Notes for Lecture 5

Today we'll move on from the slightly contrived applications of quantum algorithms we discussed previously, and we'll talk about a quantum algorithm that has serious implications for cryptography.

1 Grover Search

1.1 The Setting

Grover's algorithm allows us to solve the following promise problem (which we'll show how to extend by the end of the lecture). We have a function $f : \{0,1\}^n \to \{0,1\}$, and we know there is exactly one $s \in \{0,1\}^n$ such that f(s) = 1.

We're going to imagine that we're living in an oracle world where someone has implemented f and we can make black box access to it. Or we can also think of f as being implemented by some really complicated circuit that we don't want to open up.

How do we solve this? Classically, this problem basically takes 2^n time. We simply brute force search over all inputs. In comparison, we can use a quantum algorithm to do this in $\Theta(2^{n/2})$ time. It's important to note that this division by 2 is happening in the exponent, so it's a quadratic saving.

1.2 Motivation

Why do we care about this in cryptography?

A one way function $g : \{0,1\}^n \to \{0,1\}^n$ is a function that (informally) is easy to evaluate and hard to invert. I can take an x and compute g(x), but given g(x) I cannot (computationally) find any pre-image of g(x) (including x of course). Later on in this course we'll look at more complicated cryptographic objects, and we'll see that a lots of them really have one way functions hiding in them. All this is to say that one way functions are one of the most fundamental objects in cryptography, so understanding the hardness of breaking them is fundamental.

In practice, n is chosen to thwart known attacks. So one known attack that works for any g is to the brute force attack, where I go over all inputs x until I find a pre-image that matches the g(x) I was given. The time complexity of this is 2^n . What we do then is set n to be so large as to make this intractable, which in practice means setting n to be something like 128. For some one way functions, these brute force attacks are the best attacks that we know.

This feels extremely similar to the setup for the Grover search, but there's a slight mismatch. So for now let's assume that the one way function g is actually injective (equivalently, one-to-one), which means if I'm given y, my goal is to find the unique x such that g(x) = y. Now I define f(x) to be the function that is 1 if g(x) = y, and 0 else. Then injectivity of g means that f(x) is only 1 at a single point x, so we can now invoke Grover search.

The implication of this is that with a Grover search, I can perform inversion on (injective) one way functions in 2^{64} time. So if we're aggressive with one-way function parameter settings and we actually have n = 128, this can be broken in 2^{64} time by a quantum computer (meaning we lose security against brute force attacks). At the end of this lecture we'll remove this injectivity requirement.

1.3 Grover's Algorithm: Description 1

Step 1. I prepare the uniform superposition over all inputs

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle$$

If we recall from last time, we can easily prepare this state by starting with an n bit string that is all 0's and applying Hadamard to it: $(H|0\rangle)^{\otimes n}$.

Step 2. I alternate between doing two things.

- I apply f using a phase kickback. The map is $|x\rangle \rightarrow (-1)^{f(x)}|x\rangle$ from applying a unitary U_f (essentially the same idea we used for Deutsch's algorithm last time)
- "Grover iteration." Here we map |x⟩ → (²/_{2ⁿ} ∑_y |y⟩) |x⟩. It turns out that this is a unitary transformation, though it might not appear unitary at first glance. Step 3. Measure (and hope that I get the right answer)

How many times do we step 2? It turns out $2^{n/2}$ times is the right amount.

1.4 Grover's Algorithm: Description 2

Now we give a second description of Grover's algorithm. The point of the first description is that it's clear that we can implement it given f. The point of this second description is that it's actually equivalent to the first description, but now we write it in such a way that illuminate why each step is unitary. Step 1. Apply unitary $U_s = I - 2|s\rangle\langle s|$. This $|s\rangle\langle s|$ corresponds to a 1 on the diagonal entry corresponding to the secret input s such that f(s) = 1 and 0 everywhere else.

We can see that this is the same as step 1 in description 1, since in that description the unitary is really the identity everywhere except at x, where it is 1 - 2 = -1.

Step 2. We apply $U_{\psi} = 2|\psi\rangle\langle\psi| - I$ where $|\psi\rangle = \frac{1}{\sqrt{2^n}}\sum_x |x\rangle$. Now we can see how this is the same operation as in Step 2 in description 1. In the first term the $\frac{1}{\sqrt{2^n}}$ terms multiply and give the $\frac{2}{2^n}$ term in the description 1. Then the identity I applied to $|x\rangle$ just gives $|x\rangle$.

This is unitary since we're just substituting in ψ for s in the unitary in step 1, and then flipping the sing.

1.4.1 Intuition for Step 2

If I imagine I'm in the Fourier domain, the uniform $|\psi\rangle$ in the primal corresponds to 0 in the Fourier domain. Then flipping $|\psi\rangle$ in the primal means I'm just flipping 0 in the Fourier domain, which is what happens when you apply U_0 (taking $s = |0\rangle$ in the U_s definition in step 1).

So this mean step 2 is just $-H^{\otimes n}U_0H^{\otimes n}$. The rightmost Hadamard takes me to the Fourier domain. The U_0 is just flipping the 0, and then we apply another Hadamard to get back. The minus sign is just because in step 2 we're flipping the sign. (Note that the U_{ψ} in step 2 should really be $-U_{\psi}$, but we won't worry about the minus signs)

1.5 Picture Intuition for Grover iterations



We're looking at graphs where the horizontal axis is over the all different states, and

the y axis is the amplitude on each state.

At the very start of the Grover iterations, we're in the top left graph where the amplitude on each state is $\frac{1}{\sqrt{2^n}}$.

- (First stage of first Grover iteration) Now when I apply f, this flips the sign on the secret input. This gives the second graph on the left. Now the secret input has an amplitude of $-\frac{1}{\sqrt{2^n}}$, and everything else is left unchanged.
- (Second stage of first Grover iteration) In the second stage of a Grover iteration, recall that we flip about 0 in the Fourier domain. Intuitively, 0 corresponds to the mean (but not exactly). To see this intuition, recall the form of the Fourier transform:

$$\sum_{x} \alpha_{x} |x\rangle \to \frac{1}{\sqrt{2^{n}}} \sum_{y} |y\rangle (\sum_{x} (-1)^{x \cdot y} \alpha_{x}).$$

What is the amplitude on $|0\rangle$? We simply plug in y = 0 and get $\frac{1}{\sqrt{2^n}} \sum_x \alpha_x$, which intuitively corresponds to the mean). So we can think of what happens in the primal as a "flip about the mean". What happens is we end up at the third graph on the left side of the picture. All of the amplitudes at states not corresponding to the secret input were very slightly above the mean, and they get flipped to very slightly below the mean. On the other hand, the amplitude of the secret input gets flipped all the way up to approximately $\frac{3}{\sqrt{2^n}}$.

- (First stage of second Grover iteration). Now I flip the secret point across the mean again, and it goes to approximately $\frac{-3}{\sqrt{2^n}}$.
- (Second stage of second Grover iteration). The mean is still basically $\frac{1}{\sqrt{2^n}}$. The result is that I'm at roughly the same picture as before, except the secret point is at approximately $\frac{5}{\sqrt{2^n}}$.

After t steps, the amplitude is approximately $\frac{2t-1}{\sqrt{2^n}}$. Once this amplitude is some constant fraction, when we measure we'll get the correct answer with some reasonable probability. Then we can repeat a few times to boost correctness.

What happens if we keep going past this point? At some point, these errors are going to get too big. Also keep in mind that everything is reversible, and it turns out that if we keep doing the iterations we'll get back to where we started.

1.6 A More Careful Proof

Define $|\psi'\rangle = \frac{1}{\sqrt{2^n-1}} \sum_{x \neq s} |x\rangle$, and write the starting state as

$$|\psi\rangle = \sqrt{\frac{2^n - 1}{2^n}} |\psi'\rangle + \frac{1}{\sqrt{2^n}} |s\rangle.$$

When I apply a full Grover iteration, what happens to $|s\rangle$ is

$$|s\rangle \to U_{\psi}U_{s}|s\rangle = U_{\psi}(-|s\rangle) = (2|\psi\rangle\langle\psi| - I)(-|s\rangle) = -2|\psi\rangle\langle\psi||s\rangle + |s\rangle$$

What happens to $|\psi'\rangle$ is that U_s doesn't do anything to ψ' , so it's just

$$|\psi'\rangle \to U_{\psi}U_s|\psi'\rangle = U_{\psi}|\psi'\rangle = (2|\psi\rangle\langle\psi|-I)|\psi'\rangle = 2|\psi\rangle\langle\psi|\psi'\rangle - |\psi'\rangle$$

We note that both of these terms on the R.H.S are in the span of $(|\psi'\rangle, |s\rangle)$. So every time I apply a Grover iteration, I stay in the plane spanned by $|\psi'\rangle$ and $|s\rangle$. Also it's worth noting that $|\psi'\rangle$ and $|s\rangle$ are orthogonal, since the only place $|s\rangle$ has any weight is where $|\psi'\rangle$ has 0 weight.

I can go to another picture where I have $|s\rangle$ and $|\psi\rangle$ and $|\psi'\rangle$. Also note that we're living in the real plane. $|\psi\rangle$ and $|\psi'\rangle$ have some small angle $\theta/2$ (we'll work out the angle in just a moment).



Now when I apply U_s , what happens is I reflect about $|\psi'\rangle$. So U_s is flipping $|\psi\rangle$ to $U_s|\psi\rangle$ that has an angle $\theta/2$ below $|\psi'\rangle$.

What happens when I apply U_{ψ} ? Let's forget the overall minus sign, since we've said those don't really matter. Then I can think of this as reflecting about $|s\rangle$, and I end up at $U_{\psi}U_s|\psi\rangle$ (which is on the lefthand side, but when we account for the minus sign we're back on the righthand side).

So I keep doing this process of reflecting about ψ' and then reflecting about s (but with a minus sign), which produces the angles above.

One way to view this is to note that it's similar to what was going on before, except now the increase of θ in each step is exact.

Let $|\psi_t\rangle$ be the state after t steps. Then let θ_t be the angle from $|\psi'\rangle$. Then $\theta_t = (\frac{1}{2} + t)\theta$. Then we measure, and the probability of getting state s is $\sin^2((\frac{1}{2} + t)\theta)$. So the probability I don't get s is going to be $\cos^2((\frac{1}{2} + t)\theta)$.

So what is θ ? We know that the original vector, my probability of getting $|s\rangle$ was $\frac{1}{2^n}$ so its amplitude in the direction of s is $\frac{1}{\sqrt{2^n}}$. So we have $\sin(\theta/2) = \frac{1}{\sqrt{2^n}}$.

If we use our small θ approximations for sin, then we have $\theta \approx \frac{2}{\sqrt{2^n}}$. In order to get a high success probability I want $\sin^2((\frac{1}{2} + t)\theta)$ to be something like 1, so we need the input to sin to be around $\pi/2$. If we solve

$$(\frac{1}{2}+t)\frac{2}{\sqrt{2^n}} = \frac{\pi}{2}$$

Then I get $t = \lfloor (\pi/4)\sqrt{2^n} \rfloor$. It's important that we actually do stop here.

1.7 Many Accepting Inputs

So this all works because there was only one accepting input. But what if there are more accepting inputs? Let S be the set of r accepting inputs.

Then I define ϕ to be the superposition to be all accepting inputs, which generalizes the $|s\rangle$ from before. Basically what we can do is write $|\psi\rangle = \alpha |\psi'\rangle + \beta |\phi\rangle$. (for some α and β that we won't work out).

If we run the same analysis as before, we'll conclude that the results of both steps of the Grover iteration is some state that is in the span of $|\psi'\rangle$, $|\phi\rangle$. So just as before, we can zoom in on this plane.

The resulting picture looks identical, and the only point is that θ has changed. ϕ replaces s in the picture, where ϕ is the state where if I measure, I get one of the r accepting inputs. Now $\sin(\theta/2) = \sqrt{r/2^n}$. This means that as long as r is relatively small, θ is about $2\sqrt{r/2^n}$, so I need $|\pi/4\sqrt{2^n/r}|$ steps.

If I have more accepting inputs. We can just try all the powers of 2.