QALCO: Quantum Algorithms and Complexity **2025 - 2026**

Lecture 8: Quantum pseudorandom states

Lecturer: Alex Bredariol Grilo (alex.bredariol-grilo@lip6.fr)

1 Classical Pseudorandom Generators and Functions

Definition 1 (Classical PRG)**.** *A function*

$$G : \{0,1\}^n \to \{0,1\}^m$$

with $m > n$ is a pseudorandom generator (PRG) if:

1. *G is computable in polynomial time,*
2. *For every quantum polynomial-time distinguisher D,*

$$\left| \Pr_{k \leftarrow \{0,1\}^n}[D(G(k)) = 1] - \Pr_{y \leftarrow \{0,1\}^m}[D(y) = 1] \right| \leq negl(n).$$

The quantity $m - n$ is called the *stretch*. A fundamental fact is that stretch 1 already implies arbitrary polynomial stretch via iteration.

Time restriction. The polynomial-time restriction on D is essential: without it, exhaustive search over all seeds breaks any PRG.

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

$$F : \{0,1\}^n \times \{0,1\}^m \to \{0,1\}$$

is a pseudorandom function if for all quantum polynomial-time distinguishers D with oracle access,

$$\left| \Pr_k[D^{F_k} = 1] - \Pr_{f \leftarrow \mathcal{F}}[D^f = 1] \right| \leq negl(n),$$

where $\mathcal{F}$ is the set of all functions $\{0,1\}^m \to \{0,1\}$.

Notice that here, we assume that the distinguisher has *quantum access* to the pseudorandom/random function.

Definition 3 (One-Way Function)**.** *A function $f : \{0,1\}^n \to \{0,1\}^n$ is one-way if:*

1. *f is polynomial-time computable,*
2. *For every quantum polynomial-time algorithm A,*

$$\Pr_{x \leftarrow \{0,1\}^n}[f(A(f(x))) = f(x)] \leq negl(n).$$

Theorem 1 (Håstad–Impagliazzo–Levin–Luby)**.** *One-way functions exist if and only if pseudorandom generators exist.*

Theorem 2 (Goldreich–Goldwasser–Micali)**.** *Pseudorandom generators exist if and only if pseudorandom functions exist.*

We notice that such results were proven only in the classical setting, but they were also lifted to the post-quantum setting (i.e. the distinguisher is quantum).

Thus:

$$\text{OWF} \iff \text{PRG} \iff \text{OWF}.$$

1.1 Impagliazzo's Five Worlds

Impagliazzo described five possible "worlds" depending on the existence and strength of cryptographic primitives.

- **Algorithmica/Heuristica:** $\mathsf{P} = \mathsf{NP}$/We can solve NP-complete problems in practice. No meaningful cryptography exists.
- **Pessiland:** Average-case hardness exists but no one-way functions.
- **Minicrypt:** One-way functions exist (hence PRGs and PRFs), but public-key cryptography may not.
- **Cryptomania:** Public-key cryptography exists.

2 Quantum Pseudorandom States (PRS)

Let m be polynomial in n.

Definition 4 (Pseudorandom States)**.** *A generator*

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

is a pseudorandom state (PRS) generator if:

1. *G is quantum polynomial-time computable,*
2. *For all $t = \mathrm{poly}(n)$ and all polynomial-time quantum distinguishers D,*

$$\left| \Pr_k[D(|\psi_k\rangle^{\otimes t}) = 1] - \Pr_{|\theta\rangle \sim Haar}[D(|\theta\rangle^{\otimes t}) = 1] \right| \leq negl(n),$$

where we define $|\psi_k\rangle = G(k)$.

The parameter t is essential due to the no-cloning theorem: distinguishability may emerge only when multiple copies are available. In this class we will see the following:

1. We can break PRS with an oracle to PP.
2. OWF $\Rightarrow$ PRS
3. PRS $\overset{BB}{\not\Rightarrow}$ OWF, where BB stands for Black-box reductions. This means that if we are given a black-box that gives us PRS, we cannot construct OWFs solely from it.

3 Breaking PRS with PP

In this section we explain the key idea behind Kretschmer's result that *any* pseudorandom state generator can be broken in PP [5].

3.1 Definitions

Definition 5 (PP)**.** *The complexity class* PP *(Probabilistic Polynomial-Time) is the class of decision problems $L \subseteq \{0,1\}^*$ for which there exists a deterministic polynomial-time Turing machine M such that for every input x:*

$$x \in L \;\Rightarrow\; \Pr_r[M(x,r)=1] > \tfrac{1}{2},$$

$$x \notin L \;\Rightarrow\; \Pr_r[M(x,rr)=1] \leq \tfrac{1}{2}.$$

Theorem 3 (Toda's theorem)**.** $\mathsf{PH} \subseteq \mathsf{P}^{\mathsf{PP}}$.

Definition 6 (PostBQP)**.** *The complexity class* PostBQP *(Postselected Bounded-Error Quantum Polynomial-Time) is the class of decision problems $L \subseteq \{0,1\}^*$ for which there exists a uniform family of polynomial-size quantum circuits $\{Q_x\}$ such that for every input x:*

- *The circuit Q_x produces two designated output bits:*
  - *a* postselection *bit p,*
  - *a* decision *bit b.*
- *The probability that the postselection event occurs is nonzero:*

$$\Pr[p=1] > 0.$$

- *Conditioned on the event $p=1$, the decision bit satisfies:*

$$x \in L \;\Rightarrow\; \Pr[b=1 \mid p=1] \geq \tfrac{2}{3},$$

$$x \notin L \;\Rightarrow\; \Pr[b=1 \mid p=1] \leq \tfrac{1}{3}.$$

The probabilities are taken over the measurement outcomes of the quantum circuit, and the circuit size is polynomial in $|x|$.

Theorem 4 ([1])**.** $\mathsf{PostBQP} = \mathsf{PP}$.

4 The attack

For the attack, we are going to use the following two lemmas.

Lemma 1 (Classical Shadows [3])**.** *Let $O_1, \ldots, O_N$ be N observables acting on m qubits, and let ρ be an unknown m-qubit mixed state. For parameters $\varepsilon, \delta > 0$, there exists a quantum procedure that performs*

$$T = O\left(\frac{\log(N/\delta)}{\varepsilon^2}\right)$$

single-copy measurements of ρ in randomly chosen Clifford bases and produces classical data (called classical shadows*) from which one can estimate all expectation values*

$$\mathrm{Tr}(O_i \rho)$$

simultaneously for every $i \in [N]$ up to additive error ε with probability at least $1-\delta$.

Lemma 2. *For any fixed $|\varphi\rangle \in \mathbb{C}^N$, and let $\varepsilon > 0$. Then:*

$$\Pr_{|\psi\rangle\sim Haar}\left[|\langle\psi|\varphi\rangle|^2 \geq \varepsilon\right] \leq e^{-\varepsilon(N-1)}.$$

We are going to construct a distinguisher $\mathcal{D}$ that receives a state $|\phi\rangle^{\otimes T}$, where either $|\phi\rangle = |\psi_k\rangle$ is a pseudorandom state for a random $k \in \{0,1\}^n$, or $|\phi\rangle = |\theta\rangle$, for a Haar-random state $|\theta\rangle$. We let $b = 0$ if $\mathcal{D}$ is given copies of a pseudorandom state or $b = 1$ otherwise.

Let us define the observables $O_k = |\psi_k\rangle\langle\psi_k|$ Notice that for all k, the original state $|\phi\rangle$ was a pseudorandom state, there exists a $k \in \{0,1\}^n$ such that $Tr(O_k |\phi\rangle\langle\phi|) = 1$. On the other hand, we have that using union bound on Lemma 2 gives that us that with probability $2^n 2^{-(2^m-1)/3}$ over a Haar random $|\phi\rangle$, $Tr(O_k |\phi\rangle\langle\phi|) \leq \frac{1}{3}$ for all $k \in \{0,1\}^n$.

We can then use then Lemma 1 with $N = 2^n$, $\varepsilon = \frac{1}{3}$ and $\delta \approx 0.001$, and it gives us an (exponential) algorithm $\mathcal{C}$ that, with probability close to 0.999, uses the measurement outcomes of the state $|\phi\rangle$ and compute $Tr(O_k |\phi\rangle\langle\phi|)$ for all ks with additive error $\frac{1}{3}$. In this case, using these values, $\mathcal{C}$ can check if there exists a k such that $Tr(O_k |\theta\rangle) \geq \frac{2}{3}$, which distinguishes pseudorandom states from Haar random ones with high probability.

Notice that $\mathcal{C}$ cannot be directly used as the distinguisher since it runs in exponential time. But we will now construct another algorithm $\mathcal{D}$ that uses only the fact that $\mathcal{C}$ exists to distinguish pseudorandom states from Haar random ones.

The distinguisher $\mathcal{D}$ will pick T random Clifford circuits $C_1, ..., C_T$, and will apply C_i to the i-th copy of the received state with outcomes $s_1, ..., s_T$. We are going to construct a PostBQP algorithm $\mathcal{B}$ that, given $(C_1, ..., C_T, s_1, ..., s_T)$, distinguishes if it comes from pseudorandom states of Haar random ones with high probability. Using a query to PP, which is equivalent to PostBQP, the distinguisher $\mathcal{D}$ can learn the output of $\mathcal{B}$ on $(C_1, ..., C_T, s_1, ..., s_T)$, which finishes the proof. Thus, we focus now on constructing $\mathcal{B}$.

First, $\mathcal{B}$ constructs the following state:

$$\rho = \frac{1}{2}|0\rangle\langle 0| \otimes \mathbb{E}_k[|\psi_k\rangle\langle\psi_k|^{\otimes T}] + \frac{1}{2}|0\rangle\langle 0| \otimes \mathbb{E}_{|\theta\rangle}[|\theta\rangle\langle\theta|^{\otimes T}]. \tag{1}$$

$\mathcal{B}$ will then apply the Cliffords $C_1, ..., C_T$ on the second register, and post-select on the measurement being $(s_1, ..., s_T)$. Finally, $\mathcal{B}$ measures the first bit, gets the measurement outcome X and outputs it. We will argue that the guess from $\mathcal{B}$ is correct with probability ≥ 0.995.

For $i \in \{0,1\}$, let $p_i = \Pr[X = i | C_1, ..., C_T, s_1, ..., s_T]$. Notice that from Bayes decision rule, we have that for every function f

$$Pr[f(C_1, ..., C_T, s_1, ..., s_T) = X] \leq Pr[\arg\max_i p_i = X]. \tag{2}$$

Since $\mathcal{C}$ can use $C_1, ..., C_T, s_1, ..., s_T$ to correctly guess if the state is pseudorandom or random with probability 0.999, we have that $Pr[\arg\max_i p_i = X] \geq Pr[X = b] \geq 0.999$. We can use the law of total expectation and Markov's inequality to show that $\mathcal{B}$ makes the correct measurement with proability at least $\frac{2}{3}$ (**Exercise**).

The only remaining problem is that $\mathcal{B}$ is not currently efficient! Constructing the Haar random states from Equation (1) takes exponential time. For that, we can use the following object:

Definition 7. *An ensemble of pure states $\{(p_i, |\psi_i\rangle)\}_i$ over $\mathcal{H}$ is called an* ε-approximate state T-design *if*

$$\left\| \sum_i p_i \left(|\psi_i\rangle\langle\psi_i|\right)^{\otimes t} - \mathbb{E}_{\theta\sim\text{Haar}}\left[\left(|\theta\rangle\langle\theta|\right)^{\otimes t}\right] \right\|_1 \leq \varepsilon.$$

These objects can be constructed efficiently and unconditionally for every fixed T.

We can replace the Haar random states from Equation (1), which will incurr in some loss in the correct guess of the distinguisher, but we will still have a constant advantage.

5 Constructing PRS from PRFs

We now present the Ji–Liu–Song construction [4], whose security proof was originally proven by [2]. Let

$$F_k : \{0,1\}^m \to \{0,1\}$$

for $k \in \{0,1\}^n$ be a PRF.

Define:

$$|\psi_k\rangle = \frac{1}{2^{m/2}} \sum_{x \in \{0,1\}^m} (-1)^{F_k(x)} |x\rangle.$$

This is efficiently preparable by:

1. Preparing uniform superposition,
2. Computing $F_k(x)$ in superposition,
3. Applying a phase kickback.

We must show that for every $t = \text{poly}(n)$ and every quantum polynomial-time distinguisher D,

$$\left| \Pr_k[D(|\psi_k\rangle^{\otimes t}) = 1] - \Pr_{|\theta\rangle \sim \text{Haar}}[D(|\theta\rangle^{\otimes t}) = 1] \right| \leq \text{negl}(n).$$

The proof proceeds in two hybrids:

- Hybrid 1: Replace the PRF F_k with a truly random function f.
- Hybrid 2: Show that random phase states approximate Haar states in t-th moments.

Hybrid 1: PRF to Truly Random Function Define:

$$|\psi_f\rangle = \frac{1}{\sqrt{M}} \sum_x (-1)^{f(x)} |x\rangle$$

where $f : \{0,1\}^m \to \{0,1\}$ is uniformly random.

Lemma 3 (PRF Hybrid)**.** *For every efficient distinguisher D,*

$$\left| \Pr_k[D(|\psi_k\rangle^{\otimes t}) = 1] - \Pr_f[D(|\psi_f\rangle^{\otimes t}) = 1] \right| \leq \text{negl}(n).$$

Proof. Given oracle access to D, we build a PRF distinguisher that simulates $|\psi_k\rangle^{\otimes t}$ using superposition queries to F_k. If F_k is replaced by a random function f, the output state becomes $|\psi_f\rangle^{\otimes t}$. Any non-negligible distinguishing advantage would violate PRF security. □

Thus it suffices to show that random phase states are t-wise indistinguishable from Haar states.

Hybrid 2: From random phases to Haar random states. Let us define

$$\rho_t^{\text{phase}} = \mathbb{E}_f \left[|\psi_f\rangle \langle\psi_f|^{\otimes t} \right] \text{ and } \rho_t^{\text{Haar}} = \mathbb{E}_\theta \left[|\theta\rangle \langle\theta|^{\otimes t} \right].$$

Security reduces to bounding:

$$||\rho_t^{\text{phase}} - \rho_t^{\text{Haar}}||_1.$$

Using tools from quantum information theory, we can show that

$$\rho_t^{\text{Haar}} = \frac{\Pi_{\text{sym}}}{\text{Tr}(\Pi_{\text{sym}})},$$

where Π_{sym} is the projector onto the symmetric subspace of $(\mathbb{C}^M)^{\otimes t}$. Its dimension is $\dim(\text{Sym}^t(\mathbb{C}^M)) = \binom{M+t-1}{t}$, where $M = 2^m$.

On the other hand, we have that

$$|\psi_f\rangle^{\otimes t} = M^{-t/2} \sum_{x_1,\ldots,x_t} (-1)^{f(x_1)+\cdots+f(x_t)} |x_1, \ldots, x_t\rangle,$$

and thus

$$\rho_t^{\text{phase}} = M^{-t} \sum_{x,y} \mathbb{E}_f \left[(-1)^{\sum_i f(x_i)+\sum_i f(y_i)} \right] |x\rangle \langle y|.$$

Because f is uniformly random, the values $\{f(x)\}_x$ are independent unbiased bits, and thus

$$\mathbb{E}_f \left[(-1)^{\sum_i f(x_i)+\sum_i f(y_i)} \right] = 0$$

unless every input appears an even number of times in the multiset $\{x_1, \ldots, x_t, y_1, \ldots, y_t\}$. If we assume that each $x_i \neq x_{i'}$ for all $i \neq i'$ and analogously for all y_i's, this occurs if and only if $(y_1, \ldots, y_t)$ is a permutation of $(x_1, \ldots, x_t)$.

Notice that having a collision ($x_i = x_{i'}$ or $y_i = u_{i'}$ for some $i \neq i'$) happens with probability $O\left(\frac{t^2}{M}\right)$. Therefore, if we define $S_{m,t} = \{(x_1, \ldots, x_t) : x_i \text{ distinct}\}$, and

$$\hat{\rho}_t^{\text{phase}} = M^{-t} \sum_{x \in S_{m,t}} \sum_{\pi \in S_t} |x_1, \ldots, x_t\rangle \langle x_{\pi(1)}, \ldots, x_{\pi(t)}|,$$

we have that

$$||\rho_t^{\text{phase}} - \hat{\rho}_t^{\text{phase}}||_1^2 \leq O\left(\frac{t^2}{M}\right).$$

Let us define

$$A = \sum_{x \in S_{m,t}} |x\rangle \langle x|.$$

Then notice that

$$\hat{\rho}_t^{\text{phase}} = M^{-t} t! \, \Pi_{\text{sym}} A \Pi_{\text{sym}}$$

Since A acts almost like identity on the symmetric subspace (up to collision errors),

$$\|\hat{\rho}_t^{\text{phase}} - \rho_t^{\text{Haar}}\|_1 = \|\Pi_{\text{sym}} A \Pi_{\text{sym}} - \Pi_{\text{sym}}\|_1 = O\left(\frac{t^2}{M}\right).$$

Thus:

$$\|\rho_t^{\text{phase}} - \rho_t^{\text{Haar}}\|_1 \leq \|\rho_t^{\text{phase}} - \hat{\rho}_t^{\text{phase}}\|_1 + \|\hat{\rho}_t^{\text{phase}} - \rho_t^{\text{Haar}}\|_1 = O\left(\frac{t^2}{M}\right).$$

If $m = \omega(\log n)$ (and thus $M = 2^{\omega(\log n)}$ and $t = \text{poly}(n)$, then:

$$\frac{t^2}{M} = \text{negl}(n).$$

Thus:

$$\|\rho_t^{\text{phase}} - \rho_t^{\text{Haar}}\|_1 \leq \text{negl}(n).$$

6 Separation Between PRS and QMA = BQP

We will consider the quantum oracle $\mathcal{O} = (\mathcal{U}, \mathcal{C})$ proposed by [5] relative to which:

- $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$, and
- Pseudorandom States (PRS) and Pseudorandom Unitaries (PRUs) exist.

Notice that if $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$, then $\text{OWF}^{\mathcal{O}}$ cannot exist since $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$ can invert them. We define the oracle $\mathcal{O} = (\mathcal{U}, \mathcal{C})$ as follows

- **Quantum Oracle $\mathcal{U}$:** $\mathcal{U}$ is a sequence of unitaries $\{\mathcal{U}_n\}_{n\in\mathbb{N}}$, where each $\mathcal{U}_n$ is a direct sum of 2^n different Haar-random n-qubit unitaries. Specifically, $\mathcal{U}_n \sim \mu_{2^n}^{2^n}$, where μ_{2^n} is the Haar measure over $2^n \times 2^n$ unitaries. This means $\mathcal{U}_n$ consists of 2^n independent Haar-random unitaries.
- **Classical Oracle $\mathcal{C}$:** $\mathcal{C}$ is a language constructed deterministically and independently of $\mathcal{U}$. It is defined in stages, where for a string x, $\mathcal{C}(x) = 1$ if:
  1. x describes a quantum oracle circuit $\mathcal{V}^{\overline{\mathcal{U}},\mathcal{C}}(|\psi\rangle)$ that queries $\overline{\mathcal{U}}$ and $\mathcal{C}$,
  2. $\mathcal{V}$ runs in time at most $|x| - 1$,
  3. The average acceptance probability of $\mathcal{V}$ (as a QMA verifier) is greater than $1/2$ when averaged over $\overline{\mathcal{U}} \sim \mathcal{D}$, where $\mathcal{D}$ is the distribution of $\mathcal{U}$.

6.1 PRS/PRUs from $\mathcal{O}$

The PRU ensemble for a given length n is directly supplied by $\mathcal{U}_n$. The family $\{U_k\}_{k\in\{0,1\}^n}$ consists of the 2^n different Haar-random n-qubit unitaries in $\mathcal{U}_n$. This family has an efficient implementation relative to $\mathcal{U}$, as each U_k can be simulated with a single query to $\mathcal{U}_n$.

The security of PRUs relative to $\mathcal{O}$ relies on the fact that any polynomial-time adversary $\mathcal{A}$ with classical advice cannot distinguish a Haar-random unitary from the PRU ensemble with non-negligible advantage. This is proven using the BBBV theorem (optimality of Grover's algorithm) and the strong concentration properties of the Haar measure. The key idea is that the advantage of $\mathcal{A}$ is small with extremely high probability, even after union bounding over all choices of classical advice.

6.2 $\mathsf{BQP}^{\mathcal{O}} = \mathsf{QMA}^{\mathcal{O}}$

The intuition of the proof comes from the fact that the oracle $\mathcal{C}$ is constructed so that a $\mathsf{BQP}^{\mathcal{O}}$ machine can simulate the behavior of a $\mathsf{QMA}^{\mathcal{O}}$ verifier. Specifically, $\mathcal{C}$ allows a $\mathsf{BQP}^{\mathcal{O}}$ machine to approximate the maximum acceptance probability of a $\mathsf{QMA}^{\mathcal{O}}$ verifier $\mathcal{V}$ by:

1. Performing process tomography on $\mathcal{U}_n$ to learn a unitary transformation $\widetilde{\mathcal{U}}_n$ that is close to $\mathcal{U}_n$, for small n

2. Constructing the description of a new $\mathsf{QMA}^{\mathcal{C}}$ verifier $\mathcal{W}^{\mathcal{C}}$ that simulates $\mathcal{V}$ by replacing queries to $\mathcal{U}_n$ with $\widetilde{\mathcal{U}}_n$ or Haar-random unitaries,

3. Querying $\mathcal{C}$ with the description of $\mathcal{W}^{\mathcal{C}}$ to approximate the maximum acceptance probability of $\mathcal{W}^{\mathcal{C}}$.

Acknlowedgements

This lecture notes were prepared based on [5] and [6].

References

[1] Scott Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2063):3473–3482, 09 2005.

[2] Zvika Brakerski and Omri Shmueli. (pseudo) random quantum states with binary phase. In Dennis Hofheinz and Alon Rosen, editors, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 229–250. Springer, 2019.

[3] Hsin-Yuan (Robert) Huang, Richard Kueng, and John Preskill. Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10), 2020.

[4] Zhengfeng Ji, Yi-Kai Liu, and Fang Song. Pseudorandom quantum states. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 126–152. Springer, 2018.

[5] William Kretschmer. Quantum pseudorandomness and classical complexity. In Min-Hsiu Hsieh, editor, *16th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2021, Virtual Conference, July 5-8, 2021*, volume 197 of *LIPIcs*, pages 2:1–2:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[6] Henry Yuen. Lecture 9 - constructing pseudorandom states. `https://www.henryyuen.net/classes/spring2022/lec9-construction.pdf`.