CS59200-PRS: Pseudorandomness Nov 10th, 2025

Lecture 21: Random Phase state

Lecturer: Wei Zhan *Scribe: Xiuyu Ye*

1 Random Phase State

In this lecture we will give a construction of PRS using the following state ensemble:

Definition 1. *The n-qubit random phase state is defined as*

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

where $f\colon \{0,1\}^n \to \{0,1\}$ is a random function.

In fact, we will prove that the random phase states are approximately t-designs for exponentially large t, in the following sense.

Definition 2 (ε-approximate t-design)**.** *An ensemble $\{|\psi_k\rangle\}$ is an ε-approximate t-design if*

$$\frac{1}{2}\left\| \mathop{\mathbb{E}}_k \left[|\psi_k\rangle\langle\psi_k|^{\otimes t}\right] - \mathop{\mathbb{E}}_{|\psi\rangle\sim\mathrm{Haar}(N)} \left[|\psi\rangle\langle\psi|^{\otimes t}\right] \right\|_{\mathrm{Tr}} \le \varepsilon$$

where $\|\cdot\|_{\mathrm{Tr}}$ is the trace norm.

Definition 3 (Trace Norm)**.** *For a matrix M, let $\sigma_1 \ge \sigma_2 \ge \ldots \ge \sigma_r \ge 0$ where $r = rank(M)$ be the singular values of M. The trace norm of M is*

$$\|M\|_{\mathrm{Tr}} := \sum_{i=1}^{r} \sigma_i(M) = \max_{U,V\colon UU^\dagger=I, VV^\dagger=I} \left|\mathrm{Tr}(U^\dagger MV)\right|$$

where $\sigma_i(M)$ is the i-th singular value of matrix M.

We use trace norm in Definition 2 because it is the quantum analog of the total variation distance, and we will see later that it indeed characterizes how much two (mixed) states can be distinguished by any quantum algorithm. For now, let us just try two toy examples.

Example. For classical distributions p_1, p_2, consider the corresponding matrices $M_1 = \mathrm{Diag}(p_1)$ and $M_2 = \mathrm{Diag}(p_2)$. Then, singular values of $(M_1 - M_2)$ equal to the absolute values of diagonal entries of $(M_1 - M_2)$ and

$$\frac{1}{2}\|M_1 - M_2\|_{\mathrm{Tr}} = \frac{1}{2}|p_1 - p_2| = d_{\mathrm{TV}}(p_1, p_2).$$

Example. For pure states $|\psi_1\rangle$ and $|\psi_2\rangle$, consider their density matrices $M_1 = |\psi_1\rangle\langle\psi_1|$ and $M_2 = |\psi_2\rangle\langle\psi_2|$. Then we have

$$\frac{1}{2}\left\|M_1 - M_2\right\|_{\mathrm{Tr}} = \sqrt{1 - |\langle\psi_1|\psi_2\rangle|^2}.$$

This can be proved by examining the 2-dimensional subspace spanned by $|\psi_1\rangle$ and $|\psi_2\rangle$, which we skip here.

Theorem 1. *$\{|\psi_f\rangle\}$ is an ε-approximation t-design for $\varepsilon = O(t/\sqrt{N})$.*

Proof. Fix a function $f\colon \{0,1\}^n \to \{0,1\}$. Let $N = 2^n$. Consider

$$|\psi_f\rangle^{\otimes t} = \frac{1}{\sqrt{2^{nt}}} \sum_{x_1,\ldots,x_t \in [N]} (-1)^{f(x_1)+\ldots+f(x_t)} |x_1 x_2 \ldots x_t\rangle.$$

The corresponding density matrix is

$$|\psi_f\rangle\langle\psi_f|^{\otimes t} = \frac{1}{2^{nt}} \sum_{\substack{x_1,\ldots,x_t \in [N] \\ y_1,\ldots,y_t \in [N]}} (-1)^{f(x_1)+\ldots+f(x_t)+f(y_1)+\ldots+f(y_t)} |x_1 x_2 \ldots x_t\rangle\langle y_1 y_2 \ldots y_t|.$$

For a random function f, we know that

$$\mathop{\mathbb{E}}_f \left[(-1)^{f(x_1)+\ldots+f(x_t)+f(y_1)+\ldots+f(y_t)}\right] \neq 0$$

if and only if in the collection $(x_1, \ldots, x_t, y_1, \ldots, y_t)$, every element in $[N]$ appears even number of times. Therefore when $x_1, x_2, \ldots, x_t$ are distinct, the above expectation is non-zero only if $(y_1, \ldots, y_t)$ is a permutation of $(x_1, \ldots, x_t)$.

So we define the state

$$|\phi_f\rangle = \sqrt{\frac{1}{\binom{N}{t} \cdot t!}} \sum_{\text{Distinct } x_1,\ldots,x_t \in [N]} (-1)^{f(x_1)+\ldots+f(x_t)} |x_1 x_2 \ldots x_t\rangle,$$

which is the projection of $|\psi_f\rangle^{\otimes t}$ onto the subspace spanned by $|x_1 x_2 \ldots x_t\rangle$ with distinct $x_1, \ldots, x_t$. Since for independently random $x_1, \ldots, x_t \sim [N]$, the probability of them being not distinct is at most t^2/N by a union bound, we have

$$\begin{aligned}
&\langle\psi_f|^{\otimes t} |\phi_f\rangle \geq \sqrt{1 - \frac{t^2}{N}} \\
\implies &\frac{1}{2}\left\| |\psi_f\rangle\langle\psi_f|^{\otimes t} - |\phi_f\rangle\langle\phi_f| \right\|_{\mathrm{Tr}} \leq \frac{t}{\sqrt{N}} \\
\implies &\frac{1}{2}\left\| \mathop{\mathbb{E}}_f \left[|\psi_f\rangle\langle\psi_f|^{\otimes t}\right] - \mathop{\mathbb{E}}_f \left[|\phi_f\rangle\langle\phi_f|\right] \right\|_{\mathrm{Tr}} \leq \frac{t}{\sqrt{N}}.
\end{aligned}$$

To connect $|\phi_f\rangle$ with t copies of a Haar random state, we recall from the last lecture that

$$\mathop{\mathbb{E}}_{|\psi\rangle\sim\text{Haar}(N)}\left[|\psi\rangle\langle\psi|^{\otimes t}\right] = \frac{1}{\binom{N+t-1}{t}} \sum_{\text{multiset } X \text{ on } [N],\ |X|=t} |\text{sym}_X\rangle\langle\text{sym}_X|,$$

where

$$|\text{sym}_X\rangle = \sqrt{\frac{m_1!\dots m_N!}{t}} \cdot \sum_{\{i_1,\dots,i_t\}=X} |i_1 i_2 \dots i_t\rangle\langle i_1 i_2 \dots i_t|$$

and m_i is the multiplicity of i in X. From the reasoning above we know that in $\mathop{\mathbb{E}}_f[|\phi_f\rangle\langle\phi_f|]$ we are left with all the $|x_1\dots x_t\rangle\langle y_1\dots y_t|$ where $(y_1,\dots,y_t)$ is a permutation of $(x_1,\dots,x_t)$. Therefore when summing over all the permutations of $(x_1,\dots,x_t)$ and $(y_1,\dots,y_t)$ they combine exactly into $|\text{sym}_X\rangle\langle\text{sym}_X|$ for $X=\{x_1,\dots,x_t\}$. Therefore,

$$\mathop{\mathbb{E}}_f\left[|\phi_f\rangle\langle\phi_f|\right] = \frac{1}{\binom{n}{t}} \sum_{X\subseteq[N],|X|=t} |\text{sym}_X\rangle\langle\text{sym}_X|.$$

The two states differ only on the multisets that are not sets, which consists of at most t^2/N fraction of the multisets by the same union bound. Therefore,

$$\frac{1}{2}\left\|\mathop{\mathbb{E}}_f[|\phi_f\rangle\langle\phi_f|] - \mathop{\mathbb{E}}_{|\psi\rangle\sim\text{Haar}(N)}[|\psi\rangle\langle\psi|^{\otimes t}]\right\|_{\text{Tr}} \le \frac{t^2}{N}.$$

Finally, we can use the triangle inequality to conclude that

$$\frac{1}{2}\left\|\mathop{\mathbb{E}}_f[|\psi_f\rangle\langle\psi_f|^{\otimes t}] - \mathop{\mathbb{E}}_{|\psi\rangle\sim\text{Haar}(N)}[|\psi\rangle\langle\psi|^{\otimes t}]\right\|_{\text{Tr}} \le \frac{t}{\sqrt{N}} + \frac{t^2}{N} = O\left(\frac{t}{\sqrt{N}}\right). \qquad \square$$

Now we laim that the approximate design property directly implies the indistinguishability, through the following lemma.

Lemma 2. *For two states ρ_1 and ρ_2, and every quantum distinguisher A,*

$$|\mathbb{E}[A(\rho_1)] - \mathbb{E}[A(\rho_2)]| \le \frac{1}{2}\|\rho_1-\rho_2\|_{\text{Tr}}.$$

Proof. The quantum distinguisher can be decomposed into two parts: applying unitary operators, and a final measurement outputting 0 or 1. For the unitary part, the trace norm does not change:

$$\frac{1}{2}\left\|U\rho_1U^\dagger - U\rho_2U^\dagger\right\|_{\text{Tr}} = \frac{1}{2}\|\rho_1-\rho_2\|_{\text{Tr}}$$

since the singular values of $U(\rho_1-\rho_2)U^\dagger$ are the same as those of $\rho_1-\rho_2$.

For the measurement part, let Π be the orthogonal projection that gives outcome 1. Let $\Pi = U^\dagger DU$ be the diagonalization of Π, where $D = \mathrm{Diag}(d_1, \ldots, d_N)$ and each d_i is in $\{0, 1\}$. Then we have

$$\begin{aligned}
|\mathrm{Tr}[\Pi\rho_1] - \mathrm{Tr}[\Pi\rho_2]| &= \left|\mathrm{Tr}\left[U^\dagger DU\rho_1\right] - \mathrm{Tr}\left[U^\dagger DU\rho_2\right]\right| \\
&= \left|\mathrm{Tr}\left[DU\rho_1 U^\dagger\right] - \mathrm{Tr}\left[DU\rho_2 U^\dagger\right]\right| \\
&= \left|\sum_i d_i \left(\langle i|U\rho_1 U^\dagger|i\rangle - \langle i|U\rho_2 U^\dagger|i\rangle\right)\right| \\
&= d_{\mathrm{TV}}\left(\mathrm{diag}(U\rho_1 U^\dagger), \mathrm{diag}(U\rho_2 U^\dagger)\right).
\end{aligned}$$

The last equation is because the diagonal elements of a density matrix corresponds to a classical distribution. We can interpret the total variation distance using $\{\pm 1\}$ distinguishers instead, and thus over the diagonal matrices $V \in \{0, \pm 1\}^{N\times N}$ with diagonal entries $\{\pm 1\}$, we have

$$|\mathrm{Tr}[\Pi\rho_1] - \mathrm{Tr}[\Pi\rho_2]| = \frac{1}{2}\max_V \left|\mathrm{Tr}\left[VU\rho_1 U^\dagger\right] - \mathrm{Tr}\left[VU\rho_2 U^\dagger\right]\right| \leq \frac{1}{2}\left\|\rho_1 - \rho_2\right\|_{\mathrm{Tr}}$$

Since VU is also a unitary matrix. □

Corollary 3. *For every polynomial time quantum distinguisher A and every $t \leq \mathrm{poly}(n)$, we have*

$$\left|\mathop{\mathbb{E}}_f[A(|\psi_f\rangle\langle\psi_f|^{\otimes t})] - \mathop{\mathbb{E}}_{|\psi\rangle\sim\mathrm{Haar}(N)}[A(|\psi\rangle\langle\psi|^{\otimes t})]\right| \leq \mathrm{negl}(n).$$

Notice that this does not exactly mean that the random phase states $|\psi_f\rangle$ are PRS, as they cannot be efficiently prepared: The seed length for the random function $f : \{0, 1\}^n \to \{0, 1\}$ is already 2^n. However, we can instantiate the random phase states with a pseudorandom function, specifically some post-quantum secure PRF F. Since no polynomial time quantum algorithm can distinguish between f and F, it also cannot distinguish between $|\psi_f\rangle\langle\psi_f|^{\otimes t}$ and $|\psi_F\rangle\langle\psi_F|^{\otimes t}$ as the states are prepared by querying f or F on the uniform superposition $\frac{1}{\sqrt{2^n}}\sum_x |x\rangle$. Hence we get a PRS, assuming the existence of post-quantum PRF.