Notes for Lecture 1

1 Administrative

- MW 1:30 - 2:50, Friend Center 110.
- Course website: http://www.cs.princeton.edu/~mzhandry/2018-Fall-COS597A/
- Office Hours by appointment. Email mzhandry@princeton.edu to schedule.
- Grading will be based on homework every two to three weeks, plus a final project. Three to four homework assignments throughout the term, two weeks to do them.
- Everyone in the course should scribe 2 lectures. Homework 0 is to sign up for the scribing calendar.
- Prerequisites: nothing firm. In particular no physics is required. We'll only talk about physics today, and soon we'll transition to a mathematical model that on its surface has nothing to do with physics. One thing that'll be useful is familiarity with basic notions of complexity (P, NP, what a reduction is), but this is not required. Basic crypto knowledge is also a plus (knowing what an encryption is, what semantic security is, etc.)

2 Introduction

In traditional crypto, we have a setting where Alice wants to send a message to Bob, but she suspects someone is eavesdropping on the communication. If Alice and Bob share a secret key k, she can send a message $\mathsf{Enc}(k, m)$ to Bob and Bob runs some decrypt algorithm to recover $m \leftarrow \mathsf{Dec}(k, m)$. The idea is that the eavesdropper, who only sees $\mathsf{Enc}(k, m)$ and doesn't know the secret key k can't learn anything about m.

However, even this basic picture makes a lot of assumptions:

- Alice, Bob and Eve all obey Newtonian Physics.
- The channel obeys Newtonian physics and even the message obeys Newtonian physics (we'll get to what that means).

In the early 1900s, physicists realized that Newtonian physics was incomplete and didn't adequately describe the world around us. For example, when things are really big, Newtonian physics breaks down since the speed of light is constant, and we get relativity. When things are really small, we get quantum mechanics. Because we're not really big or really small, it didn't seem like either of these mattered for cryptography. But now it looks like we're getting to the point where quantum computers will allow us to do things that we weren't able to before.

2.1 Course Overview: 12 Week Plan

Today, we'll talk about the physical motivation of where quantum comes from. Then we'll move on to a mathematical model of quantum (2 weeks). After this we'll start looking at quantum attacks on classical crypto and defenses (4 weeks), an area referred to as *post-quantum crypto* (after someone has a quantum computer). After fall break, we'll see some quantum protocols (5 weeks). In this situation, we'll look at what happens when our smartphones are quantum computers. Maybe when we all have quantum computers we can obtain more secure protocols. We'll also look at the case where the channel we communicate over is "quantum", or when the messages themselves are "quantum". The last week of the class will be student presentations.

3 Quantum Physics

Quantum physics came out of the question: Is light a wave or particle?

In Young's double slit experiment, we have a light source on the left and a wall on the right. The first thing we do is cut a tiny hole in the wall, small enough to basically be a point. To the right of the wall we have a detector. What we see if we plot intensity vs. distance from the center is a horizontal line (the intensity of the light doesn't depend on the distance from center). If I cut a second slit in the wall, you would expect that you just get a graph with a more intense light (higher horizontal line).

But instead what you get is something that looks like a cosine curve (shifted up). This is what you would expect from wave behavior. If we look at different points on the detector, what matters is the distance from these points to the two slits. In the center, the distances are equal, and if they're both at peaks then at this point the intensity will be very high. If one is half a wavelength longer, then a peak will cancel out with a trough, and we get an intensity level in the middle.

From this experiment you would conclude that light is a wave.

In the early 1900s, we had an "ultraviolet catastrophe". Under the conventional theories of thermodynamics, we could have a black body radiating an infinite amount of energy. To resolve this, Planck posited that light was only absorbed in discrete increments called “quanta”.

Another thing we discovered in the early 1900s was the photoelectric effect. We could shine light on a material, excite electrons, and cause them to fly out. The curious thing was that the energy of the electrons depended on the frequency (color) of the light, and crucially not the intensity. This didn’t seem to make sense with the wave theory. The resolution was due to Einstein, who posited that not only is light absorbed in discrete increments, but that light itself is discrete (“photons”).

The natural thing to wonder is: if light is a particle, what’s going on in the double slit experiment? Intuitively, opening up another slit should only increase the intensity of light at any location, since we’re only increasing the chance a photon will hit any spot.

So we have experiments that conclusively demonstrate light is a wave, and other experiments that conclusively demonstrate light is a particle.

The resolution is strange: the photon is going through both slits at the same time, and can interfere with itself to construct the wave pattern we see.

3.1 A Mathematical Model (Quantum Mechanics)

We’ll define a mathematical model to explain this physics, and after that we’re going to drop all the physics. Some notation

- $|0\rangle$ (ket notation) represents a photon in the state of having gone through the top slit.
- $|1\rangle$ will be a photon having gone through the bottom slit.

If we were in classical physics, we’d be done because it would either go through the top or the bottom slit. But as we said, the photon can go through both, so its general state is something of the form

$$\alpha|0\rangle + \beta|1\rangle$$

where α and β are *amplitudes*.

My normalization condition will be that

$$|\alpha|^2 + |\beta|^2 = 1.$$

Why this? The rough physical intuition is that energy is proportional to the square of the amplitude, and the total energy is what we want to take to be 1. Some remarks:

- α, β can be negative. This actually makes sense when we look at this wave version. I can put some material at a slit to slow light down over a small region, shifting where the peaks and troughs are of the wave coming out of one slit. If one slit becomes offset from the other slit by half a wavelength, we'd represent this by saying that β is the negative of α.

- α, β can also be complex numbers, i.e. what do I do if one wave is the other shifted by a quarter of a wavelength? If the absolute value of these α, β represent the amplitudes, I need an extra dimension to represent phase shift. It turns out going to complex numbers give us this extra dimension to represent phase shifts. Recall that complex numbers have a real and imaginary component, so we can write a complex α as $\alpha_R + i\alpha_c$. The norm (squared) of α is then $|\alpha|^2 = |\alpha_r|^2 + |\alpha_c|^2$. We'll write the conjugate of α as $\alpha^* = \alpha_r - i\alpha_c$. Also it will be useful to have the fact $e^{i\theta} = \cos\theta + i\sin\theta$. $e^{i\theta}$ is the point on the unit circle at an angle of θ from 3 o'clock (counterclockwise).

This θ in the $e^{i\theta}$ will correspond to the phase shift of our waves in the double slit experiment. If I have the waves in the same phase (and equal intensity), we write this as $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$. If I want to slow down the wave going through the top slit by θ, this is the state $e^{i\theta}\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$. If I want to slow it down by half a wavelength, I set $\theta = \pi$. If I set $\theta = 2\pi$, I'm back to $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$.

3.2 Measurements and Operations

Say I measure which slit the photon went through. What happens?

If I measure it and I find out that it went through the top slit, then magically it no longer went through the bottom slit. Indeed, the interference pattern will disappear.

The fact that a measurement perturbs the system and collapses the state to $|b\rangle$ is one of the central tenants of quantum mechanics (another central tenant is that before I measure, the state is a superposition of $|0\rangle$ and $|1\rangle$).

What operations can I do on quantum states? I can put different materials on the different slits to change the phase on either $|0\rangle$ or $|1\rangle$ (multiplying either by $e^{i\theta}$ as described earlier). I can also do a bit flip that sends $|b\rangle \to |1-b\rangle$.

But phase shifts and bit flips aren't the only things we can do. A Hadamard transformation does the following:

$$\begin{aligned} |0\rangle &\to \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \\ |1\rangle &\to \frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle. \end{aligned}$$

We won't totally get into the physics of this, but here's a rough intuition.

Let's say I add a second wall to the right of the initial two-slit wall, and place two slits in this second wall as well. We place the top slit exactly in the middle of where the two initial slits are. Then we place the second slit some distance below it so the following properties hold.

Now if the photon goes through both slits, its phase is $\frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \to |0\rangle$ (Hadamard transform is applied linearly). This corresponds to the fact that if the photon goes through both slits, we'll have a peak at the top slit and a trough at the bottom slit. Label the top slit as $|0\rangle$, and the photon will always go through the top slit. If we change the phase to $\frac{1}{\sqrt{2}}|0\rangle - \frac{1}{\sqrt{2}}|1\rangle$, the Hadamard transformation takes it to $|1\rangle$. This happens because we did a phase shift on one wave, and now the peak is at $|1\rangle$ and the trough is at $|0\rangle$.

Generally, we can do any unitary transform for some matrix $U = \begin{pmatrix} U_{00} & U_{10} \\ U_{10} & U_{11} \end{pmatrix}$

$$\begin{aligned} |0\rangle &\to U_{00}|0\rangle + U_{01}|1\rangle \\ |1\rangle &\to U_{10}|0\rangle + U_{11}|1\rangle \end{aligned}$$

Why unitary? It turns out that to preserve normalization, it is sufficient for $U^\dagger U = I$ where

$$U^\dagger = \begin{pmatrix} U_{00}^* & U_{01}^* \\ U_{10}^* & U_{11}^* \end{pmatrix}$$

There are many physical realizations of this mathematical framework, so at this point we'll stop thinking about the physics.

3.3 Classical vs. Quantum

In the classical setting, we can have a weighted coin flip where heads happens with probability a and tails happens with probability b. Our normalization condition is that the probabilities add to 1: $a + b = 1, a > 0, b > 0$. What are operations? We can do things like "if the coin lands on heads, flip again" or "if the coin lands on tails, flip a coin with a different bias". I can also write the transformation as

$$S \cdot \begin{pmatrix} a \\ b \end{pmatrix}$$

where S is a stochastic matrix, meaning that its columns sum to 1 and are positive.

In the quantum setting, we have $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ where $|\alpha|^2 + |\beta|^2 = 1$. Our operations are unitary operations

$$U \cdot \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$$

where U is a unitary matrix.

Why are things different? Classically, probabilities that are positive will always add. In other words, introducing additional routes that end up at the same outcome will always increase the probability of getting that outcome. In the quantum setting, multiple paths to the same final answer can cause cancelations because amplitudes can be negative.

A final observation is that in the classical setting, operations must preserve L_1 norm. In the quantum setting, operations preserve L_2 norm. So why don't we look at L_3 norm? We'll discuss briefly next time why this probably isn't the right answer.