNY23, Qia23] particularly interesting.

Common Haar State Model. While CRQS is a quantum analogue of the common reference string model, in a similar vein, we can ask if there is a quantum analogue of the common random string model. We consider a novel model called the *common Haar state model* (CHS). In this model, every party in the system (including the adversary) receives many copies of many i.i.d Haar states. We believe that the CHS model is more pragmatic than the CRQS model owing to the fact that we do not require any structure from the common public state. This raises the possibility of avoiding a trusted setup altogether and instead we could rely upon naturally occuring physical processes to obtain the Haar states. This model was also recently introduced in an independent and concurrent recent work by Chen, Coladangelo and Sattath [CCS24] (henceforth, referred to as CCS).

There are three reasons to study this model. Firstly, this model allows us to bypass impossibility results in the plain model. For instance, as we will see later, primitives that require computational assumptions in the plain model, can instead be designed with information-theoretic security in the CHS model. Second, perhaps a less intuitive reason, is that the constructions proposed in this model can, in some cases, be adopted to obtain constructions in the plain model by instantiating the Haar states either using state designs or pseudorandom state generators (PRSGs) [JLS18]. This leads to a modular approach of designing cryptographic primitives from PRS: first design the primitive in the CHS model and then instantiate the common Haar state using PRS. Finally, this model can be leveraged to demonstrate separations between different quantum cryptographic primitives.

1.1 Our Results

We explore both feasibility results and black-box separations in the CHS model.

Feasibility Results.

Pseudorandom Function-like States with Statistical Security. We study the possibility of designing pseudorandom function-like state generators (PRFSGs), introduced by Ananth, Qian and Yuen [AQY22], with statistical security in the CHS model. Roughly speaking, a PRFSG is an efficient keyed quantum circuit that can be used to produce many pseudorandom states. We refer the reader to Appendix A of the full version [AGL24] for a detailed discussion on the different notions of pseudorandomness in the quantum world.

We are interested in designing (λ, m, n, t)-PR$\underline{\mathrm{F}}$SGs in the setting when $n \geq \lambda$ and $m = \Omega(\log(\lambda))$, where λ is the key length, m is the input length, n is the output length (and also the number of the qubits in the common Haar state) and t is the maximum number of queries that can be requested by the adversary. However, in the CHS model, we can in fact achieve statistical security.

We show the following.

Theorem 1 (Informal). *There is a statistically secure (λ, m, n, ℓ)-PR$\underline{F}$SG in the CHS model, for $m = \lambda^c$, $n \geq \lambda$ and $\ell = O\left(\frac{\lambda^{1-c}}{\log(\lambda)^{1+\varepsilon}}\right)$, for any constant $\varepsilon > 0$ and for all $c \in [0, 1)$.*

CCS is the only other work that has studied pseudorandomness in the CHS model. There are a few advantages of our result over CCS:

- Our theorem subsumes and generalizes the result of CCS who showed (λ, n, t)-PRSGs exists in their model, where the output length is larger than the key length, i.e., $n > \lambda$ and moreover, when $t = 1$ with t being the number of copies of the PRS state given to the adversary.
- Our construction, when restricted to the case of PRSGs, is slightly simpler than CCS: in CCS, on a subset of qubits of the Haar state, a random Pauli operator is applied whereas in our case a random Pauli Z operator is applied. Our construction of PR$\underline{\mathrm{F}}$SG uses the seminal Goldreich-Goldwasser-Micali approach [GGM86] to go from one-query security to many-query security.
- They propose novel sophisticated tools in their analysis whereas our analysis is arguably more elementary using well known facts about symmetric subspaces.
- Finally, we can achieve arbitrary stretch whereas it is unclear whether this is also achieved by CCS.

As a side contribution, the proof of our PRSG construction also simplifies the proof of the quantum public-key construction of Coladangelo [Col23]; this is due to the fact the core lemma proven in [Col23] is implied by the above theorem.

Interestingly, the above theorem has implications for computationally secure pseudorandomness in the plain model. Specifically, we obtain the following corollary by instantiating the CHS model using stretch PRSGs:

Corollary 1. Assuming (λ, n, ℓ)-PRSGs, there exists (λ', m, n, t)-PR$\underline{\mathrm{F}}$SGs, where $n > \lambda' > \lambda$, $m = \lambda^c$ and $\ell = O\left(\frac{\lambda^{1-c}}{\log(\lambda)^{1+\varepsilon}}\right)$, for any constant $\varepsilon > 0$ and $c \in [0, 1)$.

Prior to our work, stretch PRFSGs for super-logarithmic input length, even in the bounded query setting, was only known from one-way functions [AQY22]. This complements the work of [AQY22] who showed a construction of PRFSGs for logarithmic input length from PRSGs.

Interestingly, the state generators in both works (CCS and ours) only consume one copy of a single Haar state. In this special case, it is interesting to understand whether we can extend our result to the setting when the adversary receives $\frac{\lambda}{\log(\lambda)}$ copies or more. We show this is not possible.

Theorem 2 (Informal). *There does not exist a secure (λ, m, n, ℓ)-PR$\underline{\text{F}}$SG, for any $m \geq 1$, in the CHS model, where $n = \omega(\log(\lambda))$ and $\ell = \Omega\left(\frac{\lambda}{\log(\lambda)}\right)$.*

CCS also proved a lower bound where they showed that unbounded copy pseudorandom states do not exist. Their negative result is stronger in the sense that they rule out PRSGs who use up many copies of the Haar states from the CHRS and thus, their work gives a clean separation between 1-copy stretch PRS and unbounded copy PRS which was not known before. On the other hand, for the special case when the PR$\underline{\text{F}}$SG takes only one copy of the Haar state, we believe our result yields better parameters.

Commitments. In addition to pseudorandomness, we also study the possibility of constructing other cryptographic primitives in the CHS model. We show the following:

Theorem 3 (Informal). *There is an unconditionally secure bit commitment scheme in the CHS model.*

Both our construction and the commitments scheme proposed by CCS are different although they share strong similarities.

Black-Box Separations

LOCC Indistinguishability. We separate pseudorandom function-like states and quantum cryptographic primitives with classical communication using a variant of the CHS model. At the heart of our separations is a novel result that proves indistinguishability of identical versus independent Haar states against LOCC (local operations, classical communication) adversaries. More precisely, (A, B) is an LOCC adversary if A and B are quantum algorithms who can communicate with each other via only classical communication channels. It is important that A and B do not share any entanglement. Moreover, we restrict our attention to LOCC distinguishers which are LOCC adversaries of the form (A, B) where A does not output anything whereas B outputs a single bit. We say that a LOCC distinguisher (A, B) can distinguish two states ρ_{AB} and σ_{AB} with probability at most ε, referred to as ε-LOCC indistinguishability, where A receives the register A and B receives the register B, if $|\mathsf{Pr}[1 \leftarrow (A, B)(\rho_{\mathsf{AB}})] - \mathsf{Pr}[1 \leftarrow (A, B)(\sigma_{\mathsf{AB}})]| = \varepsilon$. Of particular interest is the case when

$$\rho_{\mathsf{AB}} = \mathop{\mathbb{E}}_{|\psi\rangle \leftarrow \mathcal{H}_n} \left[(|\psi\rangle^{\otimes t})_{\mathsf{A}} \otimes (|\psi\rangle^{\otimes t})_{\mathsf{B}} \right], \ \sigma_{\mathsf{AB}} = \mathop{\mathbb{E}}_{\substack{|\psi\rangle \leftarrow \mathcal{H}_n, \\ |\phi\rangle \leftarrow \mathcal{H}_n}} \left[(|\psi\rangle^{\otimes t})_{\mathsf{A}} \otimes (|\phi\rangle^{\otimes t})_{\mathsf{B}} \right]$$

Here, $\mathcal{H}_n$ denotes the Haar distribution on n-qubit quantum states and t is polynomial in n. A couple of works by Harrow [Har23] and Chen, Cotler, Huang and Li [CCHL22] prove that the LOCC indistinguishability of ρ_{AB} and σ_{AB} is negligible in n in the case when $t = 1$. In this work, we extend to the case when t is arbitrary.

Theorem 4. *ρ_{AB} and σ_{AB} (defined above) are ε-LOCC indistinguishable, where $\varepsilon = O\left(\frac{t^2}{2^n}\right)$.*

We also show that the above bound is tight by demonstrating an LOCC distinguisher whose distinguishing probability is $\Theta(\frac{t^2}{2^n})$.

Recently, Ananth, Kaleoglu and Yuen [AKY24] prove the indistinguishability of ρ_{AB} and σ_{AB} in the dual setting, against non-local adversaries that can share entanglement but cannot communicate.

The above theorem can easily be extended to the multi-party setting where either all the parties get (many copies of) the same Haar state or they receive i.i.d Haar states.

Separations. We use Theorem 4 to show that some quantum cryptographic primitives with classical communication are impossible in the CHS model. Let us develop some intuition towards proving such a statement. Suppose there are two or more parties participating in a quantum cryptographic protocol with classical communication in the CHS model. By definition, all the parties would receive many, say t, copies of $|\psi\rangle$, where $|\psi\rangle$ is sampled from the Haar distribution. Since the parties can only exchange classical messages, thanks to Theorem 4, without affecting correctness or security we can modify the protocol wherein for each party, say P_i, a Haar state $|\psi_i\rangle$ is sampled and t copies of $|\psi_i\rangle$ is given to P_i. From this, we can extract a quantum cryptographic primitive in the plain model since each party can sample a Haar state on its own. In conclusion, quantum cryptographic primitives with classical communication in the CHS model can be turned into their counterparts in the plain model.

This gives a natural recipe for proving impossibility results in the CHS model. We apply this recipe to obtain impossibility results for interactive key agreements and interactive commitments.

Theorem 5. *Interactive quantum key agreement and interactive quantum commitment protocols, with classical communication, are impossible in the CHS model.*

We extend the above theorem to separate interactive quantum key agreement and interactive quantum commitments from pseudorandom function-like state generators. The separations are obtained by considering a variant of the CHS model where the adversary does not get access to many copies of one Haar state but instead gets access to infinitely many input-less oracles[1]

[1] We note that [Kre21] made similar use of infinitely many oracles to prove a separation between pseudorandom states and one-way functions.

$\{\{G_{k,x}\}_{k,x\in\{0,1\}^\lambda}\}_{\lambda\in\mathbb{N}}$ such that each $G_{k,x}$ produces a copy of a Haar state $|\psi_{k,x}\rangle$. In this model, it is easy to construct pseudorandom function-like states. However, an extension of Theorem 5 rules out the possibility of interactive quantum key agreement and quantum commitments with classical communication in this variant. Thus, we have the following.

Theorem 6. *There does not exist a black-box reduction from interactive quantum key agreement and quantum commitments with classical communication to pseudorandom function-like states.*

Prior work by Chung, Goldin and Gray [CGG24] extensively studies the separations between quantum cryptographic primitives with classical communication and different quantum pseudorandomness notions. However, their framework did not capture the above result.

Prior works by [ACC+22, CLM23, LLLL24] ruled out quantum key agreements and non-interactive commitments with classical communication from post-quantum one-way functions. However, their separation was either based on a conjecture or in a restricted setting whereas our result is unconditional. This makes our result incomparable with the results from [ACC+22, CLM23, LLLL24]. Our work follows a long line of recent works [HY20, ACC+22, AHY23, CLM23, ACH+23, BGVV+23, BM+24, CM24] that make progress in understanding the landscape of black-box separations in quantum cryptography.

2 Technical Overview

2.1 Pseudorandomness in the CHS Model

Warmup: Pseudorandom State Generators (PRSGs). As a warmup, we first study 1-copy PRSG in the CHS model. Consider the following construction: $G_k(|\vartheta\rangle) := (Z^k \otimes I_{n-\lambda})|\vartheta\rangle$, where $Z^k = Z^{k_1} \otimes \cdots \otimes Z^{k_\lambda}$, $k = k_1 \cdots k_\lambda \in \{0,1\}^\lambda$ and $I_{n-\lambda}$ is an identity operator on $n-\lambda$ qubits. In other words, G_k applies a random Pauli Z operator only on the first λ qubits and does not touch the rest. Note that this construction already satisfies the stretch property (i.e. the output length is larger than the key length).

Let us consider the case when the adversary receives just one copy of $|\vartheta\rangle$ and is expected to distinguish $G_k(|\vartheta\rangle)$ versus an independent Haar state $|\varphi\rangle$. Formally, we would like to argue that the following states are close.

$$\rho := \mathop{\mathbb{E}}_{\substack{k\leftarrow\{0,1\}^\lambda \\ |\vartheta\rangle\leftarrow\mathcal{H}_n}} \left[G_k(|\vartheta\rangle) \otimes |\vartheta\rangle\langle\vartheta|\right] \text{ and } \sigma := \frac{I}{2^n} \otimes \frac{I}{2^n}.$$

By the properties of the symmetric subspace, the following holds:

$$\mathop{\mathbb{E}}_{|\vartheta\rangle\leftarrow\mathcal{H}_n}\left[|\vartheta\rangle\langle\vartheta|^{\otimes 2}\right] \approx_\varepsilon \mathop{\mathbb{E}}_{x,y\leftarrow[2^n],x^1\neq y^1}\left[\frac{1}{2}\left(|xy\rangle\langle xy| + |xy\rangle\langle yx| + |yx\rangle\langle xy| + |yx\rangle\langle yx|\right)\right],$$

where ε is negligible in n and the notation x^1 (respectively, y^1) denotes the first λ bits of x (respectively, y). Now, applying a random Z operator on the first λ qubits tantamounts to measuring the first λ qubits in the computational basis. Given the fact that $x^1 \neq y^1$, this measurement unentangles the last n qubits. Thus, the result is a state of the form $\mathbb{E}_{x,y \leftarrow [2^n], x^1 \neq y^1} \left[\frac{1}{2} |x\rangle\langle x| \otimes |y\rangle\langle y| + \frac{1}{2} |y\rangle\langle y| \otimes |x\rangle\langle x| \right]$. This state is in turn close to $\frac{I}{2^n} \otimes \frac{I}{2^n}$.

Generalizing to Many Copies of the CHS. Next, we to generalize the above approach to even when polynomially many copies of the CHS are provided. Formally, we would like to argue that the following two states are close.

$$\rho := \mathop{\mathbb{E}}_{\substack{k \leftarrow \{0,1\}^\lambda \\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[G_k(|\vartheta\rangle) \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t} \right] \text{ and } \sigma := \mathop{\mathbb{E}}_{\substack{|\varphi\rangle \leftarrow \mathcal{H}_n \\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[|\varphi\rangle\langle\varphi| \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t} \right],$$

where t is some polynomial of n. Note that, by the property of the Haar distribution, we can simplify σ to

$$\sigma = \frac{I}{2^n} \otimes \mathop{\mathbb{E}}_{T \leftarrow [0:t]^N} |T\rangle\langle T|,$$

where $|T\rangle$ is a type state[2] and $N = 2^n$. Note that by the properties of the symmetric subspace,

$$\mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathcal{H}_n} \left[|\vartheta\rangle\langle\vartheta|^{\otimes t+1} \right] \approx_\varepsilon \mathop{\mathbb{E}}_{\substack{T \leftarrow [0:t+1]^N \\ T \text{ is } \lambda\text{-prefix collision-free}}} |T\rangle\langle T|,$$

where ε is negligible in n and T is λ-prefix collision-free if $T \in \{0,1\}^N$ and for any $x, y \in T$[3] with $x \neq y$ implies $x^1 \neq y^1$, where the notation x^1 (respectively, y^1) denotes the first λ bits of x (respectively, y). Note that, any λ-prefix collision-free type T,

$$|T\rangle = \frac{1}{\sqrt{\binom{t+1}{t}}} \sum_{x \in T} |x\rangle |T \setminus \{x\}\rangle.$$

Again, applying a random Z operator on the first λ qubits tantamounts to measuring the first λ qubits in the computational basis. Given the fact that T is λ-prefix collision-free, this measurement unentangles the first n qubits. Thus, the result is a state of the form

$$\mathop{\mathbb{E}}_{\substack{T \leftarrow [0:t+1]^N \\ T \text{ is } \lambda\text{-prefix collision-free} \\ x \leftarrow T}} \left[|x\rangle\langle x| \otimes |T \setminus \{x\}\rangle\langle T \setminus \{x\}| \right].$$

This state is in turn close to $\frac{I}{2^n} \otimes \mathbb{E}_{T \leftarrow [0:t]^N} |T\rangle\langle T|$.

[2] We encourage readers unfamiliar with type states to refer to Definition 7.

[3] Since $T \in \{0,1\}^N$, we can treat it as a set, in particular the set associated to T is $\{i : T[i] = 1\}$.

Generalizing to ℓ-Copy PRSG. Finally, we generalize this ℓ-copy PRSG. Formally, we would like to argue that the following two states are close.

$$\rho := \mathop{\mathbb{E}}_{\substack{k\leftarrow\{0,1\}^\lambda\\ |\vartheta\rangle\leftarrow\mathcal{H}_n}}\left[G_k(|\vartheta\rangle)^{\otimes\ell}\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t}\right] \text{ and } \sigma := \mathop{\mathbb{E}}_{\substack{|\varphi\rangle\leftarrow\mathcal{H}_n\\ |\vartheta\rangle\leftarrow\mathcal{H}_n}}\left[|\varphi\rangle\langle\varphi|^{\otimes\ell}\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t}\right],$$

where ℓ, t is some polynomial of n. Note that, by the property of the Haar distribution, we can simplify σ to

$$\sigma = \mathop{\mathbb{E}}_{T_1\leftarrow[0:\ell]^N}|T_1\rangle\langle T_1|\otimes\mathop{\mathbb{E}}_{T_2\leftarrow[0:t]^N}|T_2\rangle\langle T_2|,$$

where $|T_1\rangle, |T_2\rangle$ are type states and $N = 2^n$. Note that, similar to the last case, we can still write,

$$\mathop{\mathbb{E}}_{|\vartheta\rangle\leftarrow\mathcal{H}_n}\left[|\vartheta\rangle\langle\vartheta|^{\otimes t+\ell}\right]\approx_\varepsilon \mathop{\mathbb{E}}_{\substack{T\leftarrow[0:t+\ell]^N\\ T\text{ is }\lambda\text{-prefix collision-free}}}|T\rangle\langle T|,$$

and any λ-prefix collision-free type T,

$$|T\rangle = \frac{1}{\sqrt{\binom{t+\ell}{\ell}}}\sum_{\substack{T_1\subset T\\ |T_1|=\ell}}|T_1\rangle|T\setminus T_1\rangle.$$

Ideally, we would want the application of $(Z^k\otimes I_{n-\lambda})^{\otimes\ell}$ to unentangle $|T_1\rangle$ from $|T\setminus T_1\rangle$. This is equivalent to measuring the first ℓ registers in the type basis. This is in general not true, not true. Hence, we settle for the next best thing, which is finding a "dense-enough"[4] subset of λ-prefix collision-free type such that $(Z^k\otimes I_{n-\lambda})^{\otimes\ell}$ to unentangle $|T_1\rangle$ from $|T\setminus T_1\rangle$. We find this subset to be "λ-prefix ℓ-fold collision-free" types.

We say that a λ-prefix collision-free type T is "λ-prefix ℓ-fold collision-free" if for all pairs of ℓ sized subsets $T_1, T_2\subset T$, $\oplus_{x\in T_1}x = \oplus_{x\in T_2}x$ only if $T_1 = T_2$. We start by noting that this subset is only "dense-enough" if $\ell = O\left(\frac{\lambda}{\log(\lambda)^{1+\varepsilon}}\right)$, for any constant $\varepsilon > 0$.[5]

Next, we show that for these λ-prefix ℓ-fold collision-free types states, applying a random $(Z^k\otimes I_{n-\lambda})^{\otimes\ell}$ is equivalent to measuring the first ℓ registers in the type basis. This is because $(Z^k\otimes I_{n-\lambda})^{\otimes\ell}$ on a type state $|T_1\rangle$ is equivalent to adding a phase of $(-1)^{k\cdot(\oplus_{x\in T_1}x)}$. Hence,

$$\begin{aligned}&\mathop{\mathbb{E}}_k\left[(Z^k\otimes I_{n-\lambda})^{\otimes\ell}\otimes I_{tn}|T\rangle\langle T|(Z^k\otimes I_{n-\lambda})^{\otimes\ell}\otimes I_{tn}\right]\\ &=\mathop{\mathbb{E}}_k\left[\frac{1}{\binom{t+\ell}{\ell}}\sum_{\substack{T_1,T_2\subset T\\ |T_1|=|T_2|=\ell}}(-1)^{k\cdot(\oplus_{x\in T_1}x\bigoplus\oplus_{y\in T_2}y)}|T_1\rangle|T\setminus T_1\rangle\langle T_2|\langle T\setminus T_2|\right],\end{aligned}$$

[4] Here, by dense-enough, we mean when picking a random type from λ-prefix collision-free, it lies in this subset with probability $1-\mathsf{negl}$.

[5] Later, in the impossibility result, we show that this is in fact the best we can hope for as a larger subset would bypass the impossibility result.

which for λ-prefix ℓ-fold collision-free types states is non-zero only if $T_1 = T_2$, giving us

$$\mathop{\mathbb{E}}_{k}\left[(Z^k \otimes I_{n-\lambda})^{\otimes \ell} \otimes I_{tn}|T\rangle\langle T|(Z^k \otimes I_{n-\lambda})^{\otimes \ell} \otimes I_{tn}\right] = \mathop{\mathbb{E}}_{\substack{T_1 \subset T\\ |T_1|=\ell}} \left[|T_1\rangle\langle T_1| \otimes |T \setminus T_1\rangle\langle T \setminus T_1|\right].$$

Over expectation over all λ-prefix ℓ-fold collision-free types states, this state is close to $\mathbb{E}_{T_1 \leftarrow [0:\ell]^N}|T_1\rangle\langle T_1| \otimes \mathbb{E}_{T_2 \leftarrow [0:t]^N}|T_2\rangle\langle T_2|$.

Limitations. To complement our result, we show that a t-copy PRSG is impossible in the CHS model, for $\ell = O\left(\frac{\lambda}{\log(\lambda)}\right)$ (for a restricted class of PRSG constructs which only takes one copy of the common Haar state). We show this by showing that the rank of σ grows much faster than the rank of ρ, hence, a simple distinguisher is a projector on the eigenspace of ρ. In particular, let $\tilde{G}_k(\vartheta)$ be the PRSG. Then define

$$\rho := \mathop{\mathbb{E}}_{\substack{k \leftarrow \{0,1\}^\lambda\\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[\tilde{G}_k(|\vartheta\rangle)^{\otimes \ell} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t}\right] \text{ and } \sigma := \mathop{\mathbb{E}}_{\substack{|\varphi\rangle \leftarrow \mathcal{H}_n\\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[|\varphi\rangle\langle\varphi|^{\otimes \ell} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t}\right]$$

Now since $\tilde{G}_k(|\vartheta\rangle)$ is a PRSG, its output is negligibly close to a pure state. This means that the rank of $\rho \leq 2^\lambda \binom{2^n+t+\ell-1}{t+\ell}$. In contrast, the rank of $\sigma = \binom{2^n+\ell-1}{\ell}\binom{2^n+t-1}{t}$. Note that, for $t = \lambda^3$ and $\ell = \lambda/\log(\lambda)$, $\mathrm{rank}(\rho)/\mathrm{rank}(\sigma) = \mathsf{negl}$. Hence, we can find a distinguisher. Here the distinguisher just projects onto the eigenspace of ρ, ρ gets accepted with probability 1 but σ gets accepted with probability negl, hence giving a disguisher. Since PRFSs imply PRSs (by setting $c = 0$), achieving an ℓ-query statistical PRFS in the CHS model for $\ell = \Omega(\lambda/\log(\lambda))$ is impossible.

Pseudorandom Function-like State Generators. Next we extend this idea from PRSGs to achieve PRFSGs. We take inspiration from the seminal Goldreich-Goldwasser-Micali approach [GGM86]. In particular, on the key $K = (k_1^0, \ldots, k_m^0, k_1^1, \ldots, k_m^1) \in \{0,1\}^{2\lambda' m}$ and the input $\mathbf{x} = (x_1, \ldots, x_m) \in \{0,1\}^m$, define the PRFSG $G_K(\mathbf{x}, |\vartheta\rangle)$ as follows: $G_K(\mathbf{x}, |\vartheta\rangle) = (Z^{\bigoplus_{i=1}^m k_i^{x_i}} \otimes I_{n-\lambda'})|\vartheta\rangle$. Formally, the following two states are close:

$$\rho := \mathop{\mathbb{E}}_{\substack{K \leftarrow \{0,1\}^{2m\lambda'}\\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[\otimes_{i=1}^q G_K(\mathbf{x}^i, |\vartheta\rangle)^{\otimes \ell_i} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t}\right],$$

and

$$\sigma := \mathop{\mathbb{E}}_{\substack{\forall i \in [q], |\varphi_i\rangle \leftarrow \mathcal{H}_n\\ |\vartheta\rangle \leftarrow \mathcal{H}_n}} \left[\otimes_{i=1}^q |\varphi_i\rangle\langle\varphi_i|^{\otimes \ell_i} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t}\right],$$

for all $\mathbf{x}^1, \ldots, \mathbf{x}^q \in \{0,1\}^m$ and $\ell_1, \ldots, \ell_q$ such that $\sum_{i=1}^q \ell_i = \ell$, for $\ell = O\left(\frac{\lambda^{1-c}}{\log(\lambda)^{1+\varepsilon}}\right)$ and $m = \lambda^c$, for any constant $\varepsilon > 0$ and $c \in [0,1)$.

Just as before, we can write σ as follows:

$$\sigma = \bigotimes_{i=1}^{q} \mathop{\mathbb{E}}_{T_i \leftarrow [0:\ell_i]^N} |T_i\rangle\langle T_i| \otimes \mathop{\mathbb{E}}_{\tilde{T} \leftarrow [0:t]^N} |\tilde{T}\rangle\langle\tilde{T}|,$$

where T_i's and $\tilde{T}$ are type states and $N = 2^n$. Note that, similar to the last case, we can still write,

$$\mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathcal{H}_n} \left[|\vartheta\rangle\langle\vartheta|^{\otimes t+\ell}\right] \approx_\varepsilon \mathop{\mathbb{E}}_{\substack{T \leftarrow [0:t+\ell]^N \\ T \text{ is } \lambda\text{-prefix } \ell\text{-fold collision-free}}} |T\rangle\langle T|,$$

and any λ-prefix ℓ-fold collision-free type T,

$$|T\rangle = \frac{1}{\sqrt{\binom{t+\ell}{\ell}}} \sum_{\substack{T_1 \subset T \\ |T_1| = \ell}} |T_1\rangle|T \setminus T_1\rangle.$$

Now, after application of one layer of $(Z^k \otimes I_{n-\lambda})^{\otimes \ell}$, we know that $|T_1\rangle$ unentagles from $|T \setminus T_1\rangle$. We extend this idea to show that even for a tensor of type states, applying $(Z^k \otimes I_{n-\lambda})^{\otimes \tilde{\ell}_i}$ on parts of each type state still unentangles each of them as long as all the type states are λ-prefix ℓ-fold collision-free type and their combined set is still λ-prefix ℓ-fold collision-free. Formally, we show the following: Let $\tilde{\ell}_1, \ldots, \tilde{\ell}_q \in \mathbb{N}$, and $t_1, \ldots, t_q \in \mathbb{N}$ such that $\sum_{i=1}^q \tilde{\ell}_i = \tilde{\ell}$ and $\sum_{i=1}^q t_i = t$. Then for any λ-prefix $\tilde{\ell}$-fold collision-free type T and any mutually disjoint sets $T_1, \ldots, T_q$ satisfying $\bigcup_{i=1}^q T_i = T$ and $|T_i| = t_i + \tilde{\ell}_i$ for all $i \in [q]$,

$$\mathop{\mathbb{E}}_{k \leftarrow \{0,1\}^n} \left[\bigotimes_{i=1}^{q} \left(\left(Z^k \otimes I_m\right)^{\otimes \ell_i} \otimes I_{n+m}^{\otimes t_i}\right) |T_i\rangle\langle T_i| \left(\left(Z^k \otimes I_m\right)^{\otimes \tilde{\ell}_i} \otimes I_{n+m}^{\otimes t_i}\right)\right]$$
$$= \bigotimes_{i=1}^{q} \mathop{\mathbb{E}}_{\substack{X_i \subset T_i \\ |X_i| = \tilde{\ell}_i}} \left[|X_i\rangle\langle X_i| \otimes |T_i \setminus X_i\rangle\langle T_i \setminus X_i|\right].$$

Hence, applying each layer $(Z^{k_i^b} \otimes I_{n-\lambda})$ unentagles all type states into two halfs. Hence, by repeated application, we get

$$\rho \approx_\varepsilon \mathop{\mathbb{E}}_{\substack{T \leftarrow [0:t+\ell]^N \\ T \text{ is } \lambda\text{-prefix } \ell\text{-fold collision-free}}} \mathop{\mathbb{E}}_{(T_1, T_2, \ldots, T_q, \hat{T})} \left[\bigotimes_{i=1}^{q} |T_i\rangle\langle T_i| \otimes |\hat{T}\rangle\langle\hat{T}|\right],$$

where $(T_1, T_2, \ldots, T_q, \hat{T})$ are sampled as follows: for $i = 1, 2, \ldots, q$, sample an ℓ_i-subset from $T \setminus (\bigcup_{j=1}^{i-1} T_j)$ uniformly and let $\hat{T} := T \setminus (\bigcup_{j=1}^{q} T_j)$. Over expectation over all λ-prefix ℓ-fold collision-free types states, this state is close to σ.

2.2 Quantum Bit Commitments

With t-copy PRSG in hand, we construct a statistically-hiding, statistically-binding commitment scheme in the CHS model. Our scheme draws inspiration from the quantum commitment scheme introduced in [MY21, MNY23] that builds quantum bit commitments from t-copy PRSG.

In particular, to commit to $b = 0$, the committer creates a superposition over all keys of the PRSG in the decommitment register and runs the PRSG in superposition over this register. The committer sets this as the commitment register. To commit to $b = 1$, the committer creates a maximally entangled state over the commitment and the decommitment register. Formally,

$$|\psi_0\rangle_{\mathsf{C}_i\mathsf{R}_i} := \frac{1}{\sqrt{2^\lambda}} \sum_{k\in\{0,1\}^\lambda} G_k(|\vartheta\rangle)_{\mathsf{C}_i} |k||0^{n-\lambda}\rangle_{\mathsf{R}_i}$$

and

$$|\psi_1\rangle_{\mathsf{C}_i\mathsf{R}_i} := \frac{1}{\sqrt{2^n}} \sum_{j\in\{0,1\}^n} |j\rangle_{\mathsf{C}_i} |j\rangle_{\mathsf{R}_i},$$

where, $(\mathsf{C}_1, \ldots, \mathsf{C}_p)$ is the commitment register and $(\mathsf{R}_1, \ldots, \mathsf{R}_p)$ is the reveal register.

To achieve hiding, our scheme relies on the pseudorandomness property of the PRSG. In particular, the commitment is very close to one where the keys are distinct for all $(\mathsf{C}_i, \mathsf{R}_i)$, in this case, one copy of PRS is indistinguishable from a maximally mixed state.[6]

Unlike the approach in [MY21], our construction is not of the canonical form [Yan2]. To achieve binding, the receiver performs multiple SWAP tests. In particular, we show that since the rank of the commitment registers is exponentially separated, multiple SWAP tests can distinguish between the two.

2.3 Black-Box Separations

LOCC Indistinguishability. The notion of LOCC indistinguishability is well-studied and is referred to as quantum data hiding by quantum information theorists [BDF+99, DLT02, EW02, Gea02, HLS05, MWW09, CLMO13, PNC14, CH14, CLM+14, HBAB19]. In this setting, there is a challenger, two (possibly entangled and mixed) bipartite quantum states ρ_{AB} and σ_{AB}, and a computationally unbounded, two-party distinguisher (Alice, Bob) who are spatially separated and without pre-shared entanglement. The challenger picks a quantum state from $\{\rho_{\mathsf{AB}}, \sigma_{\mathsf{AB}}\}$ uniformly at random and sends register A to Alice and register B to Bob respectively. The task of Alice and Bob is to distinguish whether they are given ρ_{AB} or σ_{AB} by performing local operations and communicating classically. We call such distinguishers *LOCC adversaries.*

We focus on the case where Alice and Bob each receive $t = \mathrm{poly}(\lambda)$ copies of $|\psi\rangle_{\mathsf{A}}$ and $|\phi\rangle_{\mathsf{B}}$, where $|\psi\rangle$ and $|\phi\rangle$ are either two identical or i.i.d. Haar states of length $n = \omega(\log(\lambda))$. Explicitly, the two input states are

$$\rho_{\mathsf{AB}} = \mathop{\mathbb{E}}_{|\psi\rangle\leftarrow\mathcal{H}_n} \left[|\psi\rangle\langle\psi|_{\mathsf{A}}^{\otimes t} \otimes |\psi\rangle\langle\psi|_{\mathsf{B}}^{\otimes t} \right],$$

[6] Note that this still needs multi-key security which is not trivial in the CHS model, since all the PRS generators share the same Haar state for randomness. But we prove that our construction satisfies multikey security.

$$\sigma_{\mathsf{AB}} = \mathop{\mathbb{E}}_{|\psi\rangle \leftarrow \mathcal{H}_n} \left[|\psi\rangle\langle\psi|_{\mathsf{A}}^{\otimes t} \right] \otimes \mathop{\mathbb{E}}_{|\phi\rangle \leftarrow \mathcal{H}_n} \left[|\phi\rangle\langle\phi|_{\mathsf{B}}^{\otimes t} \right].$$

Note that if global measurements are allowed, performing SWAP tests can easily distinguish them. As one of our main technical contributions, we show that for any LOCC adversary, the advantage of distinguishing ρ_{AB} from σ_{AB} is negligible in λ. Before we explain the proof, we compare our theorem with [Har23, Theorem 8]. In short, the theorems are incomparable. Our setting is stronger in the sense that the LOCC adversary both obtain polynomial copies of the input, while [Har23, Theorem 8] studies the single-copy setting. However, [Har23, Theorem 8] is more general since it holds for a family of input states, whereas the input in our setting is fixed to ρ_{AB} and σ_{AB}, which are belong to the family.

Toward the proof, we start by using the following common technique in proving LOCC indistinguishability: the set of LOCC measurements is a (proper) subset of the set of all positive partial transpose (PPT) measurements [CLM+14]. Hence, it is sufficient to upper bound the maximum distinguishing advantage over two-outcome PPT measurements, i.e., $\{M_{\mathsf{AB}}, I_{\mathsf{AB}} - M_{\mathsf{AB}}\}$ such that $0 \preceq M_{\mathsf{AB}} \preceq I_{\mathsf{AB}}$ and $0 \preceq M_{\mathsf{AB}}^{\Gamma_B} \preceq I_{\mathsf{AB}}$, where $M_{\mathsf{AB}}^{\Gamma_B}$ denote the partial transpose of M_{AB} with respect to B. Next, from the basic properties of partial transpose and trace norm, we show that the distinguishing advantage is bounded by the trace norm between $\rho_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}}$ and $\sigma_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}}$.

The most technical part of the proof is to upper bound the quantity $\left\| \rho_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} - \sigma_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} \right\|_1$. We point out that the partial transpose of a density matrix might *not* be a positive semidefinite matrix. Our first step is to expand ρ_{AB} and σ_{AB} in the *type basis* as follows:

$$\rho_{\mathsf{AB}} = \mathop{\mathbb{E}}_{T \leftarrow [0:2t]^d} \left[|T\rangle\langle T|_{\mathsf{AB}} \right],$$

$$\sigma_{\mathsf{AB}} = \mathop{\mathbb{E}}_{S_A \leftarrow [0:t]^d} \left[|S_A\rangle\langle S_A|_{\mathsf{A}} \right] \otimes \mathop{\mathbb{E}}_{S_B \leftarrow [0:t]^d} \left[|S_B\rangle\langle S_B|_{\mathsf{B}} \right],$$

where $d := 2^n$. Next, we further conditioned on the events that (1) T, S_A and S_B each have no repeated elements (2) S_A and S_B have no identical elements. From the collision bound, doing so only incurs an additional error of $O(t^2/d) = \mathsf{negl}(\lambda)$. Therefore, we can now treat T, S_A and S_B as *sets*. It suffices to prove that $\left\| \tilde{\rho}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} - \tilde{\sigma}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} \right\|_1$ is negligible in λ, where

$$\tilde{\rho}_{\mathsf{AB}} := \mathop{\mathbb{E}}_{T \leftarrow \binom{[d]}{2t}} \left[|T\rangle\langle T|_{\mathsf{AB}} \right],$$

$$\tilde{\sigma}_{\mathsf{AB}} := \mathop{\mathbb{E}}_{\substack{S_A, S_B \leftarrow \binom{[d]}{t}: \\ S_A \cap S_B = \emptyset}} \left[|S_A\rangle\langle S_A|_{\mathsf{A}} \otimes |S_B\rangle\langle S_B|_{\mathsf{B}} \right].$$

Observe that the $\tilde{\sigma}_{\mathsf{AB}}^{\Gamma_B} = \tilde{\sigma}_{\mathsf{AB}}$. To obtain a simpler expression of $\tilde{\rho}_{\mathsf{AB}}^{\Gamma_B}$, we rely on the following useful identity for bi-partitioning the type states:

$$|T\rangle_{\mathsf{AB}} = \sum_{X \in \binom{T}{t}} \frac{1}{\sqrt{\binom{2t}{t}}} |T \setminus X\rangle_{\mathsf{A}} \otimes |X\rangle_{\mathsf{B}}.$$

Hence, the partial transpose of $\tilde{\rho}_{\mathsf{AB}}$ can be written as

$$\tilde{\rho}_{\mathsf{AB}}^{\Gamma_B} = \mathop{\mathbb{E}}_{T \leftarrow \binom{[d]}{2t}} \left[\frac{1}{\binom{2t}{t}} \sum_{X,Y \in \binom{T}{t}} |T \setminus X\rangle\langle T \setminus Y|_{\mathsf{A}} \otimes |Y\rangle\langle X|_{\mathsf{B}} \right].$$

If $X = Y$, then the term is the tensor product of two *disjoint* sets $|T \setminus X\rangle\langle T \setminus X|_{\mathsf{A}} \otimes |X\rangle\langle X|_{\mathsf{B}}$. Such a term will be canceled out by the corresponding term in $\tilde{\sigma}_{\mathsf{AB}}^{\Gamma_B}$ since they have equal coefficients. Therefore, the difference between them is the following matrix with mismatched X and Y:

$$\tilde{\rho}_{\mathsf{AB}}^{\Gamma_B} - \tilde{\sigma}_{\mathsf{AB}}^{\Gamma_B} = \mathop{\mathbb{E}}_{T \leftarrow \binom{[d]}{2t}} \left[\frac{1}{\binom{2t}{t}} \sum_{X,Y \in \binom{T}{t} : X \neq Y} |T \setminus X\rangle\langle T \setminus Y|_{\mathsf{A}} \otimes |Y\rangle\langle X|_{\mathsf{B}} \right].$$

We continue to simplify it by applying a double-counting argument. Every tuple of sets (T, X, Y) uniquely determines a tuple of mutually disjoint sets (C, I, X', Y') satisfying $C = T \setminus (X \cup Y)$ (C for the complement of $X \cup Y$), $I = X \cap Y$ (I for intersection), $X' = X \setminus I$ and $Y' = Y \setminus I$. Hence, $T \setminus X = C \uplus Y'$, $Y = I \uplus Y'$, $T \setminus Y = C \uplus X'$, and $X = I \uplus X'$ where $\uplus$ denotes the disjoint union. By further classifying the summands according to $s := |C| = |I| \in \{0, 1, \ldots, t-1\}$ (note that then $|X'| = |Y'| = t - s$), we have

$$\begin{aligned}
&\left\| \tilde{\rho}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} - \tilde{\sigma}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} \right\|_1 \\
&= \frac{1}{\binom{d}{2t}\binom{2t}{t}} \left\| \sum_{s=0}^{t-1} \sum_{C \in \binom{[d]}{s}} \sum_{I \in \binom{[d] \setminus C}{s}} \sum_{\substack{X', Y' \in \binom{[d] \setminus (C \uplus I)}{t-s} : \\ X' \cap Y' = \emptyset}} |C \uplus Y'\rangle_{\mathsf{A}} |I \uplus Y'\rangle_{\mathsf{B}} \langle C \uplus X'|_{\mathsf{A}} \langle I \uplus X'|_{\mathsf{B}} \right\|_1 \\
&\leq \frac{1}{\binom{d}{2t}\binom{2t}{t}} \sum_{s=0}^{t-1} \sum_{C \in \binom{[d]}{s}} \sum_{I \in \binom{[d] \setminus C}{s}} \underbrace{\left\| \sum_{\substack{X', Y' \in \binom{[d] \setminus (C \uplus I)}{t-s} : \\ X' \cap Y' = \emptyset}} |C \uplus Y'\rangle_{\mathsf{A}} |I \uplus Y'\rangle_{\mathsf{B}} \langle C \uplus X'|_{\mathsf{A}} \langle I \uplus X'|_{\mathsf{B}} \right\|_1}_{=: K_{C,I}},
\end{aligned}$$

where the inequality follows from the triangle inequality. We observe that the matrix $K_{C,I}$ has the same structure as the adjacency matrix of *Kneser graphs*. Here, we recall the definition of Kneser graphs. For $v, k \in \mathbb{N}$, the Kneser graph $K(v, k)$ is the graph whose vertices correspond to the k-element subsets of the set $[v]$, and two vertices are adjacent if and only if the two corresponding sets are disjoint. Therefore, for every (C, I), the matrix $K_{C,I}$ is isospectral to the adjacency matrix of the Kneser graph $K(d - |C| - |I|, t - |I|)$. Finally, we employ the well-studied spectral property of Kneser graphs as a black box to obtain an $O(t^2/d) = \mathsf{negl}(\lambda)$ upper bound for $\left\| \tilde{\rho}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} - \tilde{\sigma}_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} \right\|_1$.

Furthermore, we show the tightness of the theorem by constructing an optimal LOCC distinguisher that achieves the same advantage. The strategy is simple: Alice and Bob each individually measure every copy of their input in the computational basis and obtain a total of $2t$ outcomes. Then, they output 1 if there is any collision among these $2t$ outcomes.

Impossibility Results in the CHS Model. With the LOCC Haar indistinguishability theorem in hand, we investigate the limits of the CHS model when the communication between the parties is classical. We show that the several impossibility results of information-theoretically secure schemes in the plain model can be generically lifted to the CHS model, even when the adversary does not receive any common Haar state. We emphasize that there is no classical counterpart in the CRS model. If the adversary is not given the CRS, then many information-theoretically secure schemes do exist, such as key agreements.

As common in proving impossibilities, our approach is to convert schemes in the CHS model to those in the plain model. The transform is simple: in the new scheme, the parties each sample polynomially many copies of the Haar state *independently* and run the original scheme. Crucially, despite the inconsistency in their Haar states, the new scheme still satisfies completeness thanks to the LOCC Haar indistinguishability. A caveat is that sampling Haar states is time-inefficient. However, since the impossibilities in the plain model are still valid if the (honest) algorithms in the scheme are time-inefficient, doing so is acceptable for the sake of showing impossibilities.

Separation Results. We separate many important primitives from $(\lambda, \omega(\log(\lambda)))$-PRSG. Since $(\lambda, \omega(\log(\lambda)))$-PRSGs do not exist in the CHS model, we need to "strengthen" the oracle in order to prove separations. For every security parameter $\lambda \in \mathbb{N}$, we define the oracle as $\{G_k\}_{k \in \{0,1\}^\lambda}$ where each G_k is an isometry that takes no input and outputs an i.i.d. Haar state $|\psi_k\rangle$.

Relative to this oracle, the implementation of the PRSG is straightforward: the output on k of any length $\lambda \in \mathbb{N}$ is $|\psi_k\rangle$. The security directly follows from the hardness of unstructured search. To prove the non-existence of QCCC schemes, we employ a two step approach. First, showing that a scheme with respect to this oracle can be transformed to schemes with respect to a much weaker oracle. Second, showing that this much weaker oracle does not give much extra power over the plain model. Formally: First, similar to the previous section, we show that due to the LOCC indistinguishability, the parties can sample all "large" quantum states on their own, and the correctness and security is only "polynomially" affected[7]. This means that any scheme with respect this oracle can be turned into a scheme with respect to an oracle with only short (constant times logarithmic) Haar states. Second, for short (constant times logarithmic) quantum states, we show that this oracle does not give much extra power since an adversary can learn the oracle completely. This is because for short-enough states, the adversary can run tomography on polynomial queries and learn the state with up to inverse polynomial error. Hence, the adversary can simulate

[7] Since the Haar indistinguishability has a factor of $O(t^2/d)$, as long as t^2/d is inverse-polynomial, we do not incur a lot of loss.

both parties post-selecting on a transcript to learn any secret[8]. This means that any scheme secure in the presence of this oracle can be transformed into another scheme that is secure in the plain model.

Lastly, we observe that by considering a generalized oracle, namely $\{G_{k,x}\}_{k,x\in\{0,1\}^\lambda}$, we can show that (classically accessible) PRFSGs with super-logarithmic input length exist. We can extend the impossibility of QCCC commitments to hold in the presence of the generalized oracle as well. Thus, we can separate PRFS and QCCC commitments.

3 Preliminaries

We denote the security parameter by λ. We assume that the reader is familiar with the fundamentals of quantum computing covered in [NC10].

3.1 Notation

- We use $[n]$ to denote $\{1,\ldots,n\}$ and $[0:n]$ to denote $\{0,1,\ldots,n\}$.
- For any finite set T and any integer $0 \leq k \leq |T|$, we denote by $\binom{T}{k}$ the set of all k-size subsets of T.
- For any finite set T, we use the notation $x \leftarrow T$ to indicate that x is sampled uniformly from T.
- We denote by S_t the symmetric group of degree t.
- For any set A and $t \in \mathbb{N}$, we denote by A^t the t-fold Cartesian product of A.
- For $\sigma \in S_t$ and $\mathbf{v} = (v_1,\ldots,v_t)$, we define $\sigma(\mathbf{v}) := (v_{\sigma(1)},\ldots,v_{\sigma(t)})$.
- We denote by $\mathcal{D}(H)$ the set of density matrices in the Hilbert space H.
- Let $\rho_{AB} \in \mathcal{D}(H_A \otimes H_B)$, by $\mathrm{Tr}_B(\rho_{AB}) \in \mathcal{D}(H_A)$ we denote the reduced density matrix by taking partial trace over B.
- We denote by $\mathsf{TD}(\rho,\rho') := \frac{1}{2}\|\rho - \rho'\|_1$ the trace distance between quantum states ρ,ρ', where $\|X\|_1 = \mathrm{Tr}(\sqrt{X^\dagger X})$ denotes the trace norm.
- For any matrices A, B, we write $A \preceq B$ to indicate that $B - A$ is positive semi-definite.
- For any Hermitian matrix O, the trace norm of O has the following variational definition:

$$\|O\|_1 = \max_{-I\preceq M\preceq I} \mathrm{Tr}(MO).$$

  Furthermore, if $\mathrm{Tr}(O) = 0$ then $\|O\|_1 = 2 \cdot \max_{0\preceq M\preceq I} \mathrm{Tr}(MO)$.
- We denote the Haar measure over n qubits by $\mathcal{H}_n$.
- For any matrix $M_{\mathsf{AB}} = \sum_{i,j,k,\ell} \alpha_{ijk\ell}|i\rangle\langle j|_{\mathsf{A}} \otimes |k\rangle\langle \ell|_{\mathsf{B}}$ on registers (A,B), by $M_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}}$ we denote its *partial transpose* with respect to register B, i.e., $M_{\mathsf{AB}}^{\Gamma_{\mathsf{B}}} = \sum_{i,j,k,\ell} \alpha_{ijk\ell}|i\rangle\langle j|_{\mathsf{A}} \otimes |\ell\rangle\langle k|_{\mathsf{B}}$.[9]

[8] Note that since the adversary does not need to be efficient, as long as they have the description of this oracle, they can post-select on the transcript.

[9] Note that the (partial) transpose operation needs to be defined with respect to to an orthogonal basis. Throughout this work, it is always defined with respect to to the computational basis.

3.2 Common Haar State Model

The Common Haar State (CHS) model is related to the Common Reference Quantum State (CRQS) model [MNY23]. In this model, all parties receive polynomially many copies of a *single* quantum state sampled from the Haar distribution. Recently, another work of Chen et al. [CCS24] studied a similar model called the Common Haar Random State (CHRS) model. In the CHRS model, every party receives polynomially many copies of *polynomially many* i.i.d. Haar states.

We define another variant of the CHS model called the *Keyed* Common Haar State Model. In this model, all parties (once the security parameter is set to λ) have access to the oracle (called the *Keyed* Common Haar State Oracle) $G^\lambda := \{G_k\}_{k\in\{0,1\}^\lambda}$ as follows. For every $k \in \{0,1\}^\lambda$, the oracle G_k is a Haar isometry that maps any state $|\psi\rangle$ to $|\psi\rangle|\vartheta_k\rangle$, where $|\vartheta_k\rangle$ is a Haar state of length $n(\lambda) = \omega(\log(\lambda))$.

While the above variant is harder to instantiate (hence not useful for constructions), is a natural candidate for black-box separations as seen is Sect. 9.

Pseudorandom State (PRS) Generators in the CHS Model

Definition 1 (Statistically Secure (λ, n, ℓ)-Pseudorandom State Generators in the CHS Model). *We say that a QPT algorithm G is a* statistically secure (λ, n, ℓ)-pseudorandom state generator (PRSG) *in the CHS model if the following holds:*

- ***State Generation**: For any $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^\lambda$, the algorithm G_k (where G_k denotes $G(k,\cdot)$) is a quantum channel such that for every $n(\lambda)$-qubit state $|\vartheta\rangle$,*
$$G_k(|\vartheta\rangle\langle\vartheta|) = |\vartheta_k\rangle\langle\vartheta_k|,$$
for some $n(\lambda)$-qubit state $|\vartheta_k\rangle$. We sometimes write $G_k(|\vartheta\rangle)$ for brevity.[10]
- ***ℓ-copy Pseudorandomness**: For any polynomial $t(\cdot)$ and any non-uniform, unbounded adversary $A = \{A_\lambda\}_{\lambda\in\mathbb{N}}$, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that:*
$$\Bigg| \Pr_{\substack{k\leftarrow\{0,1\}^\lambda \\ |\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)}}} \left[A_\lambda\left(G_k(|\vartheta\rangle)^{\otimes\ell(\lambda)} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)}\right) = 1\right] - \Pr_{\substack{|\varphi\rangle\leftarrow\mathcal{H}_{n(\lambda)} \\ |\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)}}} \left[A_\lambda\left(|\varphi\rangle\langle\varphi|^{\otimes\ell(\lambda)} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)}\right) = 1\right] \Bigg| \le \mathsf{negl}(\lambda).$$

[10] More generally, the generation algorithm could take multiple copies of the common Haar state as input or output a state of different size compared to the common Haar state. Here, we focus on a restricted class of generators that only require a single copy of the common Haar state as input, and the output of the generator matches the size of the common Haar states.

If G satisfies ℓ-copy pseudorandomness for every polynomial $\ell(\cdot)$ then we drop ℓ from the notation and simply denote it to be a (λ, n)-PRSG.

We define a stronger definition below called *multi-key ℓ-copy PRS generators.* Looking ahead, our construction of PRS in Sect. 4.2 satisfies this definition.

Definition 2 (Multi-key Statistically Secure (λ, n, ℓ)-Pseudorandom State Generators in the CHS Model). *We say that a QPT algorithm G is a* multi-key statistically secure (λ, n, ℓ)-pseudorandom state generator *in the CHS model if the following holds:*

- ***State Generation**: For any $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^\lambda$, the algorithm G_k (where G_k denotes $G(k, \cdot)$) is a quantum channel such that for every $n(\lambda)$-qubit state $|\vartheta\rangle$,*

$$G_k(|\vartheta\rangle\langle\vartheta|) = |\vartheta_k\rangle\langle\vartheta_k|,$$

  for some $n(\lambda)$-qubit state $|\vartheta_k\rangle$. We sometimes write $G_k(|\vartheta\rangle)$ for brevity.
- ***Multi-key ℓ-copy Pseudorandomness**: For any polynomial $t(\cdot)$, $p(\cdot)$ and any non-uniform, unbounded adversary $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that:*

$$\left| \Pr_{\substack{k_1,\ldots,k_{p(\lambda)} \leftarrow \{0,1\}^\lambda \\ |\vartheta\rangle \leftarrow \mathcal{H}_{n(\lambda)}}} \left[A_\lambda \left(\bigotimes_{i=1}^{p(\lambda)} G_{k_i}(|\vartheta\rangle)^{\otimes \ell(\lambda)} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)} \right) = 1 \right] \right.$$
$$\left. - \Pr_{\substack{|\varphi_1\rangle,\ldots,|\varphi_{p(\lambda)}\rangle \leftarrow \mathcal{H}_{n(\lambda)} \\ |\vartheta\rangle \leftarrow \mathcal{H}_{n(\lambda)}}} \left[A_\lambda \left(\bigotimes_{i=1}^{p(\lambda)} |\varphi_i\rangle\langle\varphi_i|^{\otimes \ell(\lambda)} \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)} \right) = 1 \right] \right| \leq \mathsf{negl}(\lambda).$$

If G satisfies multi-key ℓ-copy pseudorandomness for every polynomial $\ell(\cdot)$ then we drop ℓ from the notation and simply denote it to be a multi-key (λ, n)-PRSG.

Remark 1. Note that in the plain model, PRS implies multi-key PRS because the pseudorandom state generator does not share randomness for different keys. It is not clear whether this holds in the CHS model as the different executions of the pseudorandom state generator share the same common Haar state.

Pseudorandom Function-like State (PRFS) Generators in the CHS Model

Definition 3 (Statistical Selectively Secure (λ, m, n, ℓ)-PRFS Generators). *We say that a QPT algorithm G is a* statistical selectively secure (λ, m, n, ℓ)-PRFS generator *in the CHS model if the following holds:*

- ***State Generation**: For any $\lambda \in \mathbb{N}$, $k \in \{0,1\}^\lambda$ and $x \in \{0,1\}^{m(\lambda)}$, where $m(\lambda)$ is the input length, the algorithm $G_{k,x}$ (where $G_{k,x}$ denotes $G(k, x, \cdot)$) is a quantum channel such that for every $n(\lambda)$-qubit state $|\vartheta\rangle$,*

$$G_{k,x}(|\vartheta\rangle\langle\vartheta|) = |\vartheta_{k,x}\rangle\langle\vartheta_{k,x}|,$$

  for some $n(\lambda)$-qubit state $|\vartheta_{k,x}\rangle$. We sometimes write $G_{k,x}(|\vartheta\rangle)$ or $G_k(x, |\vartheta\rangle)$ for brevity.

- ***ℓ-query Selective Security**: For any polynomial $t(\cdot)$, any non-uniform, unbounded adversary $A = \{A_\lambda\}_{\lambda\in\mathbb{N}}$, and any tuple of (possibly repeated) $m(\lambda)$-bit indices $(x_1,\dots,x_{\ell(\lambda)})$, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that for all $\lambda\in\mathbb{N}$,*

$$\left| \underset{k\leftarrow\{0,1\}^{\lambda},|\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)}}{Pr} \left[A_\lambda\left(x_1,\dots,x_{\ell(\lambda)}, \bigotimes_{i=1}^{\ell(\lambda)} G(k,x_i,|\vartheta\rangle)\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)}\right)=1\right] - \underset{\substack{\forall x\in\{0,1\}^{m(\lambda)},\ |\varphi_x\rangle\leftarrow\mathcal{H}_{n(\lambda)},\\ |\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)}}}{Pr}\left[A_\lambda\left(x_1,\dots,x_{\ell(\lambda)}, \bigotimes_{i=1}^{\ell(\lambda)} |\varphi_{x_i}\rangle\langle\varphi_{x_i}|\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t(\lambda)}\right)=1\right]\right| \le \mathsf{negl}(\lambda).$$

If G satisfies ℓ-query selective security for every polynomial $\ell(\cdot)$, we drop ℓ from the notation and say that G is a (λ,m,n)-PRFS generator.

Quantum Commitments in the CHS Model

Definition 4 (Quantum Commitments in the CHS Model). *A (non-interactive) quantum commitment scheme in the CHS model is given by a tuple of the committer C and receiver R parameterized by a polynomial $p(\cdot)$, both of which are uniform QPT algorithms. Let $|\vartheta\rangle$ be the $n(\lambda)$-qubit common Haar state. The scheme is divided into two phases: the commit phase, and the reveal phase as follows:*

- *Commit phase: C takes $|\vartheta\rangle^{\otimes p(\lambda)}$ and a bit $b\in\{0,1\}$ to commit as input, generates a quantum state on registers C and R, and sends the register C to R.*
- *Reveal phase: C sends b and the register R to R. R takes $|\vartheta\rangle^{\otimes p(\lambda)}$ and $(b,\mathsf{C},\mathsf{R})$ given by C as input, and outputs b if it accepts and otherwise outputs $\bot$.*

Definition 5 (Poly-Copy Statistical Hiding). *A quantum commitment scheme (C,R) in the CHS model satisfies* poly-copy statistical hiding *if for any non-uniform, unbounded malicious receiver $R^* = \{R^*_\lambda\}_{\lambda\in\mathbb{N}}$, and any polynomial $t(\cdot)$, there exists a negligible function* $\mathsf{negl}(\cdot)$ *such that*

$$\left| \Pr\left[R^*_\lambda(|\vartheta\rangle^{\otimes t(\lambda)}, \mathrm{Tr}_{\mathsf{R}}(\sigma_{\mathsf{CR}}))=1 : \substack{|\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)},\\ \sigma_{\mathsf{CR}}\leftarrow C_,(|\vartheta\rangle^{\otimes p(\lambda)},0)}\right] - \Pr\left[R^*_\lambda(|\vartheta\rangle^{\otimes t(\lambda)}, \mathrm{Tr}_{\mathsf{R}}(\sigma_{\mathsf{CR}}))=1 : \substack{|\vartheta\rangle\leftarrow\mathcal{H}_{n(\lambda)},\\ \sigma_{\mathsf{CR}}\leftarrow C_,(|\vartheta\rangle^{\otimes p(\lambda)},1)}\right]\right| \le \mathsf{negl}(\lambda),$$

where $C_,$ is the commit phase of C.

Definition 6 (Statistical Sum-Binding). *A quantum commitment scheme (C,R) in the CHS model satisfies* statistical sum-binding *if the following holds.*

For any pair of non-uniform, unbounded malicious senders C_0^* and C_1^* that take $|\vartheta\rangle^{\otimes T(\lambda)}$ for arbitrary large $T(\cdot)$ as input and work in the same way in the commit phase, if we let p_b to be the probability that R accepts the revealed bit b in the interaction with C_b^* for $b \in \{0,1\}$, then we have

$$p_0 + p_1 \leq 1 + \mathsf{negl}(\lambda).$$

3.3 Symmetric Subspaces, Type States, and Haar States

The proofs of facts and lemmas stated in this subsection can be found in [Har13]. Let $\mathbf{v} = (v_1, \ldots, v_t) \in A^t$ for some finite set A. Let $|A| = N$. Define $\mathsf{type}(\mathbf{v}) \in [0:t]^N$ to be the *type vector* such that the i^{th} entry of $\mathsf{type}(\mathbf{v})$ equals the number of occurrences of $i \in [N]$ in $\mathbf{v}$.[11] In this work, by $T \in [0:t]^N$ we implicitly assume that $\sum_{i\in[N]} T_i = t$. For $T \in [0:t]^N$, we denote by $\mathsf{mset}(T)$ the *multiset* uniquely determined by T. That is, the multiplicity of $i \in \mathsf{mset}(T)$ equals T_i for all $i \in [N]$. We write $T \leftarrow [0:t]^N$ to mean sampling T uniformly from $[0:t]^N$ conditioned on $\sum_{i\in[N]} T_i = t$. We write $\mathbf{v} \in T$ to mean $\mathbf{v} \in A^t$ satisfies $\mathsf{type}(\mathbf{v}) = T$.

In this work, we will focus on *collision-free* types T which satisfy $T_i \in \{0,1\}$ for all $i \in [N]$. A collision-free type T can be naturally treated as a *set* and we write $\mathbf{v} \leftarrow T$ to mean sampling a uniform $\mathbf{v}$ conditioned on $\mathsf{type}(\mathbf{v}) = T$.

Definition 7 (Type States). *Let $T \in [0:t]^N$, we define the* type states*:*

$$|T\rangle := \sqrt{\frac{\prod_{i\in[N]} T_i!}{t!}} \sum_{\mathbf{v}\in T} |\mathbf{v}\rangle.$$

If T is collision-free, then it can be simplified to

$$|T\rangle = \frac{1}{\sqrt{t!}} \sum_{\mathbf{v}\in T} |\mathbf{v}\rangle.$$

Furthermore, it has the following useful expression

$$|T\rangle\langle T| = \frac{1}{t!} \sum_{\mathbf{v},\mathbf{u}\in T} |\mathbf{v}\rangle\langle \mathbf{u}| = \mathop{\mathbb{E}}_{\mathbf{v}\leftarrow T}\left[\sum_{\sigma\in S_t} |\mathbf{v}\rangle\langle \sigma(\mathbf{v})|\right]. \tag{1}$$

Lemma 1 (Average of Copies of Haar-Random States). *For all $N, t \in \mathbb{N}$, we have*

$$\mathop{\mathbb{E}}_{|\vartheta\rangle\leftarrow\mathcal{H}(\mathbb{C}^N)} |\vartheta\rangle\langle\vartheta|^{\otimes t} = \mathop{\mathbb{E}}_{T\leftarrow[0:t]^N} |T\rangle\langle T|.$$

[11] We identify $[0:t]^N$ as $[0:t]^A$.

3.4 Quantum Black-Box Reductions

We recall the definition of fully black-box reductions [RTV04,BBF13] and their quantum analogue. The definitions below are taken verbatim from [HY20].

Definition 8 (Quantum Primitives). *A quantum primitive $\mathcal{P}$ is a pair $(\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$, where $\mathcal{F}_\mathcal{P}$ is a set of quantum algorithms $\mathcal{I}$, and $\mathcal{R}_\mathcal{P}$ is a relation over pairs $(\mathcal{I}, \mathcal{A})$ of quantum algorithms $\mathcal{I} \in \mathcal{F}_\mathcal{P}$ and $\mathcal{A}$. A quantum algorithm $\mathcal{I}$ implements $\mathcal{P}$ or is an implementation of $\mathcal{P}$ if $\mathcal{I} \in \mathcal{F}_\mathcal{P}$. If $\mathcal{I} \in \mathcal{F}_\mathcal{P}$ is efficient, then $\mathcal{I}$ is an efficient implementation of $\mathcal{P}$. A quantum algorithm $\mathcal{A}$ $\mathcal{P}$-breaks $\mathcal{I} \in \mathcal{F}_\mathcal{P}$ if $(\mathcal{I}, \mathcal{A}) \in \mathcal{R}_\mathcal{P}$. A secure implementation of $\mathcal{P}$ is an implementation $\mathcal{I}$ of $\mathcal{P}$ such that no efficient quantum algorithm $\mathcal{P}$-breaks $\mathcal{I}$. The primitive $\mathcal{P}$ quantumly exists if there exists an efficient and secure implementation of $\mathcal{P}$.*

Definition 9 (Quantum Primitives Relative to Oracle). *Let $\mathcal{P} = (\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ be a quantum primitive, and O be a quantum oracle. An oracle quantum algorithm $\mathcal{I}$ implements $\mathcal{P}$ relative to O or is an implementation of $\mathcal{P}$ relative to O if $\mathcal{I}^O \in \mathcal{F}_\mathcal{P}$. If $\mathcal{I}^O \in \mathcal{F}_\mathcal{P}$ is efficient, then $\mathcal{I}$ is an efficient implementation of $\mathcal{P}$ relative to O. A quantum algorithm $\mathcal{A}$ $\mathcal{P}$-breaks $\mathcal{I} \in \mathcal{F}_\mathcal{P}$ relative to O if $(I^O, \mathcal{A}^O) \in \mathcal{R}_\mathcal{P}$. A secure implementation of $\mathcal{P}$ is an implementation $\mathcal{I}$ of $\mathcal{P}$ relative to O such that no efficient quantum algorithm $\mathcal{P}$-breaks $\mathcal{I}$ relative to O. The primitive $\mathcal{P}$ quantumly exists relative to O if there exists an efficient and secure implementation of $\mathcal{P}$ relative to O.*

Definition 10 (Quantum Fully Black-Box Reductions). *A pair (C, S) of efficient oracle quantum algorithms is a* quantum fully-black-box reduction *from a quantum primitive $\mathcal{P} = (\mathcal{F}_\mathcal{P}, \mathcal{R}_\mathcal{P})$ to a quantum primitive $\mathcal{Q} = (\mathcal{F}_\mathcal{Q}, \mathcal{R}_\mathcal{Q})$ if the following two conditions are satisfied:*

1. ***(Correctness.)*** *For every implementation $\mathcal{I} \in \mathcal{F}_\mathcal{Q}$, we have $C^\mathcal{I} \in \mathcal{F}_\mathcal{P}$.*
2. ***(Security.)*** *For every implementation $\mathcal{I} \in \mathcal{F}_\mathcal{Q}$ and every quantum algorithm $\mathcal{A}$, if $\mathcal{A}$ $\mathcal{P}$-breaks $C^\mathcal{I}$, then $S^{\mathcal{A},\mathcal{I}}$ $\mathcal{Q}$-breaks $\mathcal{I}$.*

4 Warmup: Statistical Stretch PRS Generators inthe CHS Model

We present a construction of multi-key PRS generator with statistical security in the CHS model.

Theorem 7. *There exists a multi-key (λ, n, ℓ)-statistical PRS generator in the CHS model, where $n \geq \lambda$ and $\ell = O(\lambda / \log(\lambda)^{1+\varepsilon})$ for any constant $\varepsilon > 0$.*

The proof can be found in Sect. 4.2. Later, we prove the optimality of our construction in Sect. 4.3. Specifically, we show that any (λ, n, ℓ)-statistical PRS generator cannot simultaneously satisfy $n = \omega(\log(\lambda))$ and $\ell = \Omega(\lambda / \log(\lambda))$.

4.1 Useful Lemmas

At a high level, the proof follows the template of [AGQY22, AGKL23]: we do the analysis in the symmetric subspace. First, we identify a nice property of type vectors such that (1) a randomly sampled type satisfies this property with overwhelming probability and (2) the PRS generation algorithm behaves well on every type state having this property. We identify these type vectors as ℓ-*fold collision-free* types (which are a generalization of distinct types [AGQY22, AGKL23]).

Definition 11 (ℓ-Fold n-Prefix Collision-Free Types). *Let $n, m, t, \ell \in \mathbb{N}$ such that $t \geq \ell$ and $T \in [0:t]^{2^{n+m}}$ is a type vector. We say that T is* ℓ-fold n-prefix collision-free *if for all pairs of ℓ-subsets*[12] *$\mathcal{S}, \mathcal{T} \subseteq \mathsf{mset}(T)$, the first n bits of $\bigoplus_{x \in \mathcal{S}} x \in \{0,1\}^{n+m}$ is identical to that of $\bigoplus_{y \in \mathcal{T}} y \in \{0,1\}^{n+m}$ if and only if $\mathcal{S} = \mathcal{T}$. We define $\mathcal{I}_{n,m}^{(\ell)}(t) := \{T \in [0:t]^{2^{n+m}} : T \text{ is } \ell\text{-fold } n\text{-prefix collision-free}\}$ as the set of all ℓ-fold n-prefix collision-free type vectors.*

When $t > \ell$, one can easily verify that ℓ-fold n-prefix collision-freeness implies the standard collision-freeness. Also note that when $t > 2\ell$, ℓ-fold n-prefix collision-freeness implies i-fold n-prefix collision-freeness for all $i \leq \ell$.

Next, we show that a random type is ℓ-fold n-prefix collision-free with high probability.

Lemma 2. $\Pr_{T \leftarrow [0:t]^{2^{n+m}}}[T \in \mathcal{I}_{n,m}^{(\ell)}(t)] = 1 - O(t^{2\ell}/(2^n - 2\ell))$.

Proof. First, sampling $T \leftarrow [0:t]^{2^{n+m}}$ uniformly is $O(t^2/2^{n+m})$-close to sampling a uniform collision-free T from $[0:t]^{2^{n+m}}$ by the collision bound. Furthermore, sampling a uniform collision-free T from $[0:t]^{2^{n+m}}$ is equivalent to sampling t elements $x_1, x_2, \ldots, x_t$ one by one from $\{0,1\}^{n+m}$ conditioned on them being distinct and setting T such that $\mathsf{mset}(T) = \{x_1, \ldots, x_t\}$. Hence, it suffices to show that sampling t elements $x_1, x_2, \ldots, x_t$ one by one from $\{0,1\}^{n+m}$ conditioned on them being distinct results in an ℓ-fold n-prefix collision-free set with probability $1 - O(t^{2\ell}/2^n)$.
For any two distinct ℓ-subsets of indices $\mathcal{S} \neq \mathcal{T} \subseteq [t]$, let $\mathsf{Bad}_{\mathcal{S},\mathcal{T}}$ denote the event that the first n bits of $\bigoplus_{i \in \mathcal{S}} x_i$ is the same as that of $\bigoplus_{j \in \mathcal{T}} x_j$. Then the following holds:

$$\Pr\left[\mathsf{Bad}_{\mathcal{S},\mathcal{T}} : {}^{x_1, x_2, \ldots, x_t \leftarrow \{0,1\}^{n+m}}_{x_1, x_2, \ldots, x_t \text{ are distinct}}\right] = O(1/(2^n - 2\ell)).$$

This is because we can first sample $|\mathcal{S} \cup \mathcal{T}| - 1$ elements (in $\mathcal{S} \cup \mathcal{T}$) except one with indices in $\mathcal{S} \setminus \mathcal{T}$. Then $\mathsf{Bad}_{\mathcal{S},\mathcal{T}}$ occurs only if the first n bits of the last sample is equal to the first n bits of the bitwise XOR of all other elements in $\mathcal{S}$ with all elements in $\mathcal{T}$, which happens with probability at most $O(1/(2^n - 2\ell))$.

By a union bound, we have $T \in \mathcal{I}_{n,m}^{(\ell)}(t)$ with probability at least $1 - (O(t^2/2^{n+m}) + \binom{t}{\ell}^2 \cdot O(1/(2^n - 2\ell))) = 1 - O(t^{2\ell}/(2^n - 2\ell))$. □

[12] Here we allow the subsets to contain duplicate elements.

Finally, the following two lemmas show that applying random Pauli-Z on any ℓ-fold n-prefix collision-free type state is equivalent to a "classical" probabilistic process[13].

Lemma 3. *For any* $\mathbf{v} \in \{0,1\}^{(n+m)(t+\ell)}$ *such that* $\mathsf{type}(\mathbf{v}) \in \mathcal{I}_{n,m}^{(\ell)}(t+\ell)$ *and* $\sigma \in S_{t+\ell}$*, define*

$$A_{\mathbf{v},\sigma} := \mathop{\mathbb{E}}_{k \leftarrow \{0,1\}^n} \left[\left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) |\mathbf{v}\rangle\langle\sigma(\mathbf{v})| \left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) \right].$$

Then $A_{\mathbf{v},\sigma} = |\mathbf{v}\rangle\langle\sigma(\mathbf{v})|$ *if* σ *maps* $[\ell]$ *to* $[\ell]$*; otherwise,* $A_{\mathbf{v},\sigma} = 0$*.*

Proof. Suppose $\mathbf{v} = (v_1 || w_1, \ldots, v_{t+\ell} || w_{t+\ell}) \in \{0,1\}^{(n+m)(t+\ell)}$ with $v_i \in \{0,1\}^n$ and $w_i \in \{0,1\}^m$ for all $i \in [t]$. First, a direct calculation yields:

$$\begin{aligned} &\left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) |\mathbf{v}\rangle\langle\sigma(\mathbf{v})| \left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) \\ &= (-1)^{\langle k, \bigoplus_{i=1}^{\ell} (v_i \oplus v_{\sigma(i)}) \rangle} |\mathbf{v}\rangle\langle\sigma(\mathbf{v})|. \end{aligned}$$

Therefore, after averaging over k,

$$\begin{aligned} A_{\mathbf{v},\sigma} &= \mathop{\mathbb{E}}_{k \leftarrow \{0,1\}^n} \left[(-1)^{\langle k, \bigoplus_{i=1}^{\ell} (v_i \oplus v_{\sigma(i)}) \rangle} \right] |\mathbf{v}\rangle\langle\sigma(\mathbf{v})| \\ &= \begin{cases} |\mathbf{v}\rangle\langle\sigma(\mathbf{v})| & \text{if } \bigoplus_{i=1}^{\ell} (v_i \oplus v_{\sigma(i)}) = 0 \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Since $\mathsf{type}(\mathbf{v}) \in \mathcal{I}_{n,m}^{(\ell)}(t+\ell)$, the condition $\bigoplus_{i=1}^{\ell} v_i = \bigoplus_{i=1}^{\ell} v_{\sigma(i)}$ holds if and only if the two sets $\{1, 2, \ldots, \ell\}$ and $\{\sigma(1), \sigma(2), \ldots, \sigma(\ell)\}$ are identical. □

The following lemma lies at the technical heart of this section. It states that the action of applying random Z^k on ℓ-fold n-prefix collision-free types T[14] has the following "classical" probabilistic interpretation: the output is identically distributed to first uniformly sampling an ℓ-subset X from T and then generating $|X\rangle\langle X| \otimes |T \setminus X\rangle\langle T \setminus X|$.

Lemma 4. *For any* $T \in \mathcal{I}_{n,m}^{(\ell)}(t+\ell)$*,*

$$\begin{aligned} &\mathop{\mathbb{E}}_{k \leftarrow \{0,1\}^n} \left[\left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) |T\rangle\langle T| \left(\left(Z^k \otimes I_m \right)^{\otimes \ell} \otimes I_{n+m}^{\otimes t} \right) \right] \\ &= \mathop{\mathbb{E}}_{X \leftarrow \binom{T}{\ell}} \left[|X\rangle\langle X| \otimes |T \setminus X\rangle\langle T \setminus X| \right]. \end{aligned}$$

[13] We say that this is a "classical" probabilistic process because we can write the resulting density matrix as direct sum of matrices with classical descriptions with weights chosen by a completely classical process. This means that we can simualte this process by first doing a completely classical sampling process followed by a state preparation.

[14] Since T is collision-free, we will treat it as a set.

Proof. We first use the expression in Eq. (1) on the left-hand side:

$$\mathop{\mathbb{E}}_{\mathbf{v}\leftarrow T}\left[\sum_{\sigma\in S_t}\mathop{\mathbb{E}}_{k\leftarrow\{0,1\}^n}\left[\left(\left(Z^k\otimes I_m\right)^{\otimes\ell}\otimes I_{n+m}^{\otimes t}\right)|\mathbf{v}\rangle\langle\sigma(\mathbf{v})|\left(\left(Z^k\otimes I_m\right)^{\otimes\ell}\otimes I_{n+m}^{\otimes t}\right)\right]\right]. \tag{2}$$

Then from the previous lemma (Lemma 3)

$$\begin{aligned}(2) &= \mathop{\mathbb{E}}_{\mathbf{v}\leftarrow T}\left[\sum_{\sigma_1\in S_\ell,\sigma_2\in S_t}|\mathbf{v}\rangle\langle\sigma_1\circ\sigma_2(\mathbf{v})|\right]\\ &= \mathop{\mathbb{E}}_{\mathbf{v}\leftarrow T}\left[\sum_{\sigma_1\in S_\ell}|\mathbf{v}_{[1:\ell]}\rangle\langle\sigma_1(\mathbf{v}_{[1:\ell]})|\otimes\sum_{\sigma_2\in S_t}|\mathbf{v}_{[\ell+1:\ell+t]}\rangle\langle\sigma_2(\mathbf{v}_{[\ell+1:\ell+t]})|\right]\\ &= \mathbb{E}\left[\sum_{\sigma_1\in S_\ell}|\mathbf{v}_1\rangle\langle\sigma_1(\mathbf{v}_1)|\otimes\sum_{\sigma_2\in S_t}|\mathbf{v}_2\rangle\langle\sigma_2(\mathbf{v}_2)| : \begin{array}{c}X\leftarrow\binom{T}{\ell},\\ \mathbf{v}_1\leftarrow X,\\ \mathbf{v}_2\leftarrow T\backslash X\end{array}\right]\\ &= \mathop{\mathbb{E}}_{X\leftarrow\binom{T}{\ell}}\left[|X\rangle\langle X|\otimes|T\setminus X\rangle\langle T\setminus X|\right].\end{aligned}$$

For the first equality, we use Lemma 3 and decompose $\sigma=\sigma_1\circ\sigma_2$ for some σ_1,σ_2 such that $\sigma_1(x)=x$ for all $x\in\{\ell+1,\ell+2,\cdots,\ell+t\}$ and $\sigma_2(y)=y$ for all $y\in\{1,2,\cdots,\ell\}$. Since all $\ell+1,\ell+2,\cdots,\ell+t$ are fixed points of σ_1, we can view it as an element in S_ℓ. Similarly, we view $\sigma_2(y)$ as an element in S_t. The second equality follows by denoting the first ℓ part of $\mathbf{v}$ by $\mathbf{v}_{[1:\ell]}$ and the last t part of $\mathbf{v}$ by $\mathbf{v}_{[\ell+1:\ell+t]}$. The third equality holds because sampling a tuple $\mathbf{v}$ from T is equivalent to sampling an ℓ-subset X from T followed by ordering to elements in X and $T\setminus X$. □

4.2 Construction

In this section, we assume that the length of the common Haar state satisfies $n=n(\lambda)\geq\lambda$ for all $\lambda\in\mathbb{N}$. We define the construction as follows: on input $k\in\{0,1\}^\lambda$ and a single copy of the common Haar state $|\vartheta\rangle$,

$$G_k(|\vartheta\rangle) := (Z^k\otimes I_{n-\lambda})|\vartheta\rangle.$$

Lemma 5 (ℓ-Copy Pseudorandomness). *Let G be as defined above. Let*

$$\rho := \mathop{\mathbb{E}}_{\substack{k\leftarrow\{0,1\}^\lambda\\ |\vartheta\rangle\leftarrow\mathcal{H}_n}}\left[G_k(|\vartheta\rangle)^{\otimes\ell}\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t}\right] \text{ and } \sigma := \mathop{\mathbb{E}}_{\substack{|\varphi\rangle\leftarrow\mathcal{H}_n\\ |\vartheta\rangle\leftarrow\mathcal{H}_n}}\left[|\varphi\rangle\langle\varphi|^{\otimes\ell}\otimes|\vartheta\rangle\langle\vartheta|^{\otimes t}\right].$$

Then $\mathsf{TD}(\rho,\sigma)=O\left(\frac{(\ell+t)^{2\ell}}{2^\lambda}\right)$.

Proof. We prove this via a hybrid argument:

Hybrid 1. Sample $T \leftarrow [0 : \ell{+}t]^{2^n}$. Sample $k \leftarrow \{0,1\}^\lambda$. Output $((Z^k \otimes I_{n-\lambda})^{\otimes \ell} \otimes I_n^{\otimes t})|T\rangle$.

Hybrid 2. Sample $T \leftarrow [0 : \ell + t]^{2^n}$ uniformly conditioned on $T \in \mathcal{I}^{(\ell)}_{\lambda,n-\lambda}(\ell + t)$. Sample $k \leftarrow \{0,1\}^\lambda$. Output $((Z^k \otimes I_{n-\lambda})^{\otimes \ell} \otimes I_n^{\otimes t})|T\rangle$.

Hybrid 3: Sample $T \leftarrow [0 : \ell + t]^{2^n}$ uniformly conditioned on $T \in \mathcal{I}^{(\ell)}_{\lambda,n-\lambda}(\ell + t)$. Sample a uniform ℓ-subset T_1 from T. Output $|T_1\rangle \otimes |T \setminus T_1\rangle$.

Hybrid 4. Sample $T \leftarrow [0 : \ell + t]^{2^n}$. Sample a uniform ℓ-subset T_1 from T.[15] Output $|T_1\rangle \otimes |T \setminus T_1\rangle$.

Hybrid 5. Sample a collision-free T from $[0 : \ell + t]^{2^n}$. Sample a uniform ℓ-subset T_1 from T. Output $|T_1\rangle \otimes |T \setminus T_1\rangle$.

Hybrid 6. Sample a uniform collision-free T_1 from $[0 : \ell]^{2^n}$. Sample a uniform collision-free T_2 from $[0 : t]^{2^n}$ conditioned on T_1 and T_2 have no common elements. Output $|T_1\rangle \otimes |T_2\rangle$.

Hybrid 7. Sample a uniform collision-free T_1 from $[0 : \ell]^{2^n}$. Sample a uniform collision-free T_2 from $[0 : t]^{2^n}$. Output $|T_1\rangle \otimes |T_2\rangle$.

Hybrid 8. Sample $T_1 \leftarrow [0 : \ell]^{2^n}$. Sample $T_2 \leftarrow [0 : t]^{2^n}$. Output $|T_1\rangle \otimes |T_2\rangle$.

Indistinuishability of Hybrids.

- By Lemma 2, the trace distance between Hybrid 1 and Hybrid 2 is $O((t + \ell)^{2\ell}/2^\lambda)$.
- From Lemma 4, the output of Hybrid 2 is

$$\mathop{\mathbb{E}}_{\substack{T \leftarrow [0:\ell+t]^{2^n}: \\ T \in \mathcal{I}^{(\ell)}_{\lambda,n-\lambda}(\ell+t)}} \ \mathop{\mathbb{E}}_{T_1 \leftarrow \binom{T}{\ell}} \left[|T_1\rangle\langle T_1| \otimes |T \setminus T_1\rangle\langle T \setminus T_1| \right].$$

  Hence, Hybrid 2 is equivalent to Hybrid 3.
- Again by Lemma 2, the trace distance between Hybrid 3 and Hybrid 4 is $O((t + \ell)^{2\ell}/2^\lambda)$.
- The trace distance between Hybrid 4 and Hybrid 5 is $O((t + \ell)^2/2^n)$ by the collision bound.
- Hybrid 5 and Hybrid 6 are equivalent.
- The trace distance between Hybrid 6 and Hybrid 7 is $O(t\ell/2^n)$.

[15] Since T might have collisions, T_1 is allowed to contain duplicate elements.

- Finally, the trace distance between Hybrid 7 and Hybrid 8 is $O((t^2+\ell^2)/2^n)$ by the collision bound.

This completes the proof. □

In the following, we show that our construction also satisfies multi-key ℓ-copy pseudorandomness using Lemma 5.

Lemma 6 (Multi-key ℓ-Copy Pseudorandomness). *Let G be defined as above. Let*

$$\rho := \bigotimes_{i=1}^{p} \mathop{\mathbb{E}}_{|\varphi_i\rangle \leftarrow \mathcal{H}_n} \left[|\varphi_i\rangle\langle\varphi_i|^{\otimes \ell}\right] \otimes \mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathcal{H}_n} \left[|\vartheta\rangle\langle\vartheta|^{\otimes t}\right]$$

and

$$\sigma := \mathop{\mathbb{E}}_{|\vartheta\rangle \leftarrow \mathcal{H}_n} \left[\bigotimes_{i=1}^{p} \mathop{\mathbb{E}}_{k_i \leftarrow \{0,1\}^\lambda} \left[G_{k_i}(|\vartheta\rangle)^{\otimes \ell}\right] \otimes |\vartheta\rangle\langle\vartheta|^{\otimes t}\right].$$

Then $\mathsf{TD}(\rho, \sigma) = O\left(\frac{p\cdot(p\ell+t)^{2\ell}}{2^\lambda}\right)$.

The proof of Lemma 6 can be found in the full version [AGL24].

Proof of Theorem 7. Our construction is an efficiently-implementable unitary channel and thus satisfies the state generation property. Pseudorandomness follows from Lemma 6. □

4.3 Optimality of Our PRSG Construction

In this section, if the PRS generation algorithm uses only *one copy* of the common Haar state, we show that ℓ-copy statistical PRS and multi-key ℓ-copy statistical PRS are impossible for $\ell = \Omega(\lambda/\log(\lambda))$ and $n = \omega(\log(\lambda))$.

Theorem 8. *Statistically secure (λ, n, ℓ)-PRS is impossible in the CHS model if (a) the generation algorithm uses only one copy of the common Haar state, (b) $n = \omega(\log(\lambda))$, (c) $\ell = \Omega(\lambda/\log(\lambda))$ and, (d) the length of the common Haar state is $n = \omega(\log(\lambda))$.*

The proof of Theorem 8 can be found in the full version [AGL24].

5 Statistical Stretch PRFS Generators in the CHS Model

In this section, we extend our techniques from Sect. 4.2 to construct an (λ, m, n, ℓ)-statistical PRFS in the CHS model, where $m = \lambda^c$, $\ell = \lambda^{1-c}/\log(\lambda)^{1+\varepsilon}$, the length of the common Haar state is $n \geq \lambda^{1-c}$, for any constant $\varepsilon > 0$ and $c \in [0, 1)$. In the case when $n > \lambda$, the construction satisfies stretch property. We prove the following theorem in the full version [AGL24].

Theorem 9. *There exists an (λ, m, n, ℓ)-statistical selectively secure PRFS generator in the CHS model where the length of the common Haar state is $n(\lambda)$, $m(\lambda) = \lambda^c$, $\ell = O(\lambda^{1-c}/\log(\lambda)^{1+\varepsilon})$ and $n(\lambda) \geq \lambda^{1-c}$, for any constant $\varepsilon > 0$ and for any $c \in [0, 1)$.*

Note that since a PRS can be used to computationally instantiate CHS in the plain model, the above result also gives us a way to get bounded-query long-input PRFS from PRS in the plain model. In more detail, we can start with a PRS that has stretch (i.e. $n > \lambda$) and then we can bootstrap into a PRFS for large input length at the cost of a reduction in stretch.[16]

Corollary 2. Assuming the existence of (λ, n, ℓ)-PRS, for $n > \lambda$ and $\ell = O(\lambda^{1-c}/\log(\lambda)^{1+\varepsilon})$, there exists a selectively secure $(2\lambda, m, n, \ell)$-PRFS generator with $m(\lambda) = \lambda^c$, for any constant $\varepsilon > 0$ and for any $c \in [0, 1)$.

Furthermore, since PRFS imply PRS, achieving an ℓ-query statistical PRFS in the CHS model for $\ell = \Omega(\lambda/\log(\lambda))$ is impossible from Theorem 8.

Corollary 3. (λ, m, n, ℓ)-statistical PRFS is impossible in the CHS model if (a) the generation algorithm uses only one copy of the common Haar state, (b) $\ell = \Omega(\lambda/\log(\lambda))$, (c) the length of the common Haar state is n and, (d) $n = \omega(\log(\lambda))$.

5.1 Construction

We extend the techniques used in Sect. 4.2 to construct a statistical PRFS in Fig. 1. The construction samples a uniform key for each position of the input being zero or one. Applying this to the common Haar state gives us the output of the PRFS. The details can be seen in Fig. 1. Thoughout this section, one should think of $m = \lambda^c$ and $\lambda' = \lambda^{1-c}$ for some constant $c \in [0, 1)$.

Given the common Haar state $|\vartheta\rangle$, on the key $K = (k_1^0, \ldots, k_m^0, k_1^1, \ldots, k_m^1) \in \{0,1\}^{2\lambda' m}$ and the input $\mathbf{x} = (x_1, \ldots, x_m) \in \{0,1\}^m$, define $G(K, \mathbf{x}, |\vartheta\rangle)$ as follows:

- $|\psi_{K,\mathbf{x}}\rangle = G(K, \mathbf{x}, |\vartheta\rangle) = (Z^{\bigoplus_{i=1}^m k_i^{x_i}} \otimes I_{n-\lambda'})|\vartheta\rangle$.
- Output $|\psi_{K,\mathbf{x}}\rangle$.

Fig. 1. PRFS in the CHS model

The main property of the construction that makes it a PRFS is its ability to *disentangles* any type state in $\mathcal{I}^{(\ell)}_{\lambda', n-\lambda'}(\ell + t)$ into a probabilistic mixture of disjoint subsets of the type.

[16] Formally, let G_{PRS} is a (λ, n, ℓ)-PRS and $G(k, x, |\phi\rangle)$ is (λ, m, n, ℓ)-statistical selectively secure PRFS generator in the CHS model with $n > \lambda$, $\ell = O(\lambda^{1-c}/\log(\lambda)^{1+\varepsilon})$ and $m(\lambda) = \lambda^c$, then for $K = (k_1, k_2) \in \{0,1\}^\lambda \times \{0,1\}^\lambda$ we can define $G_{PRFS}(k, x) := G(k_1, x, G_{PRS}(k_2))$ as the $(2\lambda, m, n, \ell)$-PRFS generator.

6 Quantum Commitments in the CHS Model

In this section, we construct a commitment scheme that satisfies poly-copy statistical hiding and statistical sum-biding in the CHS model. The scheme is inspired by the quantum commitment scheme proposed in [MY21, MNY23]. In contrast to the scheme in [MY21], our construction is not of the canonical form [Yan2]. To achieve binding, similar to [MNY23], the receiver needs to perform several SWAP tests. To achieve hiding, our scheme relies on the multi-key pseudorandomness property in Lemma 6.

6.1 Construction

We assume that $n(\lambda) \geq \lambda + 1$ for all $\lambda \in \mathbb{N}$. Our construction, parameterized by the polynomial $p = p(\lambda) := \lambda$, is shown in Fig. 2. In the full version [AGL24], we prove the following theorem:

Theorem 10. *The construction in Fig. 2 is a quantum commitment in the CHS model.*

Commit phase: The sender C_λ on input $b \in \{0,1\}$ does the following:

- Use p copies of the common Haar state $|\vartheta\rangle$ to prepare the state $|\varPsi_b\rangle_{\mathsf{CR}} := \bigotimes_{i=1}^{p} |\psi_b\rangle_{\mathsf{C}_i\mathsf{R}_i}$, where

$$|\psi_0\rangle_{\mathsf{C}_i\mathsf{R}_i} := \frac{1}{\sqrt{2^\lambda}} \sum_{k\in\{0,1\}^\lambda} (Z^k \otimes I_{n-\lambda})|\vartheta\rangle_{\mathsf{C}_i} |k||0^{n-\lambda}\rangle_{\mathsf{R}_i}$$

  and

$$|\psi_1\rangle_{\mathsf{C}_i\mathsf{R}_i} := \frac{1}{\sqrt{2^n}} \sum_{j\in\{0,1\}^n} |j\rangle_{\mathsf{C}_i} |j\rangle_{\mathsf{R}_i},$$

  and $\mathsf{C} := (\mathsf{C}_1, \mathsf{C}_2, \ldots, \mathsf{C}_p)$ and $\mathsf{R} := (\mathsf{R}_1, \mathsf{R}_2, \ldots, \mathsf{R}_p)$.
- Send register C to the receiver.

Reveal phase:

- The sender sends b and register R to the receiver.
- The receiver prepares the state $|\varPsi_b\rangle_{\mathsf{C}'\mathsf{R}'} = \bigotimes_{i=1}^{p} |\psi_b\rangle_{\mathsf{C}'_i\mathsf{R}'_i}$ by using p copies of the common Haar state $|\vartheta\rangle$, where $\mathsf{C}' := (\mathsf{C}'_1, \mathsf{C}'_2, \ldots, \mathsf{C}'_p)$ and $\mathsf{R}' := (\mathsf{R}'_1, \mathsf{R}'_2, \ldots, \mathsf{R}'_p)$ are receiver's registers.
- For $i \in [p]$, the receiver performs the SWAP test between registers $(\mathsf{C}_i, \mathsf{R}_i)$ and $(\mathsf{C}'_i, \mathsf{R}'_i)$.
- The receiver outputs b if all SWAP tests accept; otherwise, outputs $\perp$.

Fig. 2. Quantum commitment scheme in the CHS model

7 LOCC Indistinguishability

In this section, we prove our main technical theorem for proving impossibilities and separations in Sect. 8 and Sect. 9.

7.1 Definitions

Definition 12 (LOCC Adversaries). *An* LOCC adversary *is a tuple (A, B), where A and B are spatially separated, non-uniform, and computationally unbounded quantum algorithms without pre-shared entanglement. In addition, A and B can only perform local operations on their registers and communicate classically.*

Definition 13 (LOCC Indistinguishability). *We say that two density matrices $(\rho_{\mathsf{AB}}, \sigma_{\mathsf{AB}})$ are* ε-LOCC indistinguishable *if for any LOCC adversary (A, B) with A taking as input register A and B taking as input register B, the probability that B outputs 1 satisfies*[17]

$$|\Pr[(A, B)(\rho_{\mathsf{AB}}) = 1] - \Pr[(A, B)(\sigma_{\mathsf{AB}}) = 1]| \leq \varepsilon.$$

If $\varepsilon(\cdot)$ is negligible, then we simply say that $(\rho_{\mathsf{AB}}, \sigma_{\mathsf{AB}})$ are LOCC indistinguishable.

7.2 LOCC Haar Indistinguishability

We prove the following theorem in the full version [AGL24]:

Theorem 11 (LOCC Haar Indistinguishability). *Let $\rho_{\mathsf{AB}} := \mathbb{E}_{|\psi\rangle\leftarrow\mathcal{H}_n} |\psi\rangle\langle\psi|_{\mathsf{A}}^{\otimes t} \otimes |\psi\rangle\langle\psi|_{\mathsf{B}}^{\otimes t}$ and $\sigma_{\mathsf{AB}} := \mathbb{E}_{|\psi\rangle\leftarrow\mathcal{H}_n} \left[|\psi\rangle\langle\psi|_{\mathsf{A}}^{\otimes t}\right] \otimes \mathop{\mathbb{E}}_{|\phi\rangle\leftarrow\mathcal{H}_n} \left[|\phi\rangle\langle\phi|_{\mathsf{B}}^{\otimes t}\right]$. Then ρ_{AB} and σ_{AB} are $O(t^2/2^n)$-LOCC indistinguishable.*

7.3 An Optimal LOCC Haar Distinguisher

We present an (optimal) LOCC Haar distinguisher with advantage $\Omega(t^2/2^n)$ in the full version [AGL24]. Hence, the upper bound in Theorem 11 is tight.

8 Impossibilities of QCCC Primitives in the CHS Model

In this section, we investigate the impossibility of *statistically* secure quantum-computation classical-communication (QCCC) primitives in the CHS model. We prove the following theorem in the full version [AGL24]:

Theorem 12. *There does not exist primitive $\mathcal{P}$ in the CHS model where $\mathcal{P} \in$ {statistically secure QCCC key agreements, statistically hiding and statistically binding QCCC interactive commitments}.*

[17] Since (A, B) are allowed to communicate and we do not care about communication complexity, it is without loss of generality to assume that B outputs the bit.

9 Quantum Black-Box Separation in the QCCC Model

9.1 The Separating Oracle

As is common in black-box impossibility results, we will define oracles relative to which $\omega(\log(\lambda))$-PRSGs exist while QCCC key agreements and interactive commitments do not. We define the oracle $G := \{\{G_k\}_{k\in\{0,1\}^\lambda}\}_{\lambda\in\mathbb{N}}$ as follows. For every $\lambda \in \mathbb{N}$ and $k \in \{0,1\}^\lambda$, the oracle G_k is a Haar isometry that maps any state $|\psi\rangle$ to $|\psi\rangle|\vartheta_k\rangle$, where $|\vartheta_k\rangle$ is a Haar state of length $n(\lambda) = \omega(\log(\lambda))$. The existence of $\omega(\log(\lambda))$-PRSGs relative to G can be proven easily.

9.2 Separating QCCC Key Agreements from $(\boldsymbol{\lambda}, \boldsymbol{\omega}(\log(\boldsymbol{\lambda})))$-PRSGs

In the full version [AGL24], we prove the following theorem:

Theorem 13. *There does not exist a quantum fully black-box reduction* (C, S) *from QCCC key agreements to* $(\lambda, \omega(\log(\lambda)))$*-PRSGs such that* C *only asks classical queries to the PRSG.*

9.3 Separating QCCC Interactive Commitments from $(\boldsymbol{\lambda}, \boldsymbol{\omega}(\log(\boldsymbol{\lambda})))$-PRSGs

In the full version [AGL24], we prove the following theorem:

Theorem 14. *There does not exist a quantum fully black-box reduction* (C, S) *from QCCC Interactive Commitments to* $(\lambda, \omega(\log(\lambda)))$*-PRSGs such that* C *only asks classical queries to the PRSG.*

Acknowledgements. This work is supported by the National Science Foundation under Grant No. 2329938 and Grant No. 2341004.

References

[ACC+22] Austrin, P., Chung, H., Chung, K.-M., Fu, S., Lin, Y.-T., Mahmoody, M.: On the impossibility of key agreements from quantum random oracles. In: Annual International Cryptology Conference, pp. 165–194. Springer, Cham (2022)

[ACH+23] Afshar, A., Chung, K.-M., Hsieh, Y.-C., Lin, Y.-T., Mahmoody, M.: On the (im) possibility of time- lock puzzles in the quantum random oracle model. In: International Conference on the Theory and Application of Cryptology and Information Security, pp. 339–368. Springer, Singapore (2023). (cit. on p. 6)

[AGKL23] Ananth, P., Gulati, A., Kaleoglu, F., Lin, Y.- T.: Pseudorandom Isometries. arXiv preprint arXiv:2311.02901 [quant-ph] (2023)

[AGL24] Ananth, P., Gulati, A., Lin, Y.-T.: Cryptography in the common Haar state model: feasibility results and separations. Cryptology ePrint Archive, Paper 2024/1043 (2024). https://eprint.iacr.org/2024/1043

[AGQY22] Ananth, P., Gulati, A., Qian, L., Yuen, H.: Pseudorandom (function-like) quantum state generators: new definitions and applications. In: Theory of Cryptography Conference, pp. 237–265. Springer, Cham (2022)

[AHY23] Ananth, P., Hu, Z., Yuen, H.: On the (im) plausibility of public-key quantum money from collision-resistant hash functions. In: International Conference on the Theory and Application of Cryptology and Information Security, pp. 39–72. Springer, Singapore (2023)

[AKY24] Ananth, P., Kaleoglu, F., Yuen, H.: Simultaneous Haar Indistinguishability with applications to unclonable cryptography. In: arXiv preprint arXiv:2405.10274 (2024)

[AQY22] Ananth, P., Qian, L., Yuen, H.: Cryptography from pseudorandom quantum states. In: CRYPTO (2022)

[BBF13] Baecher, P., Brzuska, C., Fischlin, M.: Notions of black-box reductions, revisited. In: Advances in Cryptology- ASIACRYPT 2013: 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, 1–5 December 2013, Proceedings, Part I, 19, pp. 296–315. Springer, Heidelberg (2013)

[BCKM21] Bartusek, J., Coladangelo, A., Khurana, D., Ma, F.: One-way functions imply secure computation in a quantum world. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12825, pp. 467–496. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84242-0_17

[BCQ23] Brakerski, Z., Canetti, R., Qian, L.: On the computational hardness needed for quantum cryptography. In: 14th Innovations in Theoretical Computer Science Conference, ITCS 2023, p. 24. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing (2023)

[BDF+99] Bennett, C.H., et al.: Quantum nonlocality without entanglement. Phys. Rev. A **59**(2), 1070 (1999)

[BFM19] Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications. In: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, pp. 329–349 (2019)

[BGVV+23] Bouaziz, S., Grilo, A.B., Vergnaud, D., Vu, Q.-H., et al.: Towards the impossibility of quantum public key encryption with classical keys from one-way functions. In: Cryptology ePrint Archive (2023)

[BL18] Benhamouda, F., Lin, H.: k-round multiparty computation from k-round oblivious transfer via garbled interactive circuits. In: Advances in Cryptology–EUROCRYPT 2018: 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, 29 April–3 May 2018, Part II 37, pp. 500–532. Springer, Cham (2018)

[BM+24] Bouaziz, S., Muguruza, G., et al.: Quantum Pseudorandomness Cannot be Shrunk in a Black-Box Way. In: Cryptology ePrint Archive (2024)

[Bra23] Brakerski, Z.: Black-hole radiation decoding is quantum cryptography. In: Annual International Cryptology Conference, pp. 37–65. Springer (2023)

[CCHL22] Chen, S., Cotler, J., Huang, H.-Y., Li, J.: Exponential separations between learning with and without quantum memory. In: 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 574–585. IEEE (2022)

[CCS24] Chen, B., Coladangelo, A., Sattath, O.: The power of a single Haar random state: constructing and separating quantum pseudorandomness. In: arXiv preprint arXiv:2404.03295 (2024)

[CF01] Canetti, R., Fischlin, M.: Universally composable commitments. In: Advances in Cryptology—CRYPTO 2001: 21st Annual International Cryptology Conference, Santa Barbara, 19–23 August 2001 Proceedings 21, pp. 19–40. Springer, Heidelberg (2001)

[CGG24] Chung, K.-M., Goldin, E., Gray, M.: On central primitives for quantum cryptography with classical communication. arXiv preprint arXiv: 2402.17715 [cs.CR] (2024)

[CH14] Chitambar, E., Hsieh, M.-H.: Asymptotic state discrimination and a strict hierarchy in distinguishability norms. J. Math. Phys. **55**(11) (2014)

[CLM+14] Chitambar, E., Leung, D., Mančinska, L., Ozols, M., Winter, A.: Everything you always wanted to know about LOCC (but were afraid to ask). Commun. Math. Phys. **328**, 303–326 (2014)

[CLM23] Chung, K.-M., Lin, Y.-T., Mahmoody, M.: Black-box separations for non-interactive classical commitments in a quantum world. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques, pp. 144–172. Springer, Cham (2023)

[CLMO13] Childs, A.M., Leung, D., Mančinska, L., Ozols, M.: A framework for bounding nonlocality of state discrimination. Commun. Math. Phys. 323, 1121–1153 (2013)

[CLOS02] Canetti, R., Lindell, Y., Ostrovsky, R., Sahai, A.: Universally composable two-party and multi-party secure computation. In: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, pp. 494–503 (2002)

[CM24] Coladangelo, A., Mutreja, S.: On black-box separations of quantum digital signatures from pseudorandom states. arXiv preprint arXiv:2402.08194 (2024)

[Col23] Coladangelo, A.: Quantum trapdoor functions from classical one-way functions. Cryptology ePrint Archive, Paper 2023/282 (2023). https://eprint.iacr.org/2023/282

[DLT02] DiVincenzo, D.P., Leung, D.W., Terhal, B.M.: Quantum data hiding. IEEE Trans. Inf. Theory **48**(3), 580–598 (2002)

[EW02] Eggeling, T., Werner, R.F.: Hiding classical data in multipartite quantum states. Phys. Rev. Lett. **89**(9), 097905 (2002)

[Gea02] Gea-Banacloche, J.: Hiding messages in quantum data. In: J. Math. Phys. **43**(9), 4531–4536 (2002)

[GGM86] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. J. ACM **33**(4), 792–807 (1986)

[GLSV21] Grilo, A.B., Lin, H., Song, F., Vaikuntanathan, V.: Oblivious transfer is in MiniQCrypt. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12697, pp. 531–561. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77886-6_18

[GS22] Garg, S., Srinivasan, A.: Two-round multiparty secure computation from minimal assumptions. J. ACM **69**(5), 1–30 (2022)

[Har13] Harrow, A.W.: The church of the symmetric subspace. arXiv preprint arXiv:1308.6595 (2013)

[Har23] Harrow, A.W.: Approximate orthogonality of permutation operators, with application to quantum information. Lett. Math. Phys. **114**(1), 1 (2023)

[HBAB19] Halder, S., Banik, M., Agrawal, S., Bandyopadhyay, S.: Strong quantum nonlocality without entanglement. Phys. Rev. Lett. **122**(4), 040403 (2019)

[HLS05] Hayden, P., Leung, D., Smith, G.: Multiparty data hiding of quantum information. Phys. Rev. A **71**(6), 062339 (2005)

[HY20] Hosoyamada, A., Yamakawa, T.: Finding collisions in a quantum world: quantum black-box separation of collision-resistance and one-wayness. In: Advances in Cryptology– ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, 7–11 December 2020, Proceedings, Part I, 26, pp. 3–32. Springer, Cham (2020)

[JLS18] Ji, Z., Liu, Y.-K., Song, F.: Pseudorandom quantum states. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10993, pp. 126–152. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96878-0_5

[Kre21] Kretschmer, W.: Quantum pseudorandomness and classical complexity. In: Hsieh, M.-H. (ed.) 16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, 5–8 July 2021, Virtual Conference. LIPIcs, vol. 197, pp. 2:1–2:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2021). https://doi.org/10.4230/LIPIcs.TQC.2021.2

[LC97] Lo, H.-K., Chau, H.F.: Is quantum bit commitment really possible? Phys. Rev. Lett. **78**(17), 3410 (1997)

[LLLL24] Li, L., Li, Q., Li, X., Liu, Q.: How (not) to build quantum PKE in Minicrypt. arXiv preprint arXiv:2405.20295 (2024)

[May97] Mayers, D.: Unconditionally secure quantum bit commitment is impossible. Phys. Rev. Lett. **78**(17), 3414 (1997)

[MNY23] Morimae, T., Nehoran, B., Yamakawa, T.: Unconditionally secure commitments with quantum auxiliary inputs. arXiv preprint arXiv:2311.18566 [quant-ph] (2023)

[MWW09] Matthews, W., Wehner, S., Winter, A.: Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding. Commun. Math. Phys. **291**, 813–843 (2009)

[MY21] Morimae, T., Yamakawa, T.: Quantum commitments and signatures without one-way functions. arXiv preprint arXiv:2112.06369 (2021)

[MY23] Morimae, T., Yamakawa, T.: One-wayness in quantum cryptography. arXiv preprint arXiv:2210.03394 [quant-ph] (2023)

[NC10] Nielsen, M.A., Chuang, I.L.: Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press (2010). https://doi.org/10.1017/CBO9780511976667

[PNC14] Piani, M., Narasimhachar, V., Calsamiglia, J.: Quantumness of correlations, quantumness of ensembles and quantum data hiding. New J. Phys. **16**(11), 113001 (2014)

[Qia23] Qian, L.: Unconditionally secure quantum commitments with preprocessing. In: Cryptology ePrint Archive (2023)

[RTV04] Reingold, O., Trevisan, L., Vadhan, S.: Notions of reducibility between cryptographic primitives. In: Theory of Cryptography Conference, pp. 1–20. Springer, Heidelberg (2004)

[Yan2] Yan, J.: General properties of quantum bit commitments. In: International Conference on the Theory and Application of Cryptology and Information Security, pp. 628–657. Springer, Cham (2022)