Lecture Notes

Quantum Cryptography Week 6:

Quantum key distribution protocols

Contents

Last week we saw the definition of a correct and secure key distribution protocol. A *quantum key distribution* (QKD) protocol allows Alice and Bob to harness the advantages of quantum information processing to generate a shared secret key. Wiesner already suggested a quantum key distribution protocol in the 70s [Wie83]. The most well known, and indeed the oldest QKD protocol with a name is called BB84, after the inventors Bennett and Brassard [BB84].

This week we focus on the BB84 protocol. It turns out that for this protocol it is enough for us to prepare and measure single qubit quantum states. Let us thus imagine that Alice and Bob are connected by a quantum channel: unlike quantum computing, quantum cryptography is feasible *today*, at least over short distances, and it is reasonable to assume that Alice and Bob can transmit qubits over an optical fiber.

6.1 BB84 Quantum key distribution

In last week's lectures we discussed a special classical channel, where Eve is guaranteed to have some amount of noise in her attempts at intercepting bits transmitted by Alice to Bob. In such a classical protocol, if one takes away the guarantee about Eve but instead allow her to arbitrarily intercept messages on the special channel, then it is clear that there is no more security: Eve can learn all the bits of the string x. When considering a quantum channel, the classical protocol would amount to sending a string x encoded in a single fixed basis. For example, we might send x in the standard basis as $|x\rangle\langle x| = |x_1\rangle\langle x_1| \otimes \ldots \otimes |x_n\rangle\langle x_n|$. Eve, knowing the basis, can measure the transmitted quantum state to recover $x_1 \cdots x_n$ without error, copying each bit without causing any disturbance to the state. Of course, we might also encode x in a different basis, for example the Hadamard basis, as $H^{\otimes n}|x\rangle\langle x|H^{\otimes n}$. Yet, the fact remains, if Eve knows the basis, then she can copy the bits without being detected!

Exercise 6.1.1 Consider a bit b encoded in the Hadamard basis $H|b\rangle\langle b|H$. Give a measurement that recovers b (knowing it was encoded in the Hadamard basis!). Compute the post-measurement states for each possible outcome. What do you conclude? ∎

However, recall that by the no-cloning theorem presented in the Week 0 lecture notes, it is impossible to copy arbitrary qubits, i.e., qubits that could live anywhere on the Bloch sphere. This is precisely the case when Eve does not know the encoding in advance. We are thus motivated to let Alice not just choose bits x_j at random, but for each bit she will also randomly choose a basis θ_j. This gives rise to the BB84 encoding:

Definition 6.1.1 — BB84 states/encoding. The BB84 states are $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. This set of states corresponds to encoding a classical bit $x_j \in \{0,1\}$ in a randomly chosen basis $\theta_j \in \{0,1\}$ where $\theta_j = 0$ labels the standard basis, and $\theta_j = 1$ the Hadamard basis.

Note that the standard basis and the Hadamard basis are the eigenbases of the Pauli-Z and Pauli-X matrices respectively. Another similar set of states is known as the six-state encoding, which consists the eigenbases of the Pauli matrices X, Y and Z.

Definition 6.1.2 — Six-state encoding. The six-states are $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |+y\rangle, |-y\rangle\}$, where

$$|\pm y\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle). \tag{6.1}$$

This set of states corresponds to encoding a classical bit $x_j \in \{0,1\}$ in a randomly chosen basis $\theta_j \in \{0,1,2\}$, where $\theta_j = 0$ labels the standard basis, $\theta_j = 1$ the Hadamard basis, and $\theta_j = 2$ represents the eigenbasis of the Pauli Y matrix.

Both the BB84 basis and six-state basis are used frequently in quantum cryptographic protocols. This week, we will consider how BB84 states are used in a quantum key distribution protocol, focusing first on the case of noiseless transmission, i.e. the quantum channel is the identity channel, it transmits the quantum state without any errors. We first assume that Alice and Bob are connected via an classical authenticated channel (CAC), which they will use during the protocol. Later in the notes we will investigate how Alice and Bob could construct such a channel.

The BB84 protocol can be described as follows:

Protocol 1 — BB84 QKD (no noise). Outputs $k \in \{0,1\}^\ell$ to both Alice and Bob. Alice and Bob execute the following:

1. For a small real-valued parameter $\eta \ll 1$, and a large integer $n \gg 1$, Alice chooses a string $x = x_1, \ldots, x_N \in \{0,1\}^N$ uniformly at random where $N = (4+\eta)n$. She also chooses a basis string $\theta = \theta_1, \ldots, \theta_N$ uniformly at random. She sends to Bob each bit x_j by encoding it in a quantum state according to the basis θ_j: $H^{\theta_j}|x_j\rangle$.
2. Bob chooses a basis string $\tilde{\theta} = \tilde{\theta}_1, \ldots, \tilde{\theta}_N$ uniformly at random. He measures qubit j in the basis $\tilde{\theta}_j$ to obtain outcome $\tilde{x}_j$. This gives him a string $\tilde{x} = \tilde{x}_1, \ldots, \tilde{x}_N$.
3. Bob tells Alice over the CAC that he has received and measured all the qubits.
4. Alice and Bob tell each other over the CAC their basis strings θ and $\tilde{\theta}$ respectively.
5. Alice and Bob discard all rounds j in which they didn't measure in the same basis. Let $S = \{j | \theta_j = \tilde{\theta}_j\}$ denote the indices of the rounds in which they measured in the same basis. Since Alice and Bob chose $\theta, \tilde{\theta}$ at random, for large values of n, they throw away roughly $N/2 \approx 2n$ bits.
6. Alice picks a random subset[a] $T \subseteq S$ for testing and tells Bob T over the CAC. That is, Alice and Bob test roughly $|T| \approx N/4 \approx n$ bits.
7. Alice and Bob announce x_T and $\tilde{x}_T$ to each other over the CAC, where we denote by x_T the substring of x corresponding to the indices in the test set T. They compute the error rate $\delta = W/|T|$, where $W = |\{j \in T \mid x_j \neq \tilde{x}_j\}|$ is the number of errors when Alice and Bob did measure in the same basis.
8. If the error rate $\delta \neq 0$, then Alice and Bob abort the protocol. Otherwise, they proceed to denote $x_{\text{remain}} = x_{S \setminus T}$ and $\tilde{x}_{\text{remain}} = \tilde{x}_{S \setminus T}$ as the remaining bits, i.e., the bits where Alice and Bob measured in the same basis, but which they did not use for testing. The length of x_{remain} and $\tilde{x}_{\text{remain}}$ is approximately n bits.
9. Alice and Bob perform privacy amplification: Alice picks a random r, and computes $k = Ext(x_{\text{remain}}, r)$. She sends r to Bob, who computes $k = Ext(\tilde{x}_{\text{remain}}, r)$.

[a] A random subset T of S is where each element in S is included in T with probability $1/2$. By this definition, if $|S|$ is large, then $|T| \approx |S|/2$.

Note that even though steps 1 and 2 seem to take place one after the other, and you may be tempted to think that Alice and Bob require quantum storage, this is not necessarily the case. Alice can prepare the qubits one-by-one, and Bob can also measure them one-by-one. This is very appealing, since Alice and Bob only need very simple quantum devices - preparing and measuring single qubits is already enough!

Let us first investigate why the protocol is correct. If there are no errors in transmission, then whenever Bob measures in the same basis as the one chosen by Alice ($\theta_j = \tilde{\theta}_j$), then he learns the bit perfectly ($x_j = \tilde{x}_j$). If there is no eavesdropper, they will pass the test. Since $x_{\text{remain}} = \tilde{x}_{\text{remain}}$ and Bob knows r they produce the same output k as before.

But why should this protocol be secure? The intuition is that whenever Eve tries to intercept, and gain some information from the transmitted qubits, she will invariably disturb the quantum states — and Alice and Bob can detect such disturbance. It can be proven that

$$H_{\min}(X_{\text{remain}}|E) \gtrsim n\,[1 - h(\delta)]\ , \tag{6.2}$$

where $h(\delta) = -\delta \log \delta - (1-\delta)\log(1-\delta)$ is the binary entropy function, and recall from Protocol 1 that δ is the error rate between Alice and Bob's sample. We note that analyzing the sampling procedure, i.e., which qubits to test, for small values of N is an intricate problem which requires great care, as analyzed in [Pfi+16]. Here we will not dive into this but instead consider only the limit of large N.

In the case where errors in transmission occur, the error rate is always $\delta \neq 0$. This affects the correctness of Protocol 1, since Alice and Bob will get $x_{\text{remain}}, \tilde{x}_{\text{remain}}$ respectively, where $x_{\text{remain}} \neq \tilde{x}_{\text{remain}}$. In fact, when $x_{\text{remain}} \neq \tilde{x}_{\text{remain}}$, by property of the extractor Ext, with probability almost equal to 1, $Ext(x_{\text{remain}}, r) \neq Ext(\tilde{x}_{\text{remain}}, r)$. This means that almost for certain, Alice and Bob will end up with different keys!

To overcome this problem, Alice and Bob will have to perform an additional step of *information reconciliation* in the protocol. Thus instead of Protocol 1, they execute the following protocol:

Protocol 2 — BB84 QKD (with noise). Outputs $k \in \{0,1\}^\ell$ to both Alice and Bob.
Alice and Bob execute the following:

1-7. Same as Protocol 1.

8. If the error rate is larger than a certain threshold $\delta > \delta_t$, Alice and Bob abort the protocol. Otherwise, they proceed to denote $x_{\text{remain}} = x_{S\setminus T}$ and $\tilde{x}_{\text{remain}} = \tilde{x}_{S\setminus T}$ as the remaining bits, i.e., the bits where Alice and Bob measured in the same basis, but which they did not use for testing.
9. Alice and Bob perform information reconciliation: Alice sends some error correcting information C across the classical authenticated channel to Bob, and Bob corrects the errors in his string $\tilde{x}_{\text{remain}}$, so that he can obtain x_{remain} from the process as well.
10. Alice and Bob perform privacy amplification: Alice picks a random r, and computes $k = Ext(x_{\text{remain}}, r)$. She sends r to Bob, who computes $k = Ext(\tilde{x}_{\text{remain}}, r)$.

In Protocol 2, Alice and Bob allow for errors under the assumption that all the errors can be caused by a malicious Eve. They bound the amount of min-entropy Eve has about X_{remain} by (i) first invoking Eq. (6.2), and (ii) taking into account the amount of error correction information C sent across the channel from Alice to Bob. The size of k (given by the number of bits l) then depends on both δ and $|C|$.

Later on we will use the guessing game from last week in order to prove that the BB84 protocol presented aboveis secure for certain positive values of l.

6.2 Security of BB'84

To prove security of the BB'84 protocol we make two small modifications to the protocol. Although it will at first appear like these modifications give more power to the eavesdropper, they will facilitate the analysis.

6.2.1 A purified protocol

The first modification is rather benign. Consider the following two experiments. In the first experiment, Alice chooses $x, \theta \in \{0,1\}$ uniformly at random and returns $|x\rangle_\theta = H^\theta |x\rangle$, an encoding of the bit x in the basis specified by θ (the standard basis if $\theta = 0$ and the Hadamard basis if $\theta = 1$). In the second experiment, Alice first prepares an EPR pair $|\phi^+\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$. She then chooses a $\theta \in \{0,1\}$ uniformly at random and measures the first qubit in the basis $\{|0\rangle_\theta, |1\rangle_\theta\}$, obtaining outcome $x \in \{0,1\}$. She returns the second qubit.

We claim that the two experiment are absolutely equivalent. There are two things to verify. First, while in the first experiment Alice makes a choice of x uniformly at random, in the second experiment x is determined as the outcome of a measurement on the EPR pair. But we know that,

since the reduced density matrix of the EPR pair on the first qubit is the totally mixed state, any basis measurement on that qubit will return each of the two possible outcomes with probability $1/2$. So the distribution of x is identical in the two experiments.

Second, we should check that when Alice obtains outcome x by measuring the first qubit of the EPR pair in the basis θ, the qubit she returns, i.e. the second qubit of the EPR pair, is indeed projected onto the state $|x\rangle_\theta$. Again this is a property of the EPR state that is valid for any choice of basis measurement on the first qubit, so we are good — the two experiments are indeed equivalent.

Let us then consider an equivalent formulation of the BB'84 protocol in which, instead of directly preparing BB'84 states, Alice first prepares EPR pairs, keeps the first qubit of each pair to herself, and sends the second qubit to Bob. At a later stage she measures her qubit in a basis $\theta_j \in \{0,1\}$ chosen uniformly at random, and records the outcome x_j.

Thanks to the observation we made above this new formulation of the protocol is completely equivalent to the standard one. Even though it may look more complicated, the essential advantage of the new formulation is that it allows us to delay the moment in the protocol at which Alice needs to make her choice of basis. Although the difference is only conceptual, we can think of this delay as giving less power to Eve: we will now be able to easily argue that certain actions of the eavesdropper, taken early on in the protocol, could not have depended on Alice's basis choice, since the choice has not yet have been made at the time.

Here is the modified protocol in detail. For simplicity we again consider the case where there is no noise. It is called the "purified" BB'84:

Protocol 3 — Purified BB'84 (no noise).. Outputs $k \in \{0,1\}^\ell$ to both Alice and Bob.

1. For a small real-valued parameter $\eta \ll 1$, and a large integer $n \gg 1$, let $N = (4+\eta)n$. Alice prepares N EPR pairs $|\phi^+\rangle_{AB}$, and sends the second qubit of each pair to Bob. She chooses a uniformly random basis string $\theta = (\theta_1, \ldots, \theta_N) \in \{0,1\}^N$ and measures each of her qubits in the bases θ to obtain a string $x = x_1, \ldots, x_N$.
2. Bob chooses a uniformly random basis string $\tilde{\theta} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_N) \in \{0,1\}^N$. He measures the j-th qubit he received from Alice in the basis $\tilde{\theta}_j$ to obtain outcome $\tilde{x}_j$.
3. Bob tells Alice over the CAC that he received and measured all the qubits.
4. Alice and Bob exchange their basis strings θ and $\tilde{\theta}$ over the CAC.
5. Alice and Bob throw away the data from all rounds $j \in \{1, \ldots, N\}$ in which they didn't measure in the same basis. Let $S = \{j | \theta_j = \tilde{\theta}_j\}$ denote the indices in which they measured in the same basis.
6. Alice picks a random subset $T \subseteq S$ of size $T \approx |S|/2$ for testing and tells Bob T over the CAC.
7. Alice and Bob announce x_T and $\tilde{x}_T$ to each other over the CAC. They compute the error rate $\delta = W/|T|$, where $W = |\{j \in T \mid x_j \neq \tilde{x}_j\}|$ is the number of disagreements found in T. If δ is too large, they abort the protocol.
8. Let $R = S \backslash T$. Alice and Bob perform information reconciliation and privacy amplification on x_R.

The idea of considering a purified variant of the BB'84 protocol can be traced back to a different proposal for quantum key distribution put forward by Ekert in 1991 [Eke91]. Ekert's main insight was that if Alice and Bob were able to test for the presence of entanglement between their qubits, then (intuitively) by the monogamy of entanglement they would be able to certify that their systems are uncorrelated with Eve's. We will explore Ekert's protocol (and prove the intuition correct!) next week when we analyze quantum key distribution in the so-called device-independent setting.

Even though the purified protocol requires Alice to prepare EPR pairs, the formulation will only be used for the purposes of analysis. As we already discussed, from the point of view of

any eavesdropper which protocol Alice and Bob actually implement makes no difference at all, so it is perfectly fine to prove security of the purified protocol but use the original BB'84 protocol in practice. This is convenient because it is much easier to prepare single-qubit BB'84 states than to distribute EPR pairs across long distances.

6.2.2 More power to the eavesdropper

The second modification we make to the BB'84 is less benign, and will appear to give much more power to the eavesdropper. But once again it will be convenient for the analysis. Moreover, if we can prove security against stronger eavesdroppers without too much extra effort, why not do it?

The motivation for this second modification is that it is very hard to model the kinds of attacks Eve might apply to the quantum communication channel between Alice and Bob. For example, she might partially entangle herself with the qubits sent by Alice, creating a joint state ρ_{ABE} on which we have little control.

Exercise 6.2.1 Consider the case of a single EPR pair ($n = 1$), and suppose that Eve applies a CNOT on her qubit $|0\rangle_E$, controlled on the qubit B that Alice sends to Bob (Eve then forwards the qubit over to Bob). Compute the resulting joint state ρ_{ABE}. Compute the probability that Alice and Bob choose the same basis $\theta = \tilde{\theta}$ and obtain $x = \tilde{x}$. Is this a good attack? ∎

Because it is hard to model general intercepting attacks of the form described in the exercise, we will modify the protocol by allowing Eve to prepare an arbitrary pure state ρ_{ABE}, where the A and B systems are each made of N qubits, then give A to Alice, B to Bob, and keep E to herself. Alice and Bob will each measure their respective qubits using random choices of bases as in the protocol, and proceed from there on. By giving more power to Eve (she prepares the states, instead of Alice) we're preventing ourselves from thinking too hard about having a model for the attacks: in the new setup, Eve can prepare any state she likes!

This may sound crazy: if we let the eavesdropper prepare any state, then why doesn't she choose, say, $\rho_{ABE} = |000\rangle_{ABE}^{\otimes N}$? Observe that such a state would pass the "matching outputs" test when $\theta_j = \tilde{\theta}_j = 0$ (standard basis), but it would completely fail whenever $\theta_j = \tilde{\theta}_j = 1$ (Hadamard basis). So even though we're claiming Eve could prepare any state she likes, not all states will be accepted by Alice and Bob in the "matching outputs" test they perform in Step 7. How powerful is this test? Can it be used to certify that the state handed over by Eve indeed has the correct form, of being (close to) a tensor product of N EPR pairs? This may sound surprising, as the test only involves local measurements: can local measurements really detect entanglement? The answer is yes. Let's see how it works.

6.2.3 Locally implementing an entangled measurement

Suppose we modified the purified BB'84 protocol by adding an initial step as follows:

0. *Upon receiving their N respective qubits from Eve, Alice and Bob jointly measure each pair of qubits using the two-outcome POVM* $\{|\phi^+\rangle\langle\phi^+|_{AB}, \mathbb{I}_{AB} - |\phi^+\rangle\langle\phi^+|_{AB}\}$*, where* $|\phi^+\rangle_{AB}$ *denotes the EPR pair on Alice Bob's joint system. If the number of pairs of qubits that were not found to equal* $|\phi^+\rangle_{AB}$ *is larger than* δn *they abort Protocol 3. Otherwise, they proceed as usual.*

With this modification the protocol is clearly secure, and the tests performed in step 7 have become superfluous. Indeed, after the completion of step 0, Alice and Bob already have the guarantee that at least $(1-\delta)n$ of their shared pairs of qubits are perfect EPR pairs (since they are projected in the post-measurement state $|\phi^+\rangle$). In particular, any bit of the raw key obtained from measurements on these states is perfectly uniform and uncorrelated with Eve (remember that correlations generated from pure states are always perfectly monogamous).

The problem with step 0 is that it requires Alice and Bob to perform a joint entangled measurement, which they cannot implement locally. Or can they?

Exercise 6.2.2 Suppose given a tripartite state ρ_{ABE}, where A and B are each systems of a single qubit. Show that the probability that a measurement of systems A and B in the standard basis results in matching outcomes is exactly $\mathrm{Tr}(\Pi_1 \rho_{AB})$, where

$$\Pi_1 = |\phi^+\rangle\langle\phi^+| + |\Psi_{01}\rangle\langle\Psi_{01}|, \qquad \text{and} \qquad |\Psi_{01}\rangle = \frac{1}{\sqrt{2}}|00\rangle - \frac{1}{\sqrt{2}}|11\rangle. \tag{6.3}$$

Similarly, show that if the measurement is performed in the Hadamard basis then the probability of obtaining matching outcomes is $\mathrm{Tr}(\Pi_2 \rho_{AB})$, with

$$\Pi_2 = |\phi^+\rangle\langle\phi^+| + |\Psi_{10}\rangle\langle\Psi_{10}|, \qquad \text{and} \qquad |\Psi_{10}\rangle = \frac{1}{\sqrt{2}}|01\rangle + \frac{1}{\sqrt{2}}|10\rangle. \tag{6.4}$$

■

The exercise suggests that the "matching outcomes" test that Alice and Bob implement in step 7 of the original Protocol 3 is essentially equivalent to its replacement step 0 introduced above. Therefore, the security of Protocol 3 with step 0 implemented should directly imply the security of the protocol without step 0, but with step 7 instead.

To express the relation between the two steps, consider the reduced density matrix $\rho_{A_jB_j}$ of the state prepared by Eve on any two qubits for Alice and Bob, and let p_j be the probability of succeeding in the "matching outcomes" test, averaged over the choice of a uniformly random (but identical for both A_j and B_j) basis in which to perform the test. Then by expressing $\rho_{A_jB_j}$ in the Bell basis as

$$\rho_{A_jB_j} = q_{00}|\phi^+\rangle\langle\phi^+| + q_{01}|\Psi_{01}\rangle\langle\Psi_{01}| + q_{10}|\Psi_{10}\rangle\langle\Psi_{10}| + q_{11}|\Psi_{11}\rangle\langle\Psi_{11}|, \tag{6.5}$$

using the expressions for Π_1 and Π_2 obtained in Exercise 6.2.2 you can check that the condition

$$q_{00} = \langle\Psi_{00}|\rho_{A_jB_j}|\Psi_{00}\rangle \geq 2p_j - 1 \tag{6.6}$$

is satisfied. In particular, if the probability of success in the test is close to 1, say $p_j = 1 - \delta$ for some small δ, then the overlap q_{00} is correspondingly large, at least $1 - 2\delta$.

If the test performed in step 7 of Protocol 3 was really equivalent to our hypothetical step 0 projecting all pairs of qubits on EPR pairs, then we would be done with our proof of security. However, although the intuition is valid the argument is not quite complete: the two tests are not *exactly* equivalent, and making the argument precise is going to require more work.

A first distinction is that, in Protocol 3, step 7 is performed on the rounds T selected for testing, whereas it is the rounds in $R = S\backslash T$ that are used for the raw key (the outputs used for the raw key are never tested for equality, as this would leak them to Eve!). Another difficulty is that the results of the tests performed in different rounds are not independent from each other: although Alice and Bob make independent measurements, the state ρ_{ABE} prepared by Eve does not necessarily have a tensor product form.

The second distinction raises a thorny difficulty, to which we'll return in more detail next week. For now, let's concentrate on the first objection: how do we infer conditions on the qubits in rounds $j \in S\backslash T$ from results of tests performed on the qubits in rounds $j \in T$?

6.2.4 A concentration inequality

Let us summarize the situation. Suppose for simplicity that the number $|S|$ of rounds in which Alice and Bob make the same basis choice is exactly $|S| = 2n$, and that T has size $|T| = |S|/2 = n$. For

each $j \in S$, introduce an indicator random variable $Z_j \in \{0,1\}$ such that $Z_j = 1$ indicates failure in the matching outcomes test: $Z_j = 0$ if and only if $x_j = \tilde{x}_j$. With this notation the condition verified by Alice and Bob at step 7 of Protocol 3 can be written as $\sum_{j\in T} Z_j \leq \delta|T|$. In order to analyze security of their key, however, they would like to bound $\sum_{j\in S\setminus T} Z_j$. How can we do this?

The key idea is to use the fact that T is chosen as a random subset. Intuitively the average number of failures in T should be about the same as the average in the whole of S: indeed, which rounds are included in T or not is chosen at random by Alice, independently from whether the outcomes in those rounds happened to match or not.

The main tool required to make this intuition precise is called a concentration bound. There are many such bounds. The most widely used are usually referred to as the "Chernoff bound" or "Hoeffding's inequality", which is a generalized version of the Chernoff bound. If you have never heard of them, go look them up! The following is a variant of the Chernoff bound that turns out to be perfectly tuned for our scenario:

Theorem 6.2.1 — Lemma 7 in (TL15). Let $m = n + k$ and consider binary random variables $X_1, \ldots, X_m$. (The X_i may be arbitrarily correlated.) Let T be a uniformly random subset of $\{1, \ldots, m\}$ of size k. Then for any $\delta, \nu > 0$,

$$\Pr\left(\sum_{j\in T} X_j \leq \delta k \ \wedge \sum_{j\in\{1,\ldots,m\}\setminus T} X_j \geq (\delta+\nu)n \right) \leq e^{-2\nu^2 \frac{nk^2}{(n+k)(k+1)}}. \tag{6.7}$$

To see what the theorem says in our setting, set $m = |S| = 2n$ and $k = n$. Let's also choose $\nu = \delta$ for convenience. Plugging in these parameters we get the bound

$$\Pr\left(\sum_{j\in T} Z_j \leq \delta n \ \wedge \sum_{j\in S\setminus T} Z_j \geq 2\delta n \right) \leq e^{-\delta^2 n}, \tag{6.8}$$

which is valid for any choice of $\delta > 0$. This bound implies that the probability that the test performed in step 7 passes, but the outcomes obtained in the non-tested rounds $R = S\setminus T$ do not match in a fraction larger than 2δ of these rounds, is tiny — exponentially small in n! Writing ABORT to denote the event that Alice and Bob abort in Step 7 of Protocol 3, we can use Bayes' rule to rewrite the bound above as

$$\Pr\left(\sum_{j\in S\setminus T} Z_j \geq 2\delta n \,\middle|\, \neg\text{ABORT} \right) \leq \frac{e^{-\delta^2 n}}{\Pr(\neg\text{ABORT})}. \tag{6.9}$$

Writing the bound in this way points to an important subtlety in how the security of quantum key distribution is defined. As you can see, the bound is only good if $\Pr(\neg\text{ABORT})$ is not too small; if this probability was extremely tiny, then the right-hand side of Eq. (6.9) would suffer a corresponding blow-up. The probability that the protocol does not abort is not something that we can control or test, and it is natural that this probability has to be taken into account when defining security: we should always allow the protocol to have a very small probability of not aborting, in which case no claim can be made on the security. We will see a precise definition for the security of quantum key distribution next week.

Unfortunately we are still not done, due to the issue of dependencies between different tests, which may arise due to the eavesdropper preparing a state that is not in tensor product form.

If the state ρ_{ABE} *is* a tensor product across different rounds, i.e. it is of the form $\rho_{ABE} = \otimes_{j=1}^{n} \rho_{A_jB_jE_j}$ then we can complete the proof. Using that the choice of basis $\theta_j, \tilde{\theta}_j$ for $j \in R$ is uniformly random, we can conclude from the bound in Eq. (6.9) that a large fraction of $j \in R$ are such that the state $\rho_{A_jB_j}$ *would* pass the matching outcomes test, in *both* bases, with high probability

(this is because, if it were not the case, there would be a sufficiently high chance that we make a choice of basis with respect to which the state fails the test, leading to a contradiction with Eq. (6.9)). From the analysis in Section 6.2.3 we can deduce that $\rho_{A_jB_j}$ has a correspondingly large overlap with an EPR pair, and thus that the outcomes obtained by Alice and Bob when measuring in the same basis are highly correlated with one another, but (due to monogamy) have very weak correlation with Eve's system. Working out the parameters will give us a bound on the min-entropy of X_j in each round $j \in R$, which can be added up over all rounds by using the independence of different rounds.

If the state ρ_{ABE} is not a tensor product, unfortunately the analysis becomes more difficult; for instance the min-entropy does not add up easily across rounds. We will give a detailed analysis under the independence assumption next week, and we will also discuss the non-independent case in greater detail.

6.3 Authentication

Let us return to an important assumption that is always made when considering the BB'84 protocol: that the communication channel between Alice and Bob, that we've been calling the "CAC", is what its name implies — an authenticated channel. This assumption is used to guarantee that, although the eavesdropper may intercept any communication, she cannot "impersonate" Alice or Bob by sending fake messages on the channel. You can easily imagine the catastrophic consequences that such an attack could have: for example Eve could lie to Bob about Alice's choice of the test set T, making it a bit bigger; then Bob would reveal his outcomes in the bigger set, and Eve could keep them as additional side information about Alice's raw key.

This assumption is usually considered "benign", as indeed the access to an authenticated channel is a prerequisite for a large variety of cryptographic tasks. Thus in practice one imagines that any two parties Alice and Bob wanting to implement QKD have access to such a channel, that has been implemented by other means.

It is still interesting to consider how reasonable the assumption is, and how an authentication channel can be constructed. There are two main methods to achieve authentication, and each has its drawbacks. The first is to use public-key cryptography, which requires computational assumptions on the power of the eavesdropper. The second is to use private-key cryptography, which requires Alice and Bob to share a secret key... precisely the task QKD is meant to solve in the first place!

6.3.1 Everlasting security

The first method for authentication, based on public-key cryptography, can be quite efficient. We won't give any details here; the main primitive used is called a "digital signature", and it can be implemented based on any trapdoor one-way function (such as RSA, but of course with quantum adversaries you wouldn't want to rely on such an assumption!). You can check Chapter 5 in the notes [PS10] for more details.

But is it reasonable to make computational assumptions, when one of the main goals of quantum key distribution is to provide information-theoretic security? An argument in favor of this solution puts forward the property of "everlasting security". For the key generated in the protocol to be secure, it is sufficient that the CAC remains authenticated for the duration of the protocol, a few seconds at most. During this time it is crucial that Eve is not able to send fake messages. But once the protocol has ended, Alice and Bob have generated their private key, and it is no longer relevant whether the channel remains authenticated or not. So the computational assumption guaranteeing security of the authenticated channel only needs to hold for a few seconds, and the key generated in the protocol will remain forever secure: Eve has no information about it, and will not be able to gain any additional information by breaking a communication channel that is no longer in use!

6.3.2 Message authentication codes

The second method to achieve authentication is based on private-key cryptography. Let us consider a method to do this: even though it will not provide a good solution in practice (as using it to implement the CAC in QKD would require an initial shared secret key longer than the key generated by the protocol), it will still give you a good idea for the flavor of such constructions.

The main primitive used to achieve authentication is called a *message authentication code*, or MAC. A MAC is specified by a triplet of procedures:

- $Gen : () \to \mathscr{K}$ is the key generation procedure. It takes as input a security parameter n and returns a key $k \in \mathscr{K}$.
- $Tag : \mathscr{K} \times \mathscr{M} \to \mathscr{T}$ is the tagging procedure. It takes as input a key k and a message m, and returns a tag $\sigma = Tag_k(m)$.
- $Ver : \mathscr{K} \times \mathscr{M} \times \mathscr{T} \to \{0,1\}$ is the verification procedure. It takes as input a key, a message, and a (claimed) tag for the message. It returns either "1", meaning the tag is declared valid, or "0", meaning it is declared invalid.

All three procedures are required to run in polynomial time. The key generation procedure needs to be executed jointly by Alice and Bob, so that they have access to the same shared secret key k. Once this has been performed, they can separate. When Alice wants to send a message m to Bob (such as her choice of bases in the BB'84 protocol), she sends m accompanied with its tag $\sigma = Tag_k(m)$. Upon receiving (m,σ) Bob checks the tag by running the verification procedure $Ver_k(m,\sigma)$.

There are two requirements for a MAC. The first is correctness: it should always be the case that $Ver_k(m, Tag_k(m)) = 1$. The second is security. Informally, security of a MAC means that no adversary, given access to as many valid (message,tag) pairs as desired, is able to generate a message that it has not yet seen together with a valid tag for that message (except with probability negligible in the security parameter n).

Let's see a simple construction for a *one-time* MAC. A one-time MAC is a MAC that allows to tag a single message, but no more (i.e. security breaks down as soon as the adversary gets to see more than one valid (message,tag) pair generated using the same key). The construction is based on a family of two-universal hash functions. In week 4 we saw a construction of such a family $\mathscr{F} = \{f_y : \{0,1\}^n \to \{0,1\}^\ell\}$, for any $\ell \leq n$, that had size 2^{2n}. Given any such family,

- Gen returns a uniformly random $y \in \{1,\ldots,|\mathscr{F}|\}$ used to index a function f_y from $\mathscr{F}$, where ℓ is chosen as $\ell = n$.
- $Tag_y(m)$ returns $\sigma = f_y(m)$.
- $Ver_y(m,\sigma)$ returns 1 if and only if $\sigma = f_y(m)$.

This procedure is clearly correct. Why does it satisfy one-time security? Suppose given a valid (m,σ) pair. Using the property of two-universality, we know that for any $m' \neq m$ and for any σ', $\Pr_y(f_y(m) = \sigma \wedge f_y(m') = \sigma') = 2^{-2\ell}$. Given that the probability, over a random y, that $f_y(m) = \sigma$ is $2^{-\ell}$, applying Bayes' rule we get that

$$\Pr_y\left(f_y(m') = \sigma' \middle| f_y(m) = \sigma\right) = 2^{-\ell}.$$

In other words, the equation $f_y(m) = \sigma$ that the eavesdropper obtains on y when given a valid (m,σ) pair doesn't allow it to guess a valid σ', for any $m' \neq m$ of its choice, with probability more than $2^{-\ell}$, which is what a random guess would provide. Choosing $\ell \approx n$ thus gives us a construction of a one-time secure MAC.

Unfortunately, this MAC lets us authenticate messages of length n, provided we have a key of size $\log|\mathscr{F}| = 2n$. Not so useful! In general this is unavoidable, as any MAC with information-theoretic security requires a key as long as the total length of message that is to be tagged. For longer messages, one can again rely on computational assumptions (but weaker than for public-key

schemes: one-way functions are enough) to construct "many-times" MAC that allow to tag many messages with the same key. Or one can maintain information-theoretic security, but use more complicated methods for authentication that involve a multiple-round interaction between Alice and Bob. If you're interested in learning more, a good place to start is Chapter 5 in [PS10].

Acknowledgments

 The lecture notes are written by Nelly Ng, Thomas Vidick and Stephanie Wehner. We thank David Elkouss, Kenneth Goodenough, Jonas Helsen, Jérémy Ribeiro, and Jalex Stark for proofreading.

Bibliography

[BB84] C. H. Bennett and G. Brassard. “Quantum Cryptography: Public Key Distribution and Coin Tossing”. In: *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing* (1984), pages 175–179 (cited on page 3).

[Eke91] Artur K Ekert. “Quantum cryptography based on Bell’s theorem”. In: *Physical review letters* 67.6 (1991), page 661 (cited on page 6).

[Pfi+16] Corsin Pfister et al. “Sifting attacks in finite-size quantum key distribution”. In: *New Journal of Physics* 18.5 (2016), page 053001 (cited on page 5).

[PS10] Rafael Pass and Abhi Shelat. *A Course in Cryptography*. Lecture notes available at `http://www.cs.cornell.edu/courses/cs4830/2010fa/lecnotes.pdf`. 2010 (cited on pages 10, 12).

[TL15] Marco Tomamichel and Anthony Leverrier. “A rigorous and complete proof of finite key security of quantum key distribution”. In: *arXiv preprint arXiv:1506.08458* (2015) (cited on page 9).

[Wie83] Stephen Wiesner. “Conjugate Coding”. In: *SIGACT News* 15 (1983), pages 78–88 (cited on page 3).