Frontiers of Quantum Complexity and Cryptography Spring 2022

Lecture 9 - Quantum Pseudorandom States

Lecturer: Henry Yuen *Scribes: Dean Hirsch*

1 Overview

In this lecture we are going to discuss the quantum analog of a pseudo-random generator (PRG).

2 Review - Classical PRG

A PRG is a function G that takes as input a uniformly random key $k \in \{0,1\}^n$ and outputs a string $G(k) \in \{0,1\}^m$, with $m > n$, such that $G(k)$ looks indistinguishable from a random m-bit string from a polynomial verifier's point of view.

$$\overset{k}{\in\{0,1\}^n} \longrightarrow \boxed{G} \longrightarrow \overset{y}{\in\{0,1\}^m}$$

Formally:

Definition 1 (Classical PRG)**.** *$G : \{0,1\}^n \to \{0,1\}^m$ is a PRG if it is polynomial-time computable, and for all polynomial time distinguishers D the following holds:*

$$\left| \Pr_{k\sim\mathcal{U}_n}[D(G(k)) = 1] - \Pr_{y\sim\mathcal{U}_m}[D(y) = 1] \right| \leq \text{negligible}(n)$$

where $\mathcal{U}_n$ is the uniform distribution over $\{0,1\}^n$, and negligible(n) *is any function that goes to zero faster than any polynomial. That is, $f \in$ negligible(n) if $\lim_{n\to\infty} f(n) \cdot n^c = 0$ for all constants c.*

The difference $m - n$ is called the "stretch".

This definition captures the intuition that it should essentially be impossible to distinguish between $G(k)$ and a random y.

Notes and Observations

- Time restriction of D is necessary. Otherwise, there is not such thing as a PRG, since D can run on all possible inputs k in exponential time.
- We believe PRG exist, but we actually do not know for sure. If we did prove their existence, $P \neq NP$ would follow. In fact, virtually all of classical cryptography would imply the existence of a PRG.
- The stretch $(m - n)$ can be any poly(n), and in fact we can show that a classical PRG of stretch 1 can be amplified to a PRG of stretch poly(n) for any poly(n).

- In post-quantum cryptography (PQC), we assume the distinguisher D can be quantum. So a PRG in PQC is even stronger.

3 Pseudo-Random States (PRS)

We now discuss a quantum analog of a PRG.

Here we demand that several (t) copies of $|\psi_k\rangle$ are indistinguishable from t copies of a Haar random states, from the point of view of a poly-time distinguisher.

$$\underset{\in\{0,1\}^n}{k} \longrightarrow \boxed{G} \longrightarrow \underset{\in(\mathbb{C}^2)^{\otimes m}}{|\psi_k\rangle}$$

Definition 2 (PRS)**.** *A function*

$$G : \{0,1\}^n \to \mathcal{S}\left((\mathbb{C}^2)^{\otimes m}\right)$$

is a PRS generator if it can be computed in polynomial time, and for all $t = \text{poly}(n)$, for all poly-time distinguishers, D,

$$\left| \Pr_{k\sim\mathcal{U}_n}\left[D(|\psi_k\rangle^{\otimes t}) = 1\right] - \Pr_{|\theta\rangle\sim\text{Haar}(m)}\left[D(|\theta\rangle^{\otimes t}) = 1\right] \right| \leq \text{negligible}(n)$$

Question Why do we care about t? t is important because of the no-cloning theorem. In the quantum setting, it is possible to have two ensembles $|\psi_1\rangle \sim \mathcal{E}_1, |\psi_2\rangle \sim \mathcal{E}_2$ such that $|\psi_1\rangle$ is indistinguishable from $|\psi_2\rangle$ but $|\psi_1\rangle^{\otimes 2}$ *is* distinguishable from $|\psi_2\rangle^{\otimes 2}$

Example 3.

$$\mathcal{E}_1 = \{|0\rangle, |1\rangle\}, \mathcal{E}_2 = \{|+\rangle, |-\rangle\}$$

The density of one copy is $I/2$ in both cases - they maximally mixed states. However, for two copies, we can compute the densities matrices $\frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11|$ and $\frac{1}{2}|++\rangle\langle ++| + \frac{1}{2}|--\rangle\langle --|$ and observe that they are different.

As before, we observe that the polynomial-time restriction on D is necessary. This is because, with unlimited resources (time and copies), we could run state tomography on both inputs. However, it turns out that if we have t copies, where

$$2^n < \dim\left(\Pi^{2^m,t}_{\text{Sym}}\right)$$

then there still exists an (inefficient) distinguisher.

Observation 4. *Existence of a PRS implies $BQP \neq PSPACE$*

Most people believe pseudorandom states exist. For example, because they follow from the common belief that a classical PRG exists (we will show this in the second half of this class). The reverse direction, however, is not known to hold.

We also don't know how to stretch (or even shrink!) a PRS. This is because chopping off qubits from a state does not result in a pure state, as we will now see. This is a stark contrast with the classical setting.

We make the following informal claim, and show that we can formalize it to argue both that it is true and that it implies we cannot discard bits to shrink the PRS.

Claim 5. *The output of a PRS G is highly entangled.*

Formally, we measure entanglement of a m-qubit state $|\psi\rangle$, by measuring its *purity* across every partition of the bits. In more precise terms, we can trace out subsets of the qubits, and obtain a highly mixed state. Suppose we partition the qubits into two parts, A and B. Then, taking partial trace, $\rho_A = \mathrm{Tr}_B(|\psi\rangle\langle\psi|_{AB})$.

Definition 6 (purity)**.**

$$\mathrm{purity}(\rho_A) = \mathrm{Tr}(\rho_A^2)$$

If $\rho_A = |\theta\rangle\langle\theta|$ (the density matrix of a pure state), then $\mathrm{purity}(\rho_A) = 1$. In contrast, with a maximally mixed state $\rho_A = \frac{1}{\dim(A)} I$, then we see $\mathrm{purity}(\rho_A) = \frac{1}{\dim(A)} \ll 1$.

So intuitively, $\mathrm{purity}(\rho_A)$ being small implies $|\psi\rangle$ is highly entangled across the partition $A : B$.

We now make the previous claim more precise.

Claim 7. *For a randomly chosen k, and for* any *fixed cut $A : B$ of the qubits, purity*$(\rho_A^{(k)} \leq$ negligible

Proof. Suppose not. Then we show there exists an efficient poly-time distinguisher D between $|\psi_k\rangle$ and a Haar-random state. This will be our contradiction, by definition of a PRS.

The distinguisher is as follows. It takes two copies $|\psi\rangle_{AB}$, and discards the B part of both (essentially drawing from the distribution of possible remaining A-states, independently). It then uses a controlled-SWAP gate (a gate that swaps the inputs if the qubit is 1, while doing nothing if the qubit is 0) on the remaining A-parts, controlled by a $|+\rangle$ qubit. This qubit gets entangled with the two states, containing information about how different they are. The distinguisher then proceeds to pass the control qubit through a Hadamard gate, then measures it in the $\{|0\rangle, |1\rangle\}$ basis.

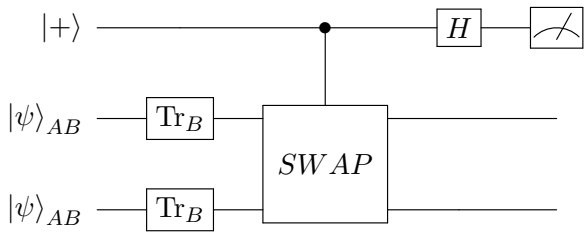

where the Tr_B "gate" means just discarding the B bits.

We claim (without proof) that

$$\Pr\left[D(|\psi\rangle^{\otimes 2}) = 1\right] = \frac{1}{2} + \frac{1}{2}\,\mathrm{purity}(\rho_A^2)$$

In contrast, inputting a true Haar-random state $|\theta\rangle$ instead, will give a probability of exactly $\frac{1}{2}$.

Thus, by assumption on ψ_{AB} being a PSR, it follows that purity(ρ_A^2) is negligible, and hence also purity(ρ_A).

□