6.S895: Quantum Cryptography

Lecture: Pseudorandom Quantum States

Lecturers: Alex Poremba **Scribe:** Vinod Vaikuntanathan

There are several reasons that notions of random and pseudo-random quantum states are interesting and useful. It is not hard to show that a random bipartite quantum state is highly entangled; a study of random and pseudorandom states is useful in understanding the notions of entanglement and more recently, *pseudoentanglement*. In cryptography, the notion of pseudorandom quantum states has emerged as a powerful lens through which to study the question of *what is the minimal assumption in quantum cryptography*. In classical cryptography, one-way functions are minimal. Pseudorandom quantum states are interesting because (a) the existence of such states implies the existence of much of quantum cryptography, far beyond key exchange; and (b) pseudorandom states could exist even when much of classical cryptography collapses (e.g. one-way functions do not exist) and even when many complexity classes collapse (e.g. $\mathcal{P} \neq \mathcal{NP}$ and even $\mathcal{BQP} = \mathcal{QMA}$.

Notation. For $d \in \mathbb{N}$, let $S(d)$ denote the set of unit vectors in $\mathbb{C}^d$, that is all $|\psi\rangle \in \mathbb{C}^d$ such that $\langle\psi|\psi\rangle = 1$. Let $U(d)$ denote the set of all d-by-d unitary matrices (acting on $\mathbb{C}^d$).

1 Haar-Random Quantum States

How should one define a "random quantum state"? Let's start by comparing classical randomness with quantum randomness, starting from a simple example. Consider the case of a random bit versus a random qubit. In the classical case, the sample space is $\Omega = \{0, 1\}$, and a random qubit gives us the state

$$\mathbb{E}_{b \leftarrow \{0,1\}}[|b\rangle\langle b|] = \frac{1}{2} \cdot |0\rangle\langle 0| + \frac{1}{2} \cdot |1\rangle\langle 1| = \frac{I}{2}$$

whereas in the quantum state, it is a random unit vector, so $\Omega = S(2)$. It turns out that

$$\mathbb{E}_{|\psi\rangle \leftarrow S(d)}[|\psi\rangle\langle\psi|] = I/2$$

So, what's the difference? Imagine you have an ensemble of states $\{(p_i, v_i)\}$, you sample a state v_i at random according to the probability distribution p_i and give t copies of this state, that is, $|v_i\rangle^{\otimes t}$. In the classical case, this is $|0\rangle^{\otimes t}$ with probability $1/2$, and $|1\rangle^{\otimes t}$ with probability $1/2$. In the quantum case, this is $|\psi\rangle^{\otimes t}$ for a "random quantum state". While we saw above that these two states are identical (have identical density matrices) when $t = 1$, these are *distinguishable* the moment t is larger than 1. This is not hard to see: in the classical case (when the state is $\frac{1}{2} \cdot (|0\rangle\langle 0| + |1\rangle\langle 1|)$ with $t = 2$, "measuring" the two states in the Z basis gives identical results, whereas in the quantum case (when the state is $|\psi\rangle^{\otimes 2}$), the measurement will produce independent random bits.

Definition 1 (Haar Measure). *The Haar measure μ_H is the unique left/right invariant measure over the unitary group $U(d)$. That is, for every "nice" function $f : \mathbb{C}^{d\times d} \to \mathbb{C}$ and every $V \in U(d)$,*

$$\mathbb{E}_{U\leftarrow U(d)}[f(U)] = \mathbb{E}_{U\leftarrow U(d)}[f(U \cdot V)] = \mathbb{E}_{U\leftarrow U(d)}[f(V \cdot U)]$$

where

$$\mathbb{E}_{U \leftarrow U(d)}[f(U)] = \int_{U(d)} f(U)\, d_{\mu_H} U$$

The definition of μ_H also extends to states in the following way:

$$\mathbb{E}_{|\psi\rangle \leftarrow S(d)}[f(|\psi\rangle\langle\psi|)] \stackrel{\text{def}}{=} \mathbb{E}_{U \leftarrow U(d)}[f(U|0\rangle\langle 0|U^\dagger)]$$

2 The Symmetric Subspace

Let $P_d(\sigma) \in \{0,1\}^{d^t \times d^t}$ be the matrix representation of the permutation $\sigma \in S_t$ acting on $[d]^t$. That is,

$$P_d(\sigma) = \sum_{i_1,\dots,i_t \in [d]} |i_{\sigma(1)},\dots,i_{\sigma(t)}\rangle\langle i_1,\dots,i_t| \ .$$

Indeed, note that

$$P_d(\sigma)|i_1,\dots,i_t\rangle = |i_{\sigma(1)},\dots,i_{\sigma(t)}\rangle$$

Let

$$\mathsf{Sym}_t(\mathbb{C}^d) = \{|\psi\rangle \in (\mathbb{C}^d)^{\otimes t} \ : \ P_d(\sigma)|\psi\rangle = |\psi\rangle \ \text{ for all } \sigma \in S_t\} \tag{1}$$

be a subspace of $(\mathbb{C}^d)^{\otimes t}$ which is invariant under the action of $P_d(\sigma)$ for any $\sigma \in S_t$.

We show a few important facts about the symmetric subspace and its relation to the Haar measure.

Lemma 2. *For any $\rho, \sigma \in S_t$,*

$$P_d(\rho)P_d(\sigma) = P_d(\rho\sigma)$$

Proof.

$$\begin{aligned}
P_d(\rho)P_d(\sigma) &= \left(\sum_{i_1,\dots,i_t \in [d]} |i_{\rho(1)},\dots,i_{\rho(t)}\rangle\langle i_1,\dots,i_t|\right)\left(\sum_{i_1,\dots,i_t \in [d]} |i_{\sigma(1)},\dots,i_{\sigma(t)}\rangle\langle i_1,\dots,i_t|\right) \\
&= \left(\sum_{i_1,\dots,i_t \in [d]} |i_{\rho(\sigma(1))},\dots,i_{\rho(\sigma(t))}\rangle\langle i_{\sigma(1)},\dots,i_{\sigma(t)}|\right)\left(\sum_{i_1,\dots,i_t \in [d]} |i_{\sigma(1)},\dots,i_{\sigma(t)}\rangle\langle i_1,\dots,i_t|\right) \\
&= \sum_{i_1,\dots,i_t \in [d]} |i_{\rho(\sigma(1))},\dots,i_{\rho(\sigma(t))}\rangle\langle i_{\sigma(1)},\dots,i_{\sigma(t)}|i_{\sigma(1)},\dots,i_{\sigma(t)}\rangle\langle i_1,\dots,i_t| \\
&= \sum_{i_1,\dots,i_t \in [d]} |i_{\rho(\sigma(1))},\dots,i_{\rho(\sigma(t))}\rangle\langle i_1,\dots,i_t| \\
&= P_d(\rho\sigma)
\end{aligned}$$

where the third equality is due to orthogonality of the basis vectors $|i_1,\dots,i_t\rangle$ and $|i'_1,\dots,i'_t\rangle$ whenever $(i_1,\dots,i_t) \neq (i'_1,\dots,i'_t)$. □

Lemma 3. *The orthogonal projector onto* $\mathsf{Sym}_t(\mathbb{C}^d)$ *is*

$$\Pi^{d,t}_{sym} = \frac{1}{t!}\sum_{\sigma \in S_t} P_d(\sigma)$$

Proof. First, note that for any $\rho \in S_t$,

$$P_d(\rho) \cdot \Pi_{sym}^{d,t} = \frac{1}{t!} \sum_{\sigma \in S_t} P_d(\rho) P_d(\sigma) = \frac{1}{t!} \sum_{\sigma \in S_t} P_d(\rho\sigma) = \frac{1}{t!} \sum_{\sigma \in S_t} P_d(\sigma) = \Pi_{sym}^{d,t} \tag{2}$$

Similarly,

$$\Pi_{sym}^{d,t} P_d(\rho) = \Pi_{sym}^{d,t}$$

It is then straightforward to check that

$$(\Pi_{sym}^{d,t})^2 = \Pi_{sym}^{d,t}$$

which is a necessary and sufficent condition for $\Pi_{sym}^{d,t}$ to be an orthogonal projector.

We now prove that $\mathsf{Im}(\Pi_{sym}^{d,t}) = \mathsf{Sym}_t(\mathbb{C}^d)$. In one direction, for any $|\psi\rangle$ and any $\rho \in S_t$,

$$P_d(\rho) \cdot \Pi_{sym}^{d,t} |\psi\rangle = \Pi_{sym}^{d,t} |\psi\rangle$$

by Equation 2, meaning that $\mathsf{Im}(\Pi_{sym}^{d,t}) \subseteq \mathsf{Sym}_t(\mathbb{C}^d)$. In the other direction, if $|\psi\rangle \in \mathsf{Sym}_t(\mathbb{C}^d)$,

$$\Pi_{sym}^{d,t} |\psi\rangle = \frac{1}{t!} \sum_{\sigma \in S_t} P_d(\sigma) |\psi\rangle = \frac{1}{t!} \cdot t! \cdot |\psi\rangle = |\psi\rangle$$

meaning that $|\psi\rangle \in \mathsf{Im}(\Pi_{sym}^{d,t})$. Thus, $\mathsf{Sym}_t(\mathbb{C}^d) \subseteq \mathsf{Im}(\Pi_{sym}^{d,t})$ as well, establishing equivalence. □

To show the next, and the most crucial, theorem, we need Schur's lemma from representation theory. We state it here in elementary language, but refer the reader to a standard textbook, e.g. [Ser77], for the (relatively simple) proof.

Theorem 4.

$$\mathbb{E}_{|\psi\rangle \leftarrow S(d)}[|\psi\rangle \langle\psi|^{\otimes t}] = \frac{\Pi_{sym}^{d,t}}{\mathsf{Tr}(\Pi_{sym}^{d,t})}$$

where $\mathsf{Tr}(\Pi_{sym}^{d,t}) = \mathsf{dim}(\mathsf{Sym}_t(\mathbb{C}^d)) = \binom{t+d-1}{t}$.

Proof. Letting

$$\rho := \mathbb{E}_{|\psi\rangle \leftarrow S(d)}[|\psi\rangle \langle\psi|^{\otimes t}] ,$$

we first note that for any unitary U,

$$U^\dagger \rho U = \rho$$

An application of Schur's lemma tells us that ρ is proportional to the identity operator on the space $[d]^n$ which is indeed $\Pi_{sym}^{d,t}$. The right normalization constant is indeed the trace of $\Pi_{sym}^{d,t}$ since the trace of the LHS is 1. □

For more facts about and applications of the symmetric subspace, we refer the reader to Harrow's excellent survey [Har13].

3 Pseudorandom Quantum States (PRS): Definition

Definition 5 (Pseudorandom Quantum State). *Let $\mathcal{K} = \{K_\lambda \subseteq \{0,1\}^\lambda\}_{\lambda \in \mathbb{N}}$ be a subset of strings that define the key-space of the PRS. For a given λ and $n \in \mathbb{N}$, a family of states*

$$\Phi_\lambda := \left\{ |\phi_k\rangle \in S(2^n) \right\}_{k \in K_\lambda}$$

is a pseudorandom quantum state if:

- *There is a QPT algorithm* Gen *such that*

$$\mathsf{Gen}(1^\lambda, k, |0\rangle) = |\phi_k\rangle$$

- *For every polynomial function* $\mathsf{poly}(\cdot)$, $t = \mathsf{poly}(n, \lambda)$, *and every QPT adversary A,*

$$\left| \Pr_{k \leftarrow K_\lambda}[A(|\phi_k\rangle \langle \phi_k|^{\otimes t})] - \Pr_{|\phi\rangle \leftarrow S(2^n)}[A(|\phi\rangle \langle \phi|^{\otimes t})] \right| = \mathrm{negl}(\lambda) \tag{3}$$

  We will also write this succinctly as

$$\mathbb{E}_{k \leftarrow K_\lambda}[|\phi_k\rangle \langle \phi_k|^{\otimes t}] \approx_c \mathbb{E}_{|\phi\rangle \leftarrow S(2^n)}[|\phi\rangle \langle \phi|^{\otimes t}] \tag{4}$$

A related notion is that of a state t-design which can be viewed as both a strengthening and weakening of a PRS: a weakening in the sense that pseudorandomness is only required to hold given an a-priori bounded polynomial number of copies of the state, namely $t = t(n)$ copies, and a strengthening in the sense that the states on both sides of equation 4 are required to be *identical* and not merely computationally indistinguishable. In this sense, the condition defining a state t-design is analogous to t-wise independence, whereas that defining a pseudorandom quantum state is analogous to (computational) pseudorandomness.

Definition 6 (State t-design). *An ensemble $\nu = \{p_i, |\psi_i\rangle\}$ over d-dimensional states is a state t-design if*

$$\mathbb{E}_{|\psi\rangle \leftarrow \nu}[|\psi\rangle \langle \psi|^{\otimes t}] = \mathbb{E}_{|\psi\rangle \leftarrow S(d)}[|\psi\rangle \langle \psi|^{\otimes t}] \tag{5}$$

A weakening of the notion of state t-design asks for equation 5 to be approximate, in the sense that the LHS are RHS are ε-close in trace distance for some $\varepsilon = \varepsilon(n)$.

4 PRS Construction

We are now ready to show how to construct a pseudo-random quantum state. We will refer to the construction as a binary phase PRS, for a reason that will become clear shortly.

Theorem 7. *If post-quantum secure one-way functions exist, then there exists a family of pseudo-random quantum states.*

Construction. Since post-quantum secure one-way functions imply post-quantum secure pseudorandom functions, we will use a pqPRF family $\mathcal{F} = \{F_k\}_{k \in \{0,1\}^\lambda}$ as a building block.

- $\mathsf{Gen}(1^\lambda)$ samples a PRF key $k \leftarrow \{0,1\}^\lambda$.
- $\mathsf{Eval}(k)$ outputs the quantum state

$$|\psi_k\rangle = \frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} (-1)^{F_k(x)} |x\rangle$$

We first show that Eval runs in quantum polynomial-time, making a single quantum query to the PRF. To prepare $|\psi_k\rangle$, Eval starts by preparing

$$H^{\otimes n} |0\rangle = \frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} |x\rangle_{\mathcal{X}}$$

Then, compute the PRF F_k in superposition to get

$$\frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} |x\rangle_{\mathcal{X}} |F_k(x)\rangle_{\mathcal{Y}}$$

Compute $I_{\mathcal{X}} \otimes Z_{\mathcal{Y}}$ (that is, Z acting on the second register) to get

$$\frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} (-1)^{F_k(x)} |x\rangle_{\mathcal{X}} |F_k(x)\rangle_{\mathcal{Y}}$$

Finally, uncompute F_k to get the PRS state (where we suppress the register name)

$$|\psi_k\rangle := \frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} (-1)^{F_k(x)} |x\rangle$$

Since all phases are binary, this construction is also called a binary phase PRS.

Security. We now prove the security of this construction. Recall that the goal is to show that

$$\mathbb{E}_{k \leftarrow K}[|\phi_k\rangle \langle \phi_k|^{\otimes t}] \quad \text{and} \quad \mathbb{E}_{|\psi\rangle \leftarrow S(d)}[|\psi\rangle \langle \psi|^{\otimes t}]$$

are computationally indistinguishable to any quantum polynomial-time distinguisher, for any $t = \mathsf{poly}(\lambda)$. Consider the following sequence of hybrid expressions, each of which defines a mixed state.

- **Hybrid** 1 is the (mixed) state

$$\mathbb{E}_{k \leftarrow K}[|\phi_k\rangle \langle \phi_k|^{\otimes t}]$$

  where F_k is the PRF.

- **Hybrid** 2 is the (mixed) state

$$\mathbb{E}_{f \leftarrow F}[|\phi_f\rangle \langle \phi_f|^{\otimes t}]$$

  where f is a uniformly random function from $\{0,1\}^n$ to $\{0,1\}$ and

$$|\phi_f\rangle = \frac{1}{\sqrt{2^n}} \cdot \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle$$

- **Hybrid** 3. To define Hybrid 3, let's write out explicitly the state in Hybrid 2 first:

$$\mathbb{E}_{f\leftarrow F}[|\phi_f\rangle\langle\phi_f|^{\otimes t}] = 2^{-nt}\cdot \sum_{x_1,\ldots,x_t,y_1,\ldots,y_t} \mathbb{E}_{f\leftarrow F}[(-1)^{f(x_1)+\ldots+f(x_t)+f(y_1)+\ldots+f(y_t)}]\,|x_1,\ldots,x_t\rangle\langle y_1,\ldots,y_t|$$

  Project this state onto the subspace of distinct $x_1,\ldots,x_t$ and distinct $y_1,\ldots,y_t$, and post-select on the event that the x_i are distinct and the y_i are distinct. This will be the state

$$\frac{1}{\binom{2^n+t-1}{t}}\cdot \sum_{\substack{x_1,\ldots,x_t \text{ distinct,}\\ y_1,\ldots,y_t \text{distinct}}} \mathbb{E}_{f\leftarrow F}[(-1)^{f(x_1)+\ldots+f(x_t)+f(y_1)+\ldots+f(y_t)}]\,|x_1,\ldots,x_t\rangle\langle y_1,\ldots,y_t| \tag{6}$$

- **Hybrid** 4 is the (mixed) state

$$\mathbb{E}_{|\psi\rangle\leftarrow S(d)}[|\psi\rangle\langle\psi|^{\otimes t}]$$

Lemma 8. *Hybrids* 1 *and* 2 *are computationally indistinguishable.*

Proof. The states in hybrid 1 and 2 can be prepared with t oracle queries either to a function f_k chosen at random from $\mathcal{F}$ or to a uniformly random function f. By the post-quantum security of the PRF family $\mathcal{F}$ (against an adversary that can make quantum superposition queries), no quantum polynomial-time distinguisher can tell these two oracles apart, which in turns implies that the two hybrids are computationally indistinguishable. □

Lemma 9. *Hybrids* 2 *and* 3 *are statistically indistinguishable.*

Proof. This follows by the gentle measurement lemma (Lemma 11) as the only difference between hybrid 3 can be obtained from hybrid 2 by projecting to the subspace of $[2^n]^t$ where all t components are different. This happens with overwhelming probability as long as $t \ll 2^{n/2}$, therefore an application of gentle measurement lemma does the job. □

Lemma 10. *Hybrids* 3 *and* 4 *are identical.*

Proof. The expectation in equation 6 is 0 whenever $x_1,\ldots,x_t$ is *not* a permutation of $y_1,\ldots,y_t$ and 1 otherwise. Then, this expectation is exactly $\frac{\Pi_{sym}^{d,t}}{\mathsf{Tr}(\Pi_{sym}^{d,t})}$ which is precisely Hybrid 4, by Theorem 4. □

5 Supporting Lemmas

Lemma 11 (Gentle Measurement Lemma)**.** *Let ρ be a state and Π be a projector such that*

$$\mathsf{Tr}(\Pi\rho) \geq 1-\varepsilon$$

for some $\varepsilon \geq 0$. Then,

$$\mathsf{TD}\left(\rho, \frac{\Pi\rho\Pi}{\mathsf{Tr}(\Pi\rho)}\right) \leq \sqrt{\varepsilon}$$

Proof. First consider the case when ρ is a pure state $|\psi\rangle\langle\psi|$. Then, the fidelity between ρ and the post-measurement state $\frac{\Pi\rho\Pi}{\mathsf{Tr}(\Pi\rho)}$ is

$$\frac{\langle\psi|\Pi|\psi\rangle\langle\psi|\Pi|\psi\rangle}{\langle\psi|\Pi|\psi\rangle} = \langle\psi|\Pi|\psi\rangle \geq 1-\varepsilon$$

where the last equality is by assumption. By the monotonicity of the fidelity function (Lemma **??**), this is true for a general mixed state as well. That is,

$$\mathsf{F}\left(\rho, \frac{\Pi\rho\Pi}{\mathsf{Tr}(\Pi\rho)}\right) \geq 1 - \varepsilon$$

Using the relation between the fidelity and trace distance (Lemma 12) finishes up the proof. □

Lemma 12 (Fidelity and Trace Distance)**.** *The following bound applies to the trace distance and the fidelity between two quantum states $\rho, \sigma \in \mathcal{S}(H)$ that live on a Hilbert space H:*

$$1 - \sqrt{\mathsf{F}(\rho,\sigma)} \leq \mathsf{TD}(\rho,\sigma) \leq \sqrt{1 - \mathsf{F}(\rho,\sigma)}$$

References

[Har13] Aram W. Harrow. The church of the symmetric subspace, 2013.

[Ser77] Jean-Pierre Serre. *Linear representations of finite groups.*, volume 42 of *Graduate texts in mathematics.* Springer, 1977.