LECTURE 25 (April 22nd)

TODAY Unitary Synthesis Lower Bound
Pseudorandom States and State Designs

RECAP Unitary Synthesis Problem

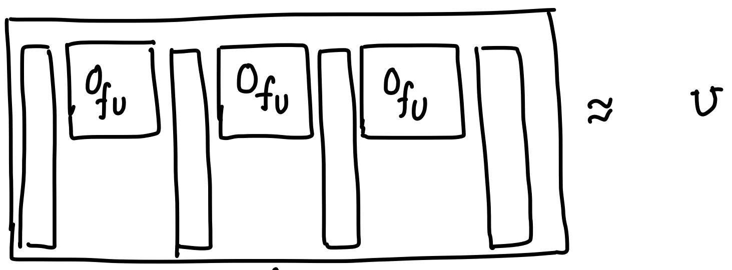

A

Is there a quantum algorithm A, a polynomial $p(n)$ and a encoding scheme that maps n-qubit unitaries U to a boolean function $f_U : \{0,1\}^{p(n)} \to \{0,1\}$ such that A makes poly(n) queries to f_U, uses poly(n) qubits of space and approximately implements U ?

Theorem No algorithm can synthesize a unitary with one-query and poly(n) ancillas.

Remark In contrast, state synthesis can be done with one-query and poly(n) ancillas, as shown in a recent work by Rosenthal

Last time we reduced it to the following problem about distinguishing two distributions on quantum states

Remark A distribution over pure quantum states is a mixed state, so one can also view the above as the problem of distinguishing two mixed states

First states $|\psi_1\rangle, \ldots |\psi_L\rangle \in (\mathbb{C}^2)^{\otimes n}$ are sampled and fixed. Here $L = 2^{n-1}$

Each state is sampled iid from the distribution $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} (-1)^{f(x)} |x\rangle$ where $f : \{0,1\}^n \to \{0,1\}$ is a uniformly random boolean function

Note the algorithm may depend on the initially sampled states

Consider the following two distributions on pure state

Distribution 1 Pick $k \in [L]$ at random and the input to the algorithm is the state $|\psi_k\rangle$. The algorithm does not know k so the corresponding mixed state is $\mathbb{E}_k |\psi_k\rangle\langle\psi_k|$

Distribution 2 The input to the algorithm is a random state sampled from the computational basis. In this case, the corresponding mixed state is the maximally mixed state $\frac{\mathbb{I}}{2^n}$

We sketched last time that if the algorithm can synthesize any unitary mapping

$$\text{span}\{|\psi_1\rangle, \ldots |\psi_L\rangle\} \text{ to } \{|1\rangle, \ldots |L\rangle\}$$

then it can distinguish the two distributions with probability $1/2$.

What does such an algorithm look like ?

The algorithm has access to an oracle f_U that might depend on $\{|\psi_1\rangle, \ldots |\psi_L\rangle\}$

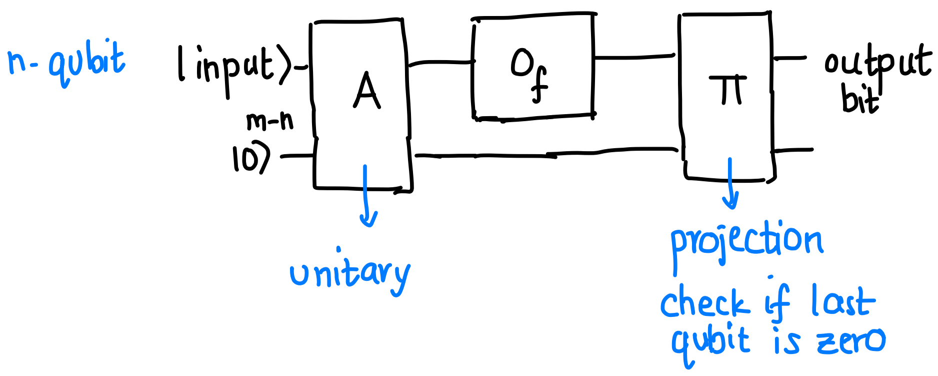

and $O_f \; |x\rangle \longrightarrow (-1)^{f(x)}|x\rangle$

where $f : \{0,1\}^{\text{poly}(n)} \longrightarrow \{0,1\}$ encodes some information about unitary U we want to synthesize

$p(|\psi\rangle, f) = \mathbb{P}[\text{ algorithm accepts on } |\psi\rangle] =$

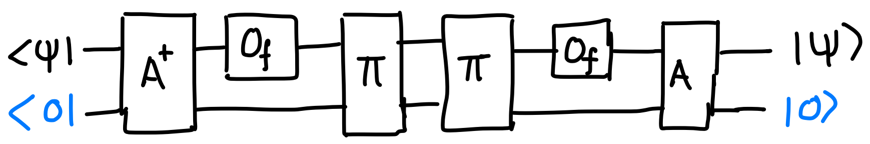

The following claim establishes that no such algorithm can distinguish Distribution 1 and 2 with $m = \text{poly}(n)$ ancillas and hence also cannot synthesize the above unitary

<u>Claim</u> w.h.p over $|\psi_1\rangle, \ldots |\psi_k\rangle$, $\max_{f:\{0,1\}^\ell \to \{0,1\}} \left| \mathbb{E}_k[p(|\psi_k\rangle, f)] - \mathbb{E}_h[p(|\psi_h\rangle, f)] \right| \lesssim \sqrt{\frac{m}{2^n}}$

We will sketch the proof of the claim in one special case.

Special case

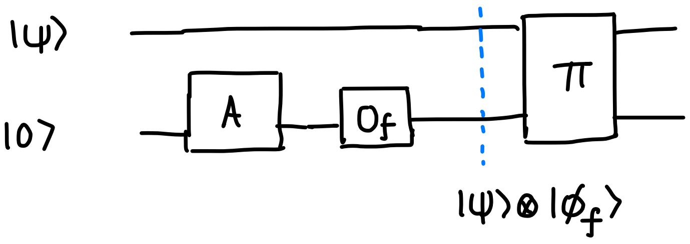

$\mathbb{P}[\text{accept}] =$

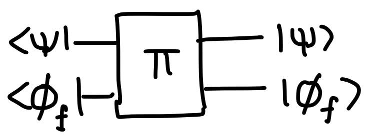

Want to show : $\mathbb{E}_{|\psi\rangle}$ [

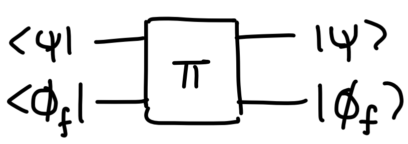

] is close under two distributions for all $f : \{\pm 1\}^\ell \to \{0,1\}$

$= \langle\phi_f| \, \mathbb{E}_\psi[M_\psi] \, |\phi_f\rangle$ where $M_\psi =$

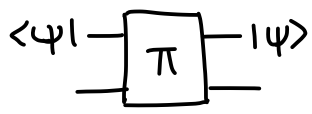

unit vector

Suffices to bound $\max_f \left| \langle\phi_f| \, \mathbb{E}_k[M_{\psi_k}] - \mathbb{E}_h[M_{\psi_h}] \, |\phi_f\rangle \right|$ w.h.p.

$\leq \left\| \mathbb{E}_k[M_{\psi_k} - \boxed{\mathbb{E}_h[M_{\psi_h}]}] \right\|_{op}$ w.h.p.

random matrix (depends on k & also $|\psi_1\rangle, \ldots |\psi_k\rangle$)

fixed matrix

Letting $B_k = M_{\psi_k} - \mathbb{E}_h[M_{\psi_h}] \quad \forall k \in [L]$

we want to bound $\| \mathbb{E}_k[B_k] \|_{op} = \left\| \frac{B_1 + \cdots + B_L}{L} \right\|_{op}$

$2^m \times 2^m$

We claim the following without proof: $B_1, \ldots B_L$ are iid random matrices with zero mean and operator norm at most 2

Matrix Chernoff Bound says that w.h.p.

$$\| B_1 + B_2 + \dots + B_L \|_{op} \lesssim \sqrt{L} \cdot \sqrt{\log \dim}$$

$$\Rightarrow \| \mathbb{E}_k [B_k] \|_{op} \leq \frac{\sqrt{\log \dim}}{\sqrt{L}} = \frac{\sqrt{m}}{\sqrt{2^{n-1}}} = \frac{\sqrt{\mathrm{poly}(n)}}{\sqrt{2^{n-1}}}$$

Thus, no algorithm of this form can distinguish the two distributions over states

The general case can essentially be reduced to the above with one small trick that we will not discuss

Pseudorandom States and Designs

We saw that the problem above reduced to distinguishing two distributions of quantum states This motivates the definition of pseudorandom states and state t-designs.

Informally, pseudorandom states and state t-designs are distributions over quantum states that can not be distinguished from a Haar random state in related but distinct ways.

First, let us revisit the classical analogoues of these objects.

Suppose we have a distribution on n bits $X_1, \dots, X_n$ which are uniform and independent One way to relax the notion of indepencence is t-wise independence

t-wise independent distribution: A distribution (on bits) $X_1 \dots X_n$ is called t-wise independent if $\forall$ every subset $S \subseteq [n]$ of size $\leq t$, the bits in S are independent, i.e, all $\leq t$-wise moments match the uniform distribution.

E.g. Let X_1, X_2 be uniform and independent random bits Then $X_1, X_2, X_3 = X_1 \oplus X_2$ is a 2-wise independent distribution.

Outputs of Pseudorandom generators (PRGs): A pseudorandom generator is a (poly-time computable) function that takes n uniform random bits and output $m > n$ random bits that look like m uniform bits to any poly-time distinguisher

Formally, $g : \{0,1\}^n \longrightarrow \{0,1\}^m$ s.t. $\forall$ any polynomial time algorithm A, we have

$$\left| \mathbb{P}_{k \in \{0,1\}^n} \left[A(g(k)) \text{ accepts} \right] - \mathbb{P}_{z \in \{0,1\}^m} \left[A(z) \text{ accepts} \right] \right| \leq \mathrm{negl}(n)$$

↳ decreases faster than any inverse polynomial

$m - n$ is called the "strectch" of the PRG.

t-wise independence is a information theoretic notion — as long as the algorithm only looks at t bits, it can't distinguish it from the uniform distribution no matter how long it takes, but t is fixed before

In contrast, in a PRG, the distinguisher can look at any poly(n) bits and it can decide how many bits to look at in a adaptive fashion, but the security is only against computationally bounded distinguishers since otherwise an exponential time distinguisher can "break" the PRG

There are several constructions of t-wise independent distributions but for PRG we actually do not know if they exist. If we could show this unconditionally, then $P \neq NP$. The best we can do is to show that PRGs exist under some cryptographic assumption such as one-way or pseudorandom functions

Remark The stretch of a PRG can be amplified even from 1 to any poly(n)

Let us now discuss the quantum analogs of these objects

First, let us remind us of the Haar measure on the sphere.

Haar measure (informal) A Haar random state $|\psi\rangle$ on n-qubits is a "uniformly random" vector on the 2^n-dimensional complex unit sphere.

The t^{th} moment of the Haar measure is the quantity

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

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,

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

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

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

<u>NEXT TIME</u> Applications & Constructions of PRS