Frontiers of Quantum Complexity and Cryptography	Spring 2022

Lecture 9 - Constructing Pseudorandom States

Lecturer: Henry Yuen — *Scribes: Yuval Efron/Melody Hsu*

1 Constructing PRS from PRG

In order to show how a PRS can be constrcuted from a PRG, we introduce the notion of pseudo-random functions(PRF).

Definition 1. *A function $F : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}$ is called a pseudo-random function if F is computable in polynomial time, and for every polynomial time(perhaps quantum) distinguisher D it holds that*

$$|| \Pr_{k \sim_R \{0,1\}^n}(D^{F_k} = 1) - \Pr_{f \sim_R \{0,1\}^{\{0,1\}^m}}(D^f = 1)|| \leq neg$$

In the above, for every $k \in \{0,1\}^n$, $F_k : \{0,1\}^m \to \{0,1\}$ is defined by $F_k(x) = F(x,k)$, and D^g for a function g denotes granting D oracle access to g.

Fact 2. *A well known result in cryptography states that PRG and PRF are equivalent primitives. In particular, if a PRG exists, so does a PRF.*

The rest of the section focuses on proving the following theorem due to Ji, Liu, Song [JLS18].

Theorem 3. *PRFs imply PRS.*

We begin by describing the construction.

Construction. Let $F : \{0,1\}^m \times \{0,1\}^n \to \{0,1\}$. We construct $G : \{0,1\}^n \to (\mathbb{C}^2)^{\otimes m}$ as follows.

1. By applying an H gate to each qubit, we prepare the uniform superposition state: $2^{-m/2} \sum\limits_{x \in \{0,1\}^m} |x\rangle$.
2. Given $k \{0,1\}^n$, we compute F_k in superposition on the above state, to obtain $|\psi_k\rangle = 2^{-m/2} \sum\limits_{x \in \{0,1\}^m} (-1)^{F_k(x)} |x\rangle$.
3. Output $|\psi_k\rangle$.

We first note that each of the above steps can be executed in quantum polynomial time, as F_k is computable in polynomial time. All that is left is proof of security. Namely we aim to prove the following.

Claim 4. *The ensemble $\{|\psi_k\rangle\}_k$ is indistinguishable from a Haar random state on m qubits, even given $poly(n)$ copies.*

Proof. Let D be a distinguisher and fix $t = poly(n)$. The proof employs a hybrid argument. Specifically, we examine D's behaviour on 3 different distributions: The first is $\{|\psi_k\rangle\}_k$, the second would be an interpolation of $\{|\psi_k\rangle\}_k$ and Haar, and the third would be a random Haar state. Formally, we consider the following experiments.

Experiment 1.

1. Sample a uniformly random $k \in \{0,1\}^n$.
2. Create t copies of $|\psi_k\rangle$.
3. Compute $D(|\psi_k\rangle^{\otimes t})$ and output the result.

Experiment 2.

1. Sample a random function $f : \{0,1\}^m \to \{0,1\}$. It is helpful to think of this step as being executed by a third party, and not the distinguisher.
2. Generate t copies of $|\psi_f\rangle = 2^{-m/2} \sum_{x\in\{0,1\}^m} (-1)^{f(x)}|x\rangle$. Denote the vector of coefficients by α, with $\alpha_x = (-1)^{f(x)}$.
3. Compute $D(|\psi_f\rangle^{\otimes t})$ and output the result.

Experiment 3.

1. Sample a random Haar state $|\theta\rangle$.
2. Compute t copies of $|\theta\rangle$
3. Compute $D(|\theta\rangle^{\otimes t})$ and output the answer.

Notice that our overarching goal is to show that the distributions produced by experiments $1(Exp_1)$ and $3(Exp_3)$ are close, we do this by showing that experiment 1 is close to experiment $2(Exp_2)$, and experiment 2 is close to experiment 3.

Observation 5. *$||Exp_1 - Exp2||_1 \leq neg$.*

The above holds from the assumption that F is a PRF.

All that is left to is to show that $||Exp_2 - Exp_3||_1 \leq neg$. We actually show that these distributions are close regardless of the chosen distinguisher, i.e. we bound the trace distance between the state distributions.

Specifically, we show that

$$||\mathop{\mathbb{E}}_f |\psi_f\rangle\langle\psi_f|^{\otimes t} - \mathbb{E}\,|\theta\rangle\langle\theta|^{\otimes t}||_1 \leq O(\frac{t^2}{2^m})$$

First, we recall that $\mathbb{E}\,|\theta\rangle\langle\theta|^{\otimes t} = \frac{\Pi_{sym}^{M,t}}{\mathrm{Tr}(\Pi)}$ where $M = 2^m$ and the nominator is the projector onto the symmetric space as we saw in previous lectures. When the parameters are clear from context, we refer to this operator as Π.

$$|\psi_f\rangle^{\otimes t} = 2^{-mt/2} \sum_{x_1,...,x_t} \alpha_{x_1}...\alpha_{x_t}|x_1,...,x_t\rangle$$

Define:

$$|\sigma\rangle = 2^{-mt/2} \sum_{x_1,...,x_t,\text{all distinct}} \alpha_{x_1}...\alpha_{x_t}|x_1,...,x_t\rangle$$

Notice that both of the above states have mt qubits. In $|\sigma\rangle$, we sum over all t-tuples of strings of length m that are pairwise distinct, denote this set by $S_{m,t}$.

With $|\sigma\rangle$ in mind, we notice the following and leave the proof as an exercise. Note that this claim essentially boils down to the question: Given t uniform random strings of length $m > n$, what is the probability that two of them are the same? If $t = poly(n)$, the probability is vanishingly small.

Claim 6. $|||\psi_f\rangle^{\otimes t} - |\sigma\rangle||_1 \leq O(\frac{t^2}{2^m})$

Thus, all that is left is to show that $||\,\mathbb{E}_f\,|\sigma\rangle\langle\sigma| - \frac{\Pi}{\mathrm{Tr}(\Pi)}||_1 \leq O(\frac{t^2}{2^m})$

We start with

$$\mathop{\mathbb{E}}_f |\sigma\rangle\langle\sigma| = 2^{-mt} \sum_{x,y\in S_{m,t}} \mathbb{E}[\alpha_{x_1}...\alpha_{x_t}\alpha_{y_1}...\alpha_{y_t}]|x_1,...,x_t\rangle\langle y_1,...,y_t|$$

Fix x, y, and note that the value of $\mathbb{E}[\alpha_{x_1}...\alpha_{x_t}\alpha_{y_1}...\alpha_{y_t}]$ can be deduced very easily. If $(x_1,...,x_t)$ is a permutation of $(y_1,...,y_t)$ then the expression equals 1, and otherwise at least one x_i is different from all strings in x, y and thus α_{x_i} will be 1 half the time and -1 half the time, which averages to 0. Thus we can continue:

$$2^{-mt} \sum_{x,y\in S_{m,t}} \mathbb{E}[\alpha_{x_1}...\alpha_{x_t}\alpha_{y_1}...\alpha_{y_t}]|x_1,...,x_t\rangle\langle y_1,...,y_t| =$$

$$2^{-mt} \sum_{x\in S_{m,t},\pi\in Sym_t} |x_1,...,x_t\rangle\langle x_{\pi(1)},...,x_{\pi(t)}| = 2^{-mt}A(\sum_{\pi\in Sym_t} P_\pi) = 2^{-mt}t!\Pi A\Pi$$

In the above Sym_t is the symmetric group on t elements, $A = \sum\limits_{x\in S_{m,t}} |x_1,...,x_t\rangle\langle x_1,...,x_t|$, and P_π is the permutation matrix defined by π. The last transition is due to $A\Pi = \Pi A$.

Now all that is left is to show that $||2^{-mt}t!\Pi A\Pi - \frac{\Pi}{\mathrm{Tr}(\Pi)}||_1 \leq O(\frac{t^2}{2^m})$. We note that the LHS is at most $2^{-mt}t!||\Pi A\Pi - \Pi||_1 + |2^{-mt} - 1/\binom{2^m+t-1}{t}| \leq O(\frac{t^2}{2^m})$. The last transition is left as an exercise to the reader, and might appear in the homework assignment. This concludes the proof that $||Exp_2 - Exp_3|| \leq neg$ and thus $||Exp_1 - Exp_3|| \leq neg$ which proves that G is a PRS, as required. $\square$

Acknowledgments. These lecture notes are based on handwritten lecture notes by Melody Hsu.

References

[JLS18] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In *CRYPTO (3)*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018.