LECTURE 26 (April 24th)

TODAY Applications & Constructions of PRS

RECAP

State t-designs A distribution over n-qubit states is called a state t-design if the t-th moments match the t-th moment of the Haar measure, i.e.

$$\mathbb{E}_{|\psi\rangle \sim t\text{-design}} \left[|\psi\rangle\langle\psi|^{\otimes t} \right] = \mathbb{E}_{|\psi\rangle \sim \text{Haar}} \left[|\psi\rangle\langle\psi|^{\otimes t} \right]$$

Note that this means that no quantum algorithm can distinguish the two no matter how much time it takes when given only t-copies of the state. t is fixed here beforehand and this is an information-theoretic notion.

Pseudorandom states (PRS) A distribution over states is called a pseudorandom state distribution if $\exists$ a poly-time quantum algorithm that takes n-bit classical input k and outputs a state $|\psi_k\rangle$ s.t. no poly-time quantum distinguisher can distinguish any poly(n) copies of $|\psi_k\rangle$ from a Haar random state i.e. $\forall\, t = \text{poly}(n)$, and for all poly-time distinguishers A,

$$\left| \mathbb{P}_{k \in \{0,1\}^n} \left[A(|\psi_k\rangle) \text{ accepts} \right] - \mathbb{P}_{|\psi\rangle \sim \text{Haar}} \left[A(|\psi\rangle) \text{ accepts} \right] \right| \leq \text{negl}(n)$$

Since the algorithm does not know k, and distributions over quantum states is a mixed state, One can equivalently think of the above problem as distinguishing two mixed states

$$\rho^{(t)}_{PRS} = \mathbb{E}_k |\psi_k\rangle\langle\psi_k|^{\otimes t} \quad \text{and} \quad \rho^{(t)}_{Haar} = \mathbb{E}_{|\psi\rangle \sim \text{Haar}} |\psi\rangle\langle\psi|^{\otimes t}$$

Single-copy PRS It is trivial to construct a PRS if $t=1$ and $|\psi\rangle$ is on n-qubits where the key-length is n. This is because here we just want a mixed state that is indistinguishable from $\rho^{(1)}_{Haar} = \frac{\mathbb{I}_n}{2^n}$, i.e. the maximally mixed state.

Such a PRS can be constructed just by outputting a random computational basis states.

However, if we require that the PRS generator outputs a state on more qubits than the key length, then this becomes a non-trivial definition.

Applications of PRS

Before talking about how to construct PRSs and the assumptions needed for that, let us look at some applications of PRS. We will not talk much about state t-design applications.

Secret Key Encryption

Alice and Bob share a secret key k and Alice wants to send Bob a bit encoded in a quantum message that no poly-time adversary can crack but Bob can still decode it.

Here is a scheme that achieves this:

Let $U_k|0^n\rangle \longrightarrow |\psi_k\rangle$ be the poly-time unitary that prepares the PRS on key k. We will assume that k is the shared secret key.

Suppose Alice wants to send a bit $b \in \{0,1\}$ to Bob.
If $b=0$, she sends $|\psi_k\rangle$ and if $b=1$, she sends a Haar random state

To decode, Bob applies $U_k^\dagger$ to the message and if he gets $|0^n\rangle$ he says $b=0$ and otherwise $b=1$

One can show that if the number of qubits in the PRS is $> n+\omega(\log n)$ where n is the key length, then this scheme is secure. Note that this only relies on single copy security.

One can also easily extend this to send multiple bit messages [Exercise]

A related notion called bit-commitment can also be built from PRS but we will not cover it here since the part that relies on PRS is similar to the above

Pseudoentanglement

Recall that a Haar random state has the maximal amount of entanglement entropy, i.e. if $|\psi\rangle$ is a n-qubit state that is Haar random, then for any bipartition of the n-qubits into two parts (A,B), the entropy of $|\psi\rangle$ across this cut is $\sim \min\{|A|, |B|\}$ with high probability.

A distribution over (n-qubit) quantum states is called pseudoentangled if it is a PRS and the entanglement entropy across every cut is $O(\log^2 n)$. The $O(\log^2 n)$ is a parameter that can be tuned but it must be $\omega(\log n)$ since a PRS it is known that PRS must have $\omega(\log n)$-entanglement entropy. [We already mentioned that PRS have some entanglement and this is a more precise version of that]

Such pseudoentangled states have applications in quantum information theory, property testing and quantum gravity.

Asssumptions Needed for PRS

Remarks 1. We know how to construct state t-designs unconditionally, but for PRS we need some sort of assumption.

This is because if we have exponentially many copies of the state, one can learn the classical description of the state by a procedure called "state tomography". This can be done in PSPACE.

Thus, showing PRS exist unconditionally implies BQP $\neq$ PSPACE

What assumptions do we need to construct PRS?

a. One-way functions

These are functions such that it is easy to compute $f(x)$ but hard to compute $f^{-1}(y)$ for a random point in the image

Almost all classical cryptography can be based on one-way functions and vice-versa

We will see a construction of PRSs based on one-way functions that are secure against quantum adversaries

b. Weaker assumptions

There is some evidence that PRS can still exist in a world where one-way functions do not

In particular, in a joint work with Kretschmer, Qian & Tal, I showed that $\exists$ a classical oracle O such that $P^O = NP^O$ but PRS exist relative to O.

Remark The above is for single-copy PRS, for multi-copy, the proof is under a conjecture.

This means that quantum cryptography and other applications of PRS, that we will discuss later might still be possible even if classical cryptography based on one-way functions is not possible

2. We don't know how to stretch or shrink a PRS. This is because removing qubits does not give a pure state and a PRS always has some entanglement (This will be an exercise). This is in contrast to the classical setting.
In fact, there is some evidence in the form of black-box separations that shrinking a PRS is not possible in a black-box way.

PRS and state design constructions

We will introduce a construction that will give us an easy way to construct both state designs and PRS, under different assumptions.

The crux of the matter is the following statement, for which we first introduce the notion of trace distance.

Trace distance

Trace distance generalizes the notion of total variation distance between two distributions to the setting of density matrices

The Trace norm of a Hermititian matrix $A = \sum_i \lambda_i |v_i\rangle\langle v_i|$ is the quantity

$$\|A\|_1 = \sum_i |\lambda_i|$$

The trace distance between ρ & σ is $\frac{1}{2}\|\rho - \sigma\|_1$.

Operationally, the probability that any measurement (possibly inefficient) distinguishes ρ from σ is exactly

$$\frac{1}{2} + \frac{1}{2}\|\rho - \sigma\|_1$$

We claim the following:

Theorem Let $f: \{0,1\}^n \to \{0,1\}$ be a random boolean function. Then, the random state

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

is a $O\left(\frac{t^2}{2^n}\right)$-approximate t-design in trace distance.

i.e. $$\left\| \mathbb{E}\, |\psi_f\rangle\langle\psi_f|^{\otimes t} - \mathbb{E}_{|\psi\rangle \sim \text{Haar}} |\psi\rangle\langle\psi|^{\otimes t} \right\|_1 \lesssim \frac{t^2}{2^n}.$$

Note: $|\psi_f\rangle$ can be prepared by making one query to the phase oracle $O_f : |x\rangle \to (-1)^{f(x)}|x\rangle$

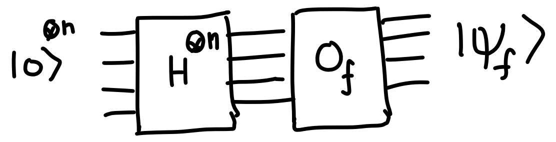

Moreover, $|\psi_f\rangle$ remains an approximate t-design even if $t \sim 2^{n/2}$, however so far it does not give efficient constructions of state designs or PRS, since O_f has exponential circuit complexity typically

Constructing state t-designs efficiently

For this we replace the random function $f : \{0,1\}^n \to \{0,1\}$ with a t-wise independent function.

What is a t-wise independent function ?

The truth table of f, i.e. $f(x_1), f(x_2), \ldots\ldots f(x_{2^n})$ is a t-wise independent bit-string

It is known how to construct these with $O(tn)$-size circuits. This gives us efficient t-designs for any fixed $t = \mathrm{poly}(n)$

Constructing PRS $\{|\psi_k\rangle\}_{k \in \{0,1\}^n}$

By definition, PRS must be efficiently computable, i.e given a key $k \in \{0,1\}^n$ there must be a poly-time quantum algorithm that generates the state $|\psi_k\rangle$ indexed by k.

To get a PRS, we need to make the following cryptographic assumption

$\boxed{\text{Existence of Pseudorandom Functions}}$
(quantum-secure PRFs)

A family of functions $\{f_k : \{0,1\}^n \to \{0,1\}\}_k$ is called a PRF if given k & x, $f_k(x)$ is efficiently computable and the output of is indistinguishable from a uniformly random function to all poly-time quantum adversaries, i.e.

$$\left| \mathbb{P}_{k \in \{0,1\}^n} \left[\mathrm{Adv}^{O_{f_k}} \text{accepts} \right] - \mathbb{P}_{f} \left[\mathrm{Adv}^{O_f} \text{accepts} \right] \right| \leq \mathrm{negl}(n)$$

This assumption is equivalent to assuming (quantum-secure) one-way functions exist.

To get a PRS, we just replace random function f with a pseudorandom function $\{f_k\}_k$

This gives us a family of states $\{|\psi_k\rangle\}_k$ that are efficiently preparable and form a PRS.

To prove the theorem we first introduce a useful concept :

Symmetric subspace Consider a quantum state on t registers each of them $d = 2^n$ dimensional. The symmetric subspace captures those states that are invariant under permuting the registers

$$\mathrm{Sym}_{d,t} = \{ |\psi\rangle \in (\mathbb{C}^d)^{\otimes t} \mid R_\sigma |\psi\rangle = |\psi\rangle \text{ for all } \sigma \in S_t \}$$

where $R_\sigma |x_1, \ldots x_t\rangle = |x_{\sigma(1)}, x_{\sigma(2)}, \ldots x_{\sigma(t)}\rangle$ is a permutation of the registers.

Example Let $t=2$, then $\frac{|12\rangle + |21\rangle}{\sqrt{2}} \in \mathrm{Sym}_{d,2}$

while $|12\rangle \notin \mathrm{Sym}_{d,2}$

If $t=1$, $\mathrm{Sym}_{d,1} = \mathbb{C}^d$

The symmetric subspace comes in the picture because of the following

Fact $\mathbb{E}_{|\psi\rangle \sim \mathrm{Haar}} |\psi\rangle\langle\psi|^{\otimes t} = \underbrace{\frac{\Pi_{\mathrm{Sym}_{d,t}}}{\dim(\Pi_{\mathrm{Sym}_{d,t}})}}_{:= \rho_{\mathrm{Sym}}}$ where $\Pi_{\mathrm{Sym}_{d,t}}$ is the projector on $\mathrm{Sym}_{d,t}$

This is the maximally mixed state on the symmetric subspace

The proof of this lemma follows from some basic representation theory which we won't cover here

Thus, our task boils down to showing

$$\mathbb{E}_f |\psi_f\rangle\langle\psi_f|^{\otimes t} \approx \rho_{\mathrm{Sym}}$$

In order to do this, we need an explicit basis for the symmetric subspace which we will introduce next time

NEXT TIME PRS analysis wrapup and Pseudorandom unitaries