6.S895: Quantum Cryptography

Lecture: Pseudo-random Quantum States and Commitments

Lecturers: Luowen Qian **Scribe:** Vinod Vaikuntanathan and Anand Natarajan

1 Applications of Pseudorandom Quantum States

We have seen so far that in the quantum world, one-way functions give us pseudorandom quantum states; one-way functions imply commitments; and commitments imply oblivious transfer and secure two-party computation.

We will introduce the notion of EFI pairs and show how to complete the picture. The high-level takeaway from today's lecture is that one can build commitments and quantum cryptography from pseudorandom quantum states (and in fact, an even weaker primitive called an EFI pair) which could exist even in a world where $\mathsf{BQP} = \mathsf{QMA}$.

Definition 1 (EFI Pairs)**.** *An EFI pair is a family of pairs of mixed states* $(\{\rho_{0,\lambda}\}_{\lambda\in\mathbb{N}}, \{\rho_{1,\lambda}\}_{\lambda\in\mathbb{N}})$ *where*

- Efficient Generation: *There is a quantum polynomial-time algorithm* $G(1^\lambda, b) = \rho_{b,\lambda}$.
- Statistical Farness: *For all* λ*,*
$$\mathsf{TD}(\rho_{0,\lambda}, \rho_{1,\lambda}) \geq 1 - 2^\lambda$$
- Computational Indistinguishability: *The following two ensembles are computationally indistinguishable:*
$$\{\rho_{0,\lambda}\}_{\lambda\in\mathbb{N}} \approx_c \{\rho_{1,\lambda}\}_{\lambda\in\mathbb{N}}$$

Classically, the analogous notion would be two distributions that are statistically far but computationally indistinguishable. A natural construction of commitments from such distribution pairs is the following: in order to commit to a bit b, the committer generates a sample from distribution D_b; to open a commitment to a bit b, she simply reveals the randomness used for sampling. This is unfortunately not binding! This is because an adversary can search for a random string that generates elements in the intersection $\mathsf{Supp}(D_0) \cap \mathsf{Supp}(D_1)$. However, by a much more sophisticated construction, such distributions can be shown to imply one-way functions (Goldreich), which in turn imply pseudorandom generators (Hastad-Impagliazzo-Levin-Luby), which in turn imply commitments (Naor). So our intuition (a priori) should be that EFI might be much weaker than quantum commitments.

- Example 1: From a (quantum) PKE:
$$(pk, \mathsf{Enc}(pk, 0)) \approx_c (pk, \mathsf{Enc}(pk, 1)$$
Note that this is overkill: we don't need the *efficient* decryptability at all to get EFI from this construction.

- Example 2: From a statistically binding commitment. The two states are the view of the receiver when committing to 0 vs. 1. Binding of the EFI pair follows from the statistical binding of the commitment which tells us that the fidelity is small, and therefore the trace distance (between the two views, and therefore the two states in the EFI pair) is large.
- Example 3: PRS, which we describe in more detail below.

Lemma 2 (PRS Implies EFI Pair)**.** *Let $G : \{0,1\}^\lambda \to S(2^n)$ is a secure PRS against t copies. Then, there is an EFI pair if*

$$\binom{2^n + t - 1}{t} \gg 2^\lambda$$

Parameter settings:

- $n \geq \lambda + 1$. $t = 1$ suffices.
- $n \geq \log_2 \lambda$. $t = \lambda + 1$ copies suffice.
- What if $n = O(1)$? Seems like you need t needs to be superpolynomial in λ acc. to the theorem statement above, but is that really necessary? Seems like

Proof. We will let

- $\rho_0 = \{G(k)^{\otimes t}\}$, i.e. $\mathbb{E}_{k \in \{0,1\}^\lambda} G(k)^{\otimes t}$. This state has rank at most 2^λ since each $G(k)$ is pure.
- $\rho_1 \propto \Pi_{sym}^{2^n,t}$. This is the maximally mixed state on the symmetric subspace, and its rank is the dimension of this subspace, which is exactly

$$\binom{2^n + t + 1}{t}.$$

.

If the condition in the theorem statement is true, then the rank of ρ_1 is much larger than the rank of ρ_0. We can show a distinguisher whose success probability is at least

$$1 - \frac{\mathsf{rank}(\rho_0)}{\mathsf{rank}(\rho_1)} \geq 1 - \frac{2^\lambda}{\binom{2^n+t-1}{t}} \geq 1/2$$

for an appropriate choice of parameters. At the same time, ρ_0 and ρ_1 are computationally indistinguishability by the t-copy security of the PRS. □

Theorem 3 (EFI Pairs Imply Commitments)**.** *The existence of an EFI pair implies the existence of a statistically binding, computationally hiding quantum commitment (to a classical bit).*

The construction can be traced back to Chailloux, Kerenidis and Rosgen from 2011.

Proof. Purify the generation algorithm G which now generates pure states $|\psi'_b\rangle_{OP}$ where $\mathsf{Tr}_P |\psi'_b\rangle = \rho_b$. The committer sends the O register to commit to b; and to reveal, it sends the P register. Hiding follows immediately from the computational indistinguishability of the EFI pair. Binding follows from the fact that the trace distance between the pure states $|\psi'_0\rangle_{OP}$ and $|\psi'_1\rangle_{OP}$ is large. □

Can we go from EFI to PRS? We don't know. Can we go from PRS to OWF? We again don't know, but we will now see evidence that this shouldn't be possible. If pqOWF exist, then $\mathsf{BQP} \not\subseteq \mathsf{QMA}$ and indeed, $\mathsf{BQP} \not\subseteq \mathsf{QCMA}$. We will show that PRUs do not imply $\mathsf{BQP} \not\subseteq \mathsf{QCMA}$, in the sense that there is an oracle world where $\mathsf{BQP} = \mathsf{QCMA}$ but PRS (and even PRU) exist.

If one-way functions do not exist, every binary phase state is breakable efficiently.

Thus, this is evidence that there is a separation between classical cryptographic assumptions like OWF, which imply classical complexity class separations, and quantum assumptions which do not have such implications.

Definition 4. *G is a pseudorandom unitary if*

- *G is efficient: There is a unitary U_k such that*

$$G(k, |\psi\rangle) = U_k |\psi\rangle$$

- *For every QPT A,*

$$|\Pr_{k\leftarrow\{0,1\}^\lambda}[A^{U_k}(1^\lambda) = 1] - \Pr_{U\leftarrow S(2^n)}[A^U(1^\lambda) = 1] = \mathsf{negl}(\lambda)$$

We can construct a PRU secure against non-adaptive queries from one-way functions; this is a very recent result of Metger et al. It is still open whether one can get security against completely adaptive queries and access to U_k as well as $U_k^\dagger$.

Claim 5. *(n,t)-PRU $\Rightarrow$ (n,t)-PRS.*

Proof. Just run the PRU on input $|0^n\rangle$. □

It is also open to construct a non-adaptive PRU from PRS. Classically, note that we do know how to get pseudorandom permutations (the analogous object to PRUs) from pseudorandom generators (the analogous object to PRFs).

Relating this to BQP and friends A QMA machine always takes classical inputs by definition (even if it takes a quantum witness state). Thus, it's not obvious how to use it to distinguish a PRS or PRU (truly quantum objects) from a random state or unitary.

Theorem 6 (Kretschmer 2021)**.** *There exists a unitary (quantum) oracle U such that $\mathsf{BQP}^U = \mathsf{QMA}^U$, yet there exists a PRU with respect to U.*

Proof. Let's try taking U to be a "keyed" Haar random unitary. Specifically, take

$$U |k\rangle |\psi\rangle = U_k |\psi\rangle ,$$

where for each $k \in \{0,1\}^{|k|}$, the unitary U_k is sampled from the Haar measure over $n \times n$ unitaries. The purpose of this is to trivialize the construction of PRUs.

Let us now add a "classical" part to this oracle, denoted P. This will be used to collapse the complexity classes. Concretely, let's take P to be an oracle for some PSPACE-complete language.

Existence of PRU: The PRU is of course $\{U_k\}_k$. We claim that

$$|\Pr[A^{U,P,U_k}(1^\lambda)] - \Pr[A^{U,P,V}(1^\lambda)]| \leq \mathrm{negl}(\lambda),$$

where V is a Haar random unitary.

We do this through hybrids. Define U' to be the "punctured" version of U, that is identical to U except on input k, where it outputs V instead. Then we have

- H_0 : $A^{U,P,U_k}(1^\lambda)$ the initial quantity.
- H_1 : $A^{U',P,U_k}(1^\lambda)$. This is close to H_0 by the Grover lower bound (since otherwise, A could find k).
- H_2 : $A^{U,P,V}(1^\lambda)$. This has exactly the same success probability as the previous.

Showing $\mathsf{BQP}^U = \mathsf{QMA}^U$ This is not so easy, because of the access to the unitary part U of the oracle. To prove this, we will need to use facts about Haar-random matrices.

Fact 7. *Let f be an L-Lipschitz real function in the Frobenius norm, i.e.*

$$|f(U) - f(V)| \leq L \cdot \|U - V\|_F.$$

Then $\forall \Delta > 0$, we the concentration bound

$$\Pr_{U \sim \mu_d}[f(U) \geq \mathbf{E}_{V \sim \mu_d}[f(V)] + \Delta] \leq \exp(-\frac{(d-2)\Delta^2}{24L^2}).$$

Note that this concentration is inverse exponential in the dimension d, which is in fact *doubly* inverse exponential in the number of qubits n.

We will apply this with f being the probability that a BQP^U machine outputs 1 on some input.

Fact 8. *If A^U makes T queries to U, then $f(U) = \Pr[A^U = 1]$ is $2T$-Lipschitz.*

Corollary 9. *Morally, $\mathsf{BQP}^U = \mathsf{BQP}$.*

Lemma 10. *For a QMA^U verifier A, define*

$$g(U) := \max_{|\phi\rangle} \Pr[A^U(|\phi\rangle) = 1].$$

Then $g(U)$ is $2T$-Lipschitz.

Note that we are working with *unnormalized* Frobenius norm, so this makes sense.

Proof. Fix U, V. Let $|\psi\rangle, |\phi\rangle$ be the maximizers.

$$\begin{aligned}
&|\max \Pr[A^U(\cdot)] - \max \Pr[A^V(\cdot)]| && (1)\\
&= |\Pr[A^U(|\psi\rangle) = 1] - \Pr[A^V(|\phi\rangle) = 1]| && (2)\\
&= \max\{\Pr[A^U(|\psi\rangle) = 1] - \Pr[A^V(|\phi\rangle) = 1], \Pr[A^V(|\phi\rangle) = 1] - \Pr[A^U(|\psi\rangle) = 1]\} && (3)\\
&\leq \max\{\Pr[A^U(|\psi\rangle) = 1] - \Pr[A^V(|\psi\rangle) = 1], \Pr[A^V(|\phi\rangle) = 1] - \Pr[A^U(|\phi\rangle) = 1]\} \quad \leq 2T\|U - V\|_F. && (4)
\end{aligned}$$

□

We are now in a position to prove

$$\mathsf{BQP}^{U,P} = \mathsf{QMA}^{U,P}.$$

Given a T-query verifier $V^{U,P}$. First, as a technicality, use unitary tomography on U for all input lengths up to $100 \log T$. The second step is for all U of length bigger than $100 \log T$, replace it by approximate T-designs. The concentration bound shows us that the new V can be replaced by one that does not query U at all. Finally, we know that with only the classical oracle P and not the unitary oracle U, the two classes are equal. □