CMSC 657: Introduction to Quantum Information Processing
Lecture 28

Instructor: Daniel Gottesman

Fall 2024

1 Quantum Key Distribution

1.1 Cryptography and the One-Time Pad

We return now to cryptography. As we discussed earlier in the class, current cryptographic systems mostly use computational security, and for some systems, such as RSA, the underlying computational assumption would be violated by a quantum computer, making it possible to break those cryptographic systems. However, recall that the one-time pad is information-theoretically secure, meaning it cannot possibly be broken regardless of Eve's computational power.

- **Private-Key System:** Alice and Bob share a random bit string k, unknown to Eve. k is n bits long.
- **Ciphertext:** Given message m, the ciphertext is $e = m \oplus k$, the bitwise XOR of k and m.
- **Decryption:** $e \mapsto e \oplus k = (m \oplus k) \oplus k = m$.

The security of the one-time pad does not depend on her computational power or any other assumption except: Alice and Bob have private domains which Eve cannot see into; and the shared private key must be completely random, as long as the message, and secret from Eve. This last requirement means that the private key cannot be reused.

This last set of requirements makes setting up the key for a one-time pad rather difficult. It requires a trusted courier or prior physical meeting to set up the initial secret key, and the key will run out and need to be renewed if you send a lot of messages or long messages. Computationally secure systems let you use the same key for a longer period of time. The advantage of a public key system is that it lets you talk securely with someone you have not met before.

Quantum key distribution (QKD) uses quantum mechanics to achieve similar properties. It lets Alice and Bob generate a secure private key over an insecure quantum channel, plus an authenticated public classical channel. The key can then be used with the one-time pad to send messages with information-theoretic security. It can be used to generate more key if Alice and Bob have been talking for a long time and used up their old key, and it can generate new key between Alice and Bob that have not previously met, provided there is a way for them to be sure that they are talking to the right people.

The set up we will consider for QKD is that Alice and Bob share an insecure quantum channel from Alice to Bob (or sometimes both ways). Eve can do whatever she likes with it, subject to the laws of quantum mechanics. She can measure it, she can take the qubits out of the channel and replace them with something else, anything she wants. Alice and Bob also have classical channels in both directions; these channels can be read by Eve but not changed (that is, the channels are *authenticated*).

1.2 BB84 Protocol

The BB84 protocol is named after Bennett and Brassard, the inventors.

1. Alice generates N random bit pairs (a_i, r_i).

2. Each pair specifies a one-qubit state. Alice creates the states and sends them to Bob over the quantum channel.

   - a_i specifies the basis. 0 is the Z basis, 1 is the X basis. r_i specifies which bit in the basis we use. Thus:

a_i	r_i	state
0	0	$\|0\rangle$
0	1	$\|1\rangle$
1	0	$\|+\rangle$
1	1	$\|-\rangle$

3. Bob chooses N random bases b_i and measures qubit i in basis b_i, getting result s_i.

4. Alice and Bob announce a_i and b_i over the classical channel. They do *not* announce r_i and s_i. If $a_i = b_i$, they keep bit i. If $a_i \neq b_i$, they discard bit i. The bits r_i and s_i that they keep are the *raw key.*

   - At this point, if everything is perfect, Alice and Bob should have $r_i = s_i$ for all the bits they keep. However, in a realistic system, there will be errors, so they don't expect to have r_i and s_i agree completely.

     More importantly, Eve could be eavesdropping, and eavesdropping introduces errors. For instance, suppose Alice sends $|0\rangle$ for qubit i. Eve does not yet know the basis, so perhaps she chooses to measure this bit in the X basis instead. She gets a result and creates the corresponding qubit and sends it to Bob. If Bob chooses $b_i = 1$, the bit is discarded anyway, but if Bob chooses $b_i = 0$, Alice and Bob will keep bit i. However, Bob's measurement result in the Z basis is random and independent of Alice's bit $r_i = 0$. Thus, this bit has a 50% chance of error.

     Indeed, one can show that any attempt at eavesdropping by Eve has a chance to cause errors. The more information Eve learns, the more disturbance she causes in the state.

5. Alice and Bob choose a random subset, consisting of say 50% of the remaining bits, and announce r_i and s_i for bits in the subset. They compare and calculate the error rate. If the error rate is above some value determined from the security proof, they abort the protocol and discard the key. Otherwise they continue to the next step.

   - QKD lets Alice and Bob *detect* the presence of an eavesdropper. If Eve chooses not to learn very much, Alice and Bob may not detect her, but she also doesn't break the protocol. However, if she learns a lot, Alice and Bob will almost certainly detect her and abort the protocol. The key by itself doesn't tell Eve anything. QKD is vulnerable to denial of service attacks where Eve doesn't let Alice and Bob create a secret key, but it protects against situations where Eve learns the secret key but Alice and Bob do not know that.

6. Alice and Bob calculate the parities of bits corresponding to the parity checks of some classical error-correcting code. They announce the parities and compare, using the results of the comparisons to correct any errors, disagreements in their bit strings. They discard one bit per parity announced.

7. Alice and Bob perform *privacy amplification:* They choose random subsets of their current lists of bits and calculate the parity of each subset. They keep the parities secret and use them to form the *final key.* The remaining bits are discarded.

   - We need this step because Eve might have learned a small amount about the key in the earlier steps. The experimental error rate might be lower than we expect, so Eve could be hiding her eavesdropping to make it appear like natural noise. Thus, Eve could have learned a small fraction of the raw key. Privacy amplification washes away most of Eve's information, since when we take

a parity, Eve must know something about all the bits in order to have much information about the parity.

In the end, we get the security condition that, for any attack by Eve, either there is a large chance (exponentially close to 1) she will be caught and Alice and Bob abort the protocol, or with high probability (exponentially close to 1), Eve has an exponentially small amount of information (much less than a single bit) about the final key.

1.3 Attacks on QKD

A full proof of security of a QKD protocol is a bit challenging, since you must take into account everything that Eve can do, including things like massive entangled measurements on all the qubits sent by Alice. Nevertheless, many different security proofs exist for BB84. The security proof then tells you the maximum error rate Alice and Bob can tolerate before Eve can overcome the privacy amplification and learn information about the final key. Depending on the error correction and privacy amplification procedures Alice and Bob use, they can tolerate up to an error rate of about 19% in step **??**, although the higher the error rate, the more bits they have to use for correcting errors and for privacy amplification.

However, that is not the end of the story on security, since a proof is only as good as the assumptions made for the theorem. In particular, the devices used in the theoretical proof are idealizations and the real devices are slightly different. Eve can sometimes exploit those differences to attack and even break a QKD system.

1.3.1 Denial of Service Attack

One clear drawback of QKD is that it is very vulnerable to denial of service attacks. Simply by creating noise, Eve can cause Alice and Bob to abort their key generation attempt. They have no real recourse except to try again later, hoping that Eve has gone away.

1.3.2 Photon-Number Splitting Attack

One major class of attacks is the photon number splitting attack. The real photon sources that Alice uses might be weak coherent states rather than true single-photon states. That means that sometimes there are two photons rather than just one. Eve therefore has the following attack: For each mode passing through the quantum channel, Eve can measure the number of photons. This is not something Bob checks in BB84, so she can do this as much as she likes without getting caught by Bob. If the state happens to be a two-photon state, then Eve puts it through a beam splitter, which will send one photon on to Bob and let her keep the other photon, at least some of the time (and she can again tell when by measuring the photon number). Both photons have the same polarization because they are in the same mode created by Alice. She can store her copy of the photon in a quantum memory, and then later when Alice announces which basis she used, Eve is ready to measure in that basis, getting full information about that particular bit without creating any extra error rate.

Now, when the state is a very weak coherent state, most of the time there is no photon and only rarely are there two photons, so Eve only gains a small amount of information and that is wiped out by the privacy amplification. However, in real quantum channels, there is some loss of photons along the way. Alice and Bob will generate the raw key from only those photons that actually arrive. Eve can modfiy the characteristics of the channel (remember she can do *anything* allowed by quantum mechanics) so that one-photon states are more likely to get lost and two-photon states are less likely to. That means that the fraction of the raw key bits known by Eve could potentially be quite a bit larger than one might expect. Over longer distances (where the loss rate is larger), QKD can therefore become insecure against this kind of attack.

One solution is to modify the protocol. For instance, by occasionally sending "decoy states" with a different strength than the main ones used in the protocol and monitoring loss rates on them as a function of their strength, Alice and Bob can estimate the loss rate for different photon numbers and defeat this attack.

1.3.3 Detector Blinding Attack

Another attack on QKD takes advantage of properties of the detectors. If bright light is shined on a photodetector, it is "blinded" and won't detect any photons for some amount of time. In some QKD configurations, Eve can potentially send to Bob light much brighter than Alice ever sends and blind some of Bob's detectors, specifically ones that correspond to measuring in a particular basis. Eve then has an attack where she chooses a random basis to mesaure in and blinds Bob's detectors corresponding to the other basis. Then if she guessed right about the basis, she will learn the value of the key bit, and if she guessed wrong, then the detector is blinded and Bob doesn't receive a photon. Again, Eve has an opportunity to gain information without being detected, making the protocol insecure. This has to be taken into account in the design of the system, for instance by monitoring the strength of the light entering Bob's detectors.