7 Quantum key distribution protocols

Week 6, Lecture 6.1,Lecture 1: Intro to QKD
Week 6, Lecture 6.2,Lecture 1: Definitions and concepts in QKD

In the previous chapter we saw the definition of a correct and secure key distribution protocol, and we studied a simple example of such a protocol that works in a restricted setting. In this chapter we tackle the real thing: we construct a *quantum* key distribution (QKD) protocol, that obtains security using only a public quantum channel (and, as always, a CAC)! Informally, a QKD protocol allows two honest users Alice and Bob to harness the advantages of quantum information processing to generate a shared secret key. The most well known, and indeed the first QKD protocol that was discovered is called BB'84, after its inventors Bennett and Brassard and the year in which their paper describing the protocol was published. In this chapter we describe the BB'84 protocol and we introduce the main ideas for showing that the protocol is secure.

7.1 BB'84 Quantum key distribution

Week 6, Lecture 6.3,Lecture 1: BB84 states and six state protocol states
Week 6, Lecture 6.4,Lecture 1: The BB84 protocol

In the previous chapter we presented a key distribution protocol, Protocol **??**, that could be used when Alice and Bob have access to a very special classical channel: a channel such that Eve receives a noisy copy of each bit communicated over the channel, with some guaranteed minimal amount of noise $q > 0$. This assumption facilitated the analysis, but it is not realistic as in general there is no way to tell that noise has to be applied for Eve (and not for Bob). Our main idea in this chapter is to use a quantum channel to "simulate" the guarantee provided by this classical channel. As in the previous chapter, in addition to a public quantum channel we assume that Alice and Bob have access to a classical authenticated channel (CAC).

7.1.1 The BB'84 encoding

Suppose first that we use the quantum channel exactly as a classical channel. At the first step, Alice sends the string $x = x_1 \cdots x_n$ as a quantum state by encoding it in the standard basis. This can be written as $|x\rangle\langle x| = |x_1\rangle\langle x_1| \otimes \ldots \otimes |x_n\rangle\langle x_n|$. Since the basis is fixed it is public knowledge: Eve knows it as well and she can measure the transmitted quantum state in the standard basis to immediately recover x without error. In other words, she can easily copy x on the fly and later use her copy to correctly guess Alice and Bob's entire key. Clearly this is not any better than sending x directly over the CAC, and it doesn't work: the message has no privacy at all. Obviously it won't work either if we use any other fixed basis to encode the information.

However, recall that by the no-cloning theorem from Chapter **??** it is impossible to copy *arbitrary*

qubits, i.e. qubits that are not deterministically prepared in a fixed basis. This suggests an idea: in addition to her random string x, for each bit x_j of x Alice will randomly choose a basis $\theta_j \in \{0,1\}$ to encode the bit, where as usual 0 is used to designate the standard basis and 1 the Hadamard basis. She will then send the bit x_j encoded in the basis θ_j, which we denote $|x_j\rangle_{\theta_j} = H^{\theta_j}|x_j\rangle$ with H the Hadamard matrix. Since it is so important, let's record this encoding using a definition.

Definition 7.1.1 (BB'84 encoding) *The BB'84 states are $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. We can write each BB'84 state in the form $|x\rangle_\theta$ with $x \in \{0,1\}$ denoting the encoded bit and $\theta \in \{0,1\}$ the encoding basis, where $\theta = 0$ labels the standard basis and $\theta = 1$ the Hadamard basis.*

The intuition for using this encoding is that since the eavesdropper does not know the basis in which the bits are encoded, she cannot measure them directly. Moreover, by the no-cloning theorem she cannot copy them perfectly as quantum states either. Indeed, we already used this intuition in Chapter **??** when we considered the use of BB'84 states for the problem of quantum money. But now our the setting is different: for example, what if Eve simply keeps the state, and replaces it by some kind of dummy state that she forwards to Bob? In that case, can Bob detect that he did not receive the correct information? For this there should be some kind of check that Bob can perform in the protocol. Finding such a check raises some other issues. For example, does Bob know the correct basis? If so, how did he learn it? If not, how does he recover x? We will describe the BB'84 protocol soon, but you may wish to pause for a moment and think for yourself how you would build on the idea of using BB'84 states to design a complete QKD protocol.

While you're thinking, let us observe that the standard basis and the Hadamard basis are the eigenbases of the Pauli-Z and Pauli-X matrices respectively. A more complicated set of states than the BB'84 states that one could use consists of the eigenbases of the Pauli matrices X, Y and Z. This is known as the six-state encoding.

Definition 7.1.2 (Six-state encoding) *The six states are $\{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |+y\rangle, |-y\rangle\}$, where*

$$|\pm y\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle \pm i|1\rangle\big)\,. \tag{7.1}$$

We can write each such state as $|x\rangle_\theta$ with $x \in \{0,1\}$ the encoded bit and $\theta \in \{0,1,2\}$ the encoding basis, where $\theta = 0$ labels the standard basis, $\theta = 1$ the Hadamard basis, and $\theta = 2$ the eigenbasis of the Pauli Y matrix.

Both the four BB'84 states and the six states are used frequently in quantum cryptographic protocols. Here for simplicity we focus on the BB'84 states.

7.1.2 The BB'84 protocol

We are ready to describe the BB'84 protocol! We first give a simplified version of the protocol that can be used when the users expect that their quantum channel is noiseless, meaning that in the absence of any eavesdropping Bob can expect to perfectly receive any qubit sent by Alice. In practice there will always be errors on the communication channel, because no transmission can ever be perfect. We will later modify the protocol to allow this.

Protocol 1 (BB'84 QKD (no noise)) *The protocol depends on a large integer n publicly chosen by the users (intuitively n is the number of bits of key that they expect to generate at the end). Let $N = 4n$. Alice and Bob execute the following.*

1 *Alice chooses a string* $x = x_1, \ldots, x_N \in \{0,1\}^N$ *uniformly at random and a basis string* $\theta = \theta_1, \ldots, \theta_N \in \{0,1\}^N$ *uniformly at random. Alice sends to Bob each bit* x_j *by encoding it in a quantum state according to the basis* θ_j *as* $|x_j\rangle_{\theta_j} = H^{\theta_j}|x_j\rangle$.

2 *Bob chooses a basis string* $\tilde{\theta} = \tilde{\theta}_1, \ldots, \tilde{\theta}_N \in \{0,1\}^N$ *uniformly at random. For each* $j = 1, \ldots, N$ *he measures the* j*-th qubit received from Alice in the basis* $\tilde{\theta}_j$ *to obtain an 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 *Let* $S = \{j | \theta_j = \tilde{\theta}_j\}$ *denote the indices of the rounds in which Alice and Bob measured in the same basis. Alice and Bob discard the information for all rounds not in* S.

6 *Alice picks a random subset* $T \subseteq S$ *by flipping a fair coin for each* $i \in S$ *to decide if it is selected in* T*. Alice tells Bob what* T *is over the CAC.*

7 *"Matching outcomes" test: 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*, and similarly for* $\tilde{x}_T$*. They compute the error rate* $\delta = \frac{1}{|T|}|\{j \in T \mid x_j \neq \tilde{x}_j\}|$.

8 *If* $\delta \neq 0$ *then Alice and Bob abort the protocol. Otherwise, they proceed to denote* $x_{\mathrm{remain}} = x_{S \setminus T}$ *and* $\tilde{x}_{\mathrm{remain}} = \tilde{x}_{S \setminus T}$ *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 in the previous step.*

9 *Alice and Bob return* x_{remain} *and* $\tilde{x}_{\mathrm{remain}}$ *as their output, respectively.*

Remark 7.1.1 *The first two steps of the protocol are described as taking place one after the other. However, in an actual execution of the protocol 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 enough, and no quantum storage is required.*

There's a lot going on in this protocol! The most important step is the simplest, step 3: this step guarantees that Bob has measured all his qubits *before* the basis choices θ and $\tilde{\theta}$ are announced publicly. As we will see, this is crucial for security.

Exercise 7.1.1 Show that if Alice announces θ at any step prior to step 3 then the protocol is completely insecure. Namely, there is a way for an eavesdropper Eve, having only access to the CAC and the communication on the quantum channel, to learn both users' entire output in the protocol.

Let's do a quick "back of the envelope" calculation to estimate the number of output bits that are produced in this protocol. At step 5, since Alice and Bob chose $\theta, \tilde{\theta}$ at random we expect that on average they will discard roughly $|S| \approx N/2 = 2n$ bits. The size of T will be $|T| \approx |S|/2 \approx n$ bits, and so the length of x_{remain} and $\tilde{x}_{\mathrm{remain}}$ is also approximately n bits.[1]

Now, is this protocol correct and secure? Informally, based on the fact that Alice and Bob obtained exactly the same outcomes $x_T = \tilde{x}_T$ we expect that it should also be the case that $x_{\mathrm{remain}} = \tilde{x}_{\mathrm{remain}}$, and so the protocol should be correct. Furthermore, the same condition should intuitively guarantee that Eve has learned very little information about x. This is because due to the no-cloning principle any "copying" that she might have attempted while the qubits were flying from Alice to Bob in step 1 would have been detected, because it couldn't have depended on the secret choice of θ. (Here we use the assumption that Alice has a "secure lab", as described in Box **??**! Otherwise Eve could peak into it and see x and θ right away.)

[1] The *key rate* of this protocol, defined as the ratio of the expected number of key bits produced divided by the total number of qubits exchanged, is approximately $\frac{1}{4}$. In practice one can perform various optimizations to improve this, such as using fewer rounds for the matching outcomes test in step 7 and biasing Alice and Bob's choice of basis to increase the likelihood that they make the same choice. For clarity we give the simplest possible formulation of the protocol, without such optimizations.

Of course this is just intuition and we have to make it precise. But first let's give a more realistic protocol that accounts for the fact that even without any eavesdropper Alice and Bob can't expect to receive exactly the same strings — there will always be some kind of error on their quantum communication channel. To keep a correct protocol we allow an error rate $\delta > 0$ at step 8 and introduce an additional step of *information reconciliation*. Moreover, at the last step Alice and Bob also perform *privacy amplification*. This leads to the following protocol.

Protocol 2 (BB'84 QKD (with noise)) *The protocol depends on a small constant $\delta_{max} > 0$ and a large integer n publicly chosen by the users. Let $N = 4(1+C\delta_{max})n$, where C is a large constant that can be determined from the security analysis. Let Ext be a two-universal extractor and H a parity check matrix for a good classical error-correcting code. Alice and Bob execute the following:*

1-7. Same as Protocol 1.

8. If the error rate is $\delta > \delta_{max}$ then Alice and Bob abort the protocol. Otherwise they set $x_{\mathrm{remain}} = x_{S\setminus T}$ and $\tilde{x}_{\mathrm{remain}} = \tilde{x}_{S\setminus T}$ respectively.

9. Alice and Bob perform information reconciliation: Alice sends some error correcting information $c = Hx_{\mathrm{remain}}$ across the classical authenticated channel to Bob and Bob corrects the errors in his string $\tilde{x}_{\mathrm{remain}}$ to obtain a corrected string $\hat{x}_{\mathrm{remain}}$.

10. Alice and Bob perform privacy amplification: Alice picks a random seed r and computes $k_A = \mathrm{Ext}(x_{\mathrm{remain}}, r)$. She sends r to Bob, who computes $k_B = \mathrm{Ext}(\hat{x}_{\mathrm{remain}}, r)$.

The string x_{remain} obtained by Alice at step 8 is called the *raw key*. It is called like this because, after that point, only classical post-processing operations are performed: first, information reconciliation and then, privacy amplification. Note that we have not made precise the parameters of the information reconciliation subprotocol (the choice of H) or the privacy amplification subprotocol (the choice of Ext). We will discuss these later.

7.1.3 Correctness and security

In this section we sketch arguments for the correctness and security of the final BB'84 protocol, Protocol 2, focusing on the intuition. Making these arguments precise will occupy the remainder of the chapter.

Let us first argue correctness. Based on observing a certain error rate δ in step 7 the users can conclude that the strings x_{remain} and also $\tilde{x}_{\mathrm{remain}}$ are likely to match in about $(1-\delta)$ fraction of positions. Intuitively this is because the set T used for testing is chosen uniformly at random and so the fraction of errors inside T and outside it should be approximately the same; we will explain how to prove this formally in Section 7.4. Based on this estimate the users can decide on the exact parameters for the information reconciliation protocol to guarantee that, after information reconciliation, it holds that $\Pr(x_{\mathrm{remain}} \neq \hat{x}_{\mathrm{remain}}) \leq \varepsilon_c$ where ε_c is the target correctness error.

We now consider the secrecy condition. Let $h(\delta) = -\delta\log\delta - (1-\delta)\log(1-\delta)$ be the binary entropy function. The crux of the argument is to prove that at the end of step 8 of the protocol, conditioned on not having aborted in that step, the following bound holds

$$H_{\min}(X_{\mathrm{remain}}|E) \gtrsim n\ (1 - h(\delta_{\max}))\ , \tag{7.2}$$

where E designates all information available to the eavesdropper. In this equation the symbol $\gtrsim$ designates that we are ignoring lower-order terms (growing less fast than any linear function of n) and a logarithmic dependence on the probability of not aborting, which we will discuss later.

We will sketch a proof of (7.2) in Section 7.4, and generalize it in the next chapter. Let us see why it

is sufficient to show that the protocol is secret. Taking into account the additional information leaked to Eve when performing information reconciliation we get by the chain rule (Box **??**) that

$$H_{\min}(X_{\text{remain}}|EC) \gtrsim n\,(1 - h(\delta_{\max})) - |C|\ . \tag{7.3}$$

Based on the discussion of information reconciliation protocols from the previous chapter (Section **??**) we know that it is possible to ensure that $|C| \approx h(\delta_{\max})n$. Using what we know about privacy amplification (Section **??**) it is then possible to choose parameters for the extractor used to perform privacy amplification in the last step of the protocol so as to obtain an output key k_A that is ℓ bits long and ε_s-secret, where

$$\ell \approx (1 - 2h(\delta_{\max}))\, n - 2\log(1/\varepsilon_s)\ .$$

Quiz 7.1.1 *In the analysis of our cryptographic protocols we generally imagine that the adversary Eve is "all powerful." Consider a scenario where Eve has placed a transmitter in the random number generators of Alice and Bob, such that she can find out what are the random bits that Alice and Bob generate. Can Alice and Bob be guaranteed security against Eve in this case?*

a) Yes, it is possible for Alice and Bob to be secure against Eve even in this case.
b) Quantum Mechanics allows Alice and Bob to check whether such a transmitter has been placed in the random number generators.
c) $\longrightarrow$ ***No, one of the crucial assumptions for the security of quantum cryptographic protocols is that Eve has no access to the labs of Alice and Bob. That is, it is not possible for Alice and Bob to be secure against Eve in this case.***

Quiz 7.1.2 *Consider the following scenario. First, Alice prepares an eigenstate of the Pauli matrix X. Second, Eve measures this state uniformly at random in one of the three bases (i.e. in each of them with probability $p_i = \frac{1}{3}$): standard, Hadamard and the Y-basis. Third, Eve sends the post-measurement state to Bob. Bob then measures again in the X-basis. What is the probability that Bob's post-measurement state is the same state as the one that Alice prepared?*

a) $\frac{1}{3}$
b) $\frac{1}{2}$
c) $\longrightarrow \frac{2}{3}$
d) $\frac{3}{4}$

Quiz 7.1.3 *Alice and Bob run the BB'84 protocol but without the step where Bob announces the receipt of the states. Later in the testing stage they find out that their error rate is zero. They conclude that the quantum channel must be noise-free and that is there is no eavesdropper. Hence, omitting the step of confirmation of receipt of the states by Bob did not lead to any compromise of security in this case. Is the reasoning of Alice and Bob correct?*

a) Yes
b) $\longrightarrow$ ***No***

Before we look in more detail into showing both requirements, ε_c-correctness and ε_s-security, we pause to make sure that we understand what we mean when we say that Eq. (7.2) should hold "conditioned on not having aborted in step 8." In general we can always represent the state of the entire system of interest, which for our purposes consists of Alice and her random choices, Bob and his random choices, and Eve's quantum state, as a giant quantum state ρ_{ABE} where the A part also contains x and θ, the B part contains $\hat{x}$ and $\hat{\theta}$, etc. At step 8 we can imagine that each of the users initializes a special

"abort" register, and depending on their classical information they either write '0' (for "not abort") or '1' (for "abort") in that register. Because the classical communication channel is authenticated we know that at this step of the protocol both users make exactly the same decision. "Conditioned on not aborting" means that we measure the "abort" register for both users and post-select on the result being '0' for both of them. Here "post-select" means that we renormalize the state, exactly as if the outcome '0' had been 'forced'. The resulting state is the one on which (7.2) is evaluated. The following example will make this operation of postselection clear.

Example 7.1.1 *Suppose that Alice, Bob and Eve share the pure state*

$$|\psi\rangle_{ABE} = \frac{1}{\sqrt{3}}\big(|00\rangle_A |00\rangle_B |0\rangle_E + |01\rangle_A |01\rangle_B |0\rangle_E + |10\rangle_A |10\rangle_B |1\rangle_E \big) .$$

You can see that this state is in a superposition of three states, such that that A and B always have the same information, and E has a bit that is equal to their first bit. Now suppose that Alice and Bob, for some reason, each decide to abort in case the parity of their two bits is equal to 0. To determine the state "conditioned on not aborting" we first evaluate the abort condition in a new register A' for Alice and B' for Bob to get

$$|\psi'\rangle_{ABE} = \frac{1}{\sqrt{3}}\big(|00\rangle_A |1\rangle_{A'} |00\rangle_B |1\rangle_{B'} |0\rangle_E + |01\rangle_A |0\rangle_{A'} |01\rangle_B |0\rangle_{B'} |0\rangle_E + |10\rangle_A |0\rangle_{A'} |10\rangle_B |0\rangle_{B'} |1\rangle_E \big) .$$

Finally, we imagine measuring both A' and B' and forcing the outcome to a 0. After renormalization, the state is

$$|\psi_{\text{notabort}}\rangle_{ABE} = \frac{1}{\sqrt{2}}\big(|01\rangle_A |0\rangle_{A'} |01\rangle_B |0\rangle_{B'} |0\rangle_E + |10\rangle_A |0\rangle_{A'} |10\rangle_B |0\rangle_{B'} |1\rangle_E \big) .$$

This is the state "conditioned on not aborting". ∎

In general, assuming that Alice and Bob follow the correct actions of the protocol and that Eve has some given strategy, there is a well-defined probability of the protocol aborting in step 8. This is not a parameter that is known by the users (unless they repeat the protocol many times, but then they wouldn't know if Eve does the same thing each time or not), but it is a well-defined number. This number will appear in the security proofs. Intuitively, this is because if the probability of aborting is very close to 1 then it means that Eve is doing something pretty crazy, and Alice and Bob will detect this crazyness with probability close to 1. However, if by lack of luck they do not detect anything then we really can't guarantee any secrecy. This is a common feature of most cryptographic protocols: there is always a chance that things go wrong, and our goal as protocol designers is to minimize this chance. In other words, we want to obtain good security guarantees for probabilities of aborting that are as close to 1 as we can manage.

On a more technical level, the probability of not aborting will arise in the analysis precisely because the entropy on the left-hand side of the inequality (7.2) is evaluated on the state of the users and Eve at step 8, conditioned on not aborting. Due to a very large renormalization in case that the probability of not aborting is very small, the inequality that we are able to prove in our security analysis will get worse and worse as the probability of not aborting gets smaller.

7.2 A modified protocol

Week 6, Lecture 6.4,Lecture 2: Security in BB84

To facilitate the task of showing security for the BB'84 protocol, which we tackle in the next section, 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 make the analysis simpler.

7.2.1 The purified protocol

Week 6, Lecture 6.5,Lecture 1: Purifying protocols using entanglement

The first modification is straightforward. 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. In the second experiment Alice first prepares an EPR pair $|\mathrm{EPR}\rangle = \frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$. 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 an outcome $x \in \{0,1\}$. She returns the second qubit.

We claim that the two experiments are absolutely equivalent. To show this 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 $|x\rangle_\theta$. In Exercise **??** we showed that this is the case for both the standard and Hadamard bases, and in fact it is a property of the EPR state that is valid for any choice of basis measurement on the first qubit.[2] So it is true — 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. This new formulation of the protocol is completely equivalent to the standard one. Even though it may look more complicated, an important advantage of the new formulation is that it allows us to delay the moment in the protocol when Alice needs to make her choice of basis. We can think of this delay as giving less power to Eve: we will now be able to more 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 had not yet been made at the time.

Here is the modified protocol in detail. It is called the "purified" BB'84 protocol.

Protocol 3 (Purified BB'84) *Choose parameters as in Protocol 2. Perform the following:*

1 *Alice prepares N EPR pairs $|\mathrm{EPR}\rangle_{AB}$ and sends the second qubit of each pair to Bob.*
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 an outcome $\tilde{x}_j$.*
3 *Bob tells Alice over the CAC that he received and measured all the qubits.*

[2] Up to a transpose, which in the case of complex coefficients, such as for the Y eigenbasis, amounts to exchanging basis elements.

4 Alice 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 $x = x_1, \ldots, x_N$. Alice and Bob exchange their basis strings θ and $\tilde{\theta}$ over the CAC.

5–10 Same as Protocol 2.

Notice how we "pushed" Alice's choice of string x and measurement bases θ all the way from step 1 to step 4 of the protocol, without in fact changing anything about the actual outcomes of the protocol or the eavedropper's power.

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. 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.

Remark 7.2.1 *Even though the purified protocol requires Alice to prepare EPR pairs, this formulation will only be used for the purposes of analysis. 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.*

7.2.2 More power to the eavesdropper

The second modification we make to the BB'84 protocol is less benign, and will appear to give much more power to the eavesdropper. But we will see that it is also very 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 that 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, the mathematicians in charge of showing security, have little control.

Exercise 7.2.1 Consider the case of a single EPR pair ($n = 1$). Suppose that Eve initializes an extra qubit in the state $|0\rangle_E$ and applies a CNOT on it controlled on the qubit B that Alice sends to Bob in step 1 of the protocol (Eve then forwards the qubit B to Bob). Compute the joint state ρ_{ABE} that is created by this operation. 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 state ρ_{ABE}, where the A and B systems are each made of N qubits, and then give A to Alice, B to Bob, and keep E to herself. The protocol from step 2 onwards is unchanged: Alice and Bob will each measure their respective qubits using random choices of bases 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 outcomes" test from step 7 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 allowing Eve to prepare any state she likes, not all states will be accepted by Alice and Bob in the protocol. For example, you can calculate that the state ρ_{ABE} defined above

succeeds with probability about $(3/4)^{N/4}$: this is because roughly $N/4$ rounds are used for testing, and for each such round there is a probability $1/2$ that the basis is the Hadamard basis in which case the probability of a matching outcome is $1/2$.

This simple example shows that the "matching outcomes" test must play an essential role in the security analysis. 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? If we manage to show this then we'll be in good shape, because in the modified protocol the only step in which Eve can really have a chance to influence the quantum information exchanged by the users is step 1. It may sound surprising that we would be able to achieve this, as the test only involves local measurements: can local measurements detect entanglement? The answer is yes, and we'll soon see how it works.

7.3 Security of BB'84 key distribution

Let's show security! Based on our knowledge of privacy amplification, to show ε_s-security it suffices to prove an equation on the min-entropy of the form given in Eq. (7.2). Because it is the crux of the security proof, let's re-state the inequality here:

$$H_{\min}(X_{\text{remain}}|E) \gtrsim \kappa\, n\ . \tag{7.4}$$

Recall that in this equation, X_{remain} denotes the classical string in Alice's possession at step 8 of the protocol, and is called the raw key. E denotes all the information available to the eavesdropper Eve: her quantum state, which she created at step 1, as well as all the information exchanged by the users over the CAC. Finally, n is the expected length of X_{remain} and κ is a coefficient which we hope to show is as close to 1 as possible. After information reconciliation and privacy amplification, the users will be left with approximately $\kappa n - h(\delta_{\max})n - 2\log(1/\varepsilon_s)$ bits of key, where $h(\delta_{\max})n$ is the maximum number of bits used for information reconciliation and $2\log(1/\varepsilon_s)$ the number of bits lost due to privacy amplification.

We will give three different methods to show (7.4). Each of the methods has its advantages and disadvantages, and each gives a different insight on *why* the protocol is secure. The most intuitive, but quantitatively weakest, method is the one from the next section, which gives an interpretation of the matching outcomes test as a test for EPR pairs. In Section 7.3.2 we give a method based on the tripartite guessing game from Chapter **??**, whose major advantage is that it shows security under "general attacks," which we will define later. Finally in Section 7.3.3 we give a method based on entropic uncertainty relations that, when fully worked out, gives the best guarantees on the secret key rate; in particular it will let us get $\kappa = 1 - h(\delta)$ in (7.4).

7.3.1 Locally implementing a Bell basis measurement

Week 6, Lecture 6.5,Lecture 2: Implementing a Bell basis measurement locally

We start with a relatively informal argument for security, that will help us build intuition on the role played by the "matching outcomes" test. We will see that this test corresponds to a "virtual" projection of the state shared by Alice and Bob in an EPR pair. First, let's convince ourselves that this is all we need. For this, suppose that 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 $\{|\mathrm{EPR}\rangle\langle\mathrm{EPR}|_{AB}, \mathbb{I}_{AB} - |\mathrm{EPR}\rangle\langle\mathrm{EPR}|_{AB}\}$*, where* $|\mathrm{EPR}\rangle_{AB}$ *denotes an EPR pair on Alice and Bob's joint system. If the number of pairs of qubits that were not found to equal* $|\mathrm{EPR}\rangle_{AB}$ *is larger than* $\delta_{\max}N$ *then they abort. Otherwise, they proceed as usual.*

With this modification the protocol is immediately secure. Indeed, after the completion of step 0 Alice and Bob have the guarantee that at least $(1-\delta_{\max})N$ of their shared pairs of qubits are perfect EPR pairs (since they are projected in the post-measurement state $|\mathrm{EPR}\rangle$). Clearly, any bit of the raw key obtained from measurements on these states is perfectly uniform and uncorrelated with Eve. In this situation, getting a bound on the min-entropy such as (7.4) does not pose any difficulty.

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 7.3.1 Suppose we are 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

$$\Pi_1 = |\mathrm{EPR}\rangle\langle\mathrm{EPR}| + |\psi_{01}\rangle\langle\psi_{01}|, \quad \text{and} \quad |\psi_{01}\rangle = \frac{1}{\sqrt{2}}|00\rangle - \frac{1}{\sqrt{2}}|11\rangle\,. \tag{7.5}$$

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

$$\Pi_2 = |\mathrm{EPR}\rangle\langle\mathrm{EPR}| + |\psi_{10}\rangle\langle\psi_{10}|, \quad \text{and} \quad |\psi_{10}\rangle = \frac{1}{\sqrt{2}}|01\rangle + \frac{1}{\sqrt{2}}|10\rangle\,. \tag{7.6}$$

Now suppose that ρ_{AB} is any state such that

$$\frac{1}{2}\,\mathrm{Tr}\left(\Pi_1\rho_{AB}\right) + \frac{1}{2}\,\mathrm{Tr}\left(\Pi_2\rho_{AB}\right) \geq 1-\delta\,,$$

for some $\delta \geq 0$. Using the above, show that the fidelity

$$F\big(\rho_{AB}, |\mathrm{EPR}\rangle\langle\mathrm{EPR}|\big) \;=\; \sqrt{\langle\mathrm{EPR}|\,\rho_{AB}\,|\mathrm{EPR}\rangle} \;\geq\; \sqrt{1-2\delta}\,.$$

[Hint: to show this, imagine measuring ρ_{AB} *in the Bell basis. What can you say about the probability of each of the four possible outcomes?]*

The exercise suggests that the "matching outcomes" test that Alice and Bob implement in step 7 of Protocol 3 can play a similar role to the imaginary step 0 introduced above, because high success in the test implies high fidelity with an EPR pair. Therefore, the security of Protocol 3 with step 0 implemented should imply the security of the protocol without step 0, but with step 7 instead.

This sketch of a security proof provides the right intuition for security, and it can be worked out precisely. Rather than pursuing this route, we give two other arguments, each with its own advantages and disadvantages.

Quiz 7.3.1 *Recall our notation for the Bell states:*
$|\psi_{00}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$, $\quad |\psi_{01}\rangle = \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle)$,
$|\psi_{10}\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$, $\quad |\psi_{11}\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$
Suppose that Alice and Bob both measure their qubits in the standard basis, but want to select for opposite outcomes ($|0\rangle$ *and* $|1\rangle$ *respectively or vice versa). Which of the following projectors* Π *correspond to this scenario? (Recall that the matching outcomes test is equivalent to a projection onto a subspace spanned by Bell states.)*

a) $|\psi_{01}\rangle\langle\psi_{01}| + |\psi_{10}\rangle\langle\psi_{10}|$
b) $|\psi_{01}\rangle\langle\psi_{01}| + |\psi_{11}\rangle\langle\psi_{11}|$

c) $\longrightarrow$ $|\psi_{10}\rangle\langle\psi_{10}| + |\psi_{11}\rangle\langle\psi_{11}|$
d) $|\psi_{10}\rangle\langle\psi_{10}| - |\psi_{11}\rangle\langle\psi_{11}|$
e) None of the above, but some other linear combination of $|\psi_a\rangle\langle\psi_a|$.
f) No linear combination of $|\psi_a\rangle\langle\psi_a|$.

Quiz 7.3.2 *Suppose now that Alice measures her qubit in the standard basis while Bob measures his in the Hadamard basis. They want to select for the "same" outcome, i.e.* $(|0\rangle, |+\rangle)$ *or* $(|1\rangle, |-\rangle)$. *Which of the following projectors* Π *corresponds to this scenario?*

a) $|\psi_{01}\rangle\langle\psi_{01}| + |\psi_{10}\rangle\langle\psi_{10}|$
b) $|\psi_{01}\rangle\langle\psi_{01}| + |\psi_{11}\rangle\langle\psi_{11}|$
c) $|\psi_{10}\rangle\langle\psi_{10}| + |\psi_{11}\rangle\langle\psi_{11}|$
d) $|\psi_{10}\rangle\langle\psi_{10}| - |\psi_{11}\rangle\langle\psi_{11}|$
e) None of the above, but some other linear combination of $|\psi_a\rangle\langle\psi_a|$
f) $\longrightarrow$ ***No linear combination of*** $|\psi_a\rangle\langle\psi_a|$

Quiz 7.3.3 *Let Alice and Bob share* n *qubit pairs in the state* $\rho_{AB}^{\otimes n}$, *and suppose that the matching outcomes test succeeds with probability exactly* $p_j = 0.95$ *on each of the* n *pairs. What is the largest value of* n *for which the overlap* $\langle\psi_{00}|^{\otimes n}\rho_{AB}^{\otimes n}|\psi_{00}\rangle^{\otimes n}$ *is guaranteed to exceed* $1/2$*?*

a) 2
b) 4
c) $\longrightarrow$ ***6***
d) 8

7.3.2 Security from the tripartite guessing game

Our second proof of security leverages the tripartite guessing game from Chapter **??**. Let's remember roughly how that game proceeds; for details see Section **??**. In the game, Eve prepares an arbitrary state ρ_{ABE} such that A, B are one qubit each, and gives A to Alice and B to Bob. Alice and Bob choose a random basis $\Theta \in \{0, 1\}$, measure their qubit in that basis, and give Θ to Eve. They win if Alice and Bob's outcomes are equal, and furthermore Eve is also able to guess the same outcome by performing a measurement, which may depend on Θ, on E.

Sounds familiar? Well sure it does! This is exactly one round of the matching outcomes test, with the addition that we now also ask Eve to predict Alice and Bob's matching outcomes. Recall that in Section **??** we showed that the maximum success probability in this game is $p_{\text{succ}} \leq \frac{1}{2} + \frac{1}{2\sqrt{2}}$. Let's investigate what this means for our QKD protocol. First let's reformulate the maximum success probability in the tripartite guessing game as a guessing probability. Let X_A be Alice's outcome, X_E Eve's guess, and Ω the event that Alice and Bob's outcomes match. Then we have

$$\begin{aligned} p_{\text{succ}} &= p\big(\Omega \wedge (X_A = X_E)\big) \\ &= p(X_A = X_E|\Omega)p(\Omega) \\ &= p_{guess}(X_A|E\Theta\Omega)p(\Omega) \ . \end{aligned}$$

Here for the second line we used Bayes' rule, and for the third line we used that X_E is Eve's best guess for X_A, given the information available to her: her quantum state in register E, and the choice of basis Θ. Shuffling terms around, we get the bound

$$p_{guess}(X_A|E\Theta\Omega) \ = \ \frac{p_{\text{succ}}}{p(\Omega)} \ \leq \ \Big(\frac{1}{2} + \frac{1}{2\sqrt{2}}\Big)\frac{1}{p(\Omega)} \ .$$

Using the relation between guessing probability and min-entropy,

$$H_{\min}(X_A|E\Theta\Omega) \;\geq\; -\log\Big(\frac{1}{2}+\frac{1}{2\sqrt{2}}\Big) - \log\frac{1}{p(\Omega)}\;.$$

This is good progress! As we saw earlier, obtaining lower bounds on the conditional min-entropy of Alice's raw key, given the information available to the eavesdropper, is the most important step is showing security. The derivation above was done considering a single round, but by using the n-round version of the guessing game we can similarly get the bound

$$H_{\min}(X_A|E\Theta\Omega) \;\geq\; -n\log\Big(\frac{1}{2}+\frac{1}{2\sqrt{2}}\Big) - \log\frac{1}{p(\Omega)}\;, \tag{7.7}$$

where now X_A is Alice's n-bit outcome string, and Ω is the event that *all* Alice and Bob's outcomes match.

How does (7.7) compare to our target bound (7.4)? First of all, a minor difference is that we wrote the basis choice Θ explicitly, whereas in (7.4) it is included in E; this is just a question of notation.[3] Another minor difference is that here, the bound is on the entire X_A, not only $X_{\rm remain}$. This is easy to deal with because $X_{\rm remain}$ is a sub-string of X_A, and so by the data-processing inequality a lower bound of the min-entropy of the latter implies a lower bound on the min-entropy of the former.

A more important difference is that here, we are also conditioning on Ω, which is the probability that $X_A = X_B$. What do we know about this probability? If we considered the version of the protocol with $\delta_{\max} = 0$, meaning that Alice and Bob abort as soon as they see a difference, then the Ω would roughly be the same as the probability of not aborting. As expected, our bound on the entropy depends on how likely is this event. The difficulty is that the non-abort condition in the real protocol requires that X_A and X_B are close, but not identical. To deal with this we would have to define a new guessing game in which the winning condition is that Alice and Bob's outcomes match in $(1-\delta)$ fraction of positions, and Eve's outcome matches Alice's outcome (always). This game is a bit harder to analyze, but it can be done. The result is that the bound on the success probability becomes

$$p_{\rm win} \;\leq\; \Big(2^{h(\delta_{\max})}\Big(\frac{1}{2}+\frac{1}{2\sqrt{2}}\Big)\Big)^n\;,$$

which will make an additional $h(\delta_{\max})$ appear on the right-hand side of (7.7).[4]

Finally, we comment on the coefficient in front of n in (7.7). This coefficient is $\log(1/2+1/2\sqrt{2}) \approx 0.16$, not 1. This is a bit disappointing – it means that our analysis only guarantees that, at best, we will be able to obtain one bit of secure key for every (approximately) 5 rounds of communication in the protocol. Unfortunately this is a limitation of our method — as we saw, it *is* possible to succeed with probability $\frac{1}{2}+\frac{1}{2\sqrt{2}}$ in the guessing game, not only $\frac{1}{2}$ as would be needed to obtain a coefficient of 1. In principle we could do better by requiring that Alice and Bob have matching outcomes with high probability; this is because in the optimal strategy for the game they only agree with probability $\frac{1}{2}+\frac{1}{2\sqrt{2}}$, but in the QKD protocol we expect them to agree with a higher probability. As it turns out, once we do that the repeated version of the game, with n qubits, becomes much harder to analyze. In the next section we explore a different approach which uses this observation and gives a better constant (but has its own drawbacks!).

[3] Recall that in (7.4) E denotes all the information available to the eavesdropper. This includes the basis information Θ, which was exchanged over the CAC by the users.

[4] As expected, the bound is bigger than the one we had before, because the game is now easier.

Box 7.1 The i.i.d. assumption

When analyzing a multi-round cryptographic protocol it is common to make an assumption known as the *i.i.d. assumption*, where i.i.d. stands for "identically and independently distributed".

This assumption is composed of two parts, the "identically" and the "independently". We start with "independently". In the case of the purified BB'84 protocol this is the assumption that Eve prepares the state ρ_{ABE} as a tensor product $\rho_{ABE} = \rho_{A_1B_1E_1} \otimes \cdots \otimes \rho_{A_NB_NE}$. This assumption can be used to derive a bound on the general min-entropy, for all rounds together, from a bound on the min-entropy on each round, which would be obtained from the qualitative considerations in Section 7.3.1 or the quantitative arguments in Section 7.3.3. For a general quantum state the min-entropy does not simply add up across rounds. Informally this is because the min-entropy is related to the guessing probability, and in general Eve can make any global measurement on all her quantum information to predict all bits x at once.

The second component is "identically". This means that we also implicitly assume that each $\rho_{A_iB_iE_i}$ is equal. This assumption can be used to obtain information about some of the rounds (the rounds used to get the key) based on data gathered in other rounds (the test rounds).

While the i.i.d. assumption is commonly made and can greatly simplify the analysis, it is important to realize that it is only this, an assumption, and that in general the eavesdropper may not respect it!

7.3.3 Security from uncertainty relations

Week 6, Lecture 6.6,Lecture 1: Security from a guessing game
Week 6, Lecture 6.6,Lecture 2: A concentration inequality

To get the best quantitative bound on the security of the protocol, i.e. the largest possible κ in (7.4), we use *entropic uncertainty relations*. These relations will help us measure very precisely the trade-off between the success probability in the matching outcomes test, which is a measure of correlation between A and B, and the conditional min-entropy, which measures correlation between A and E. The main drawback of our approach is that it will require us to make an assumption about the eavesdropper's behavior which is referred to as the "i.i.d.", for "independently and identically distributed", assumption. This assumption is explained in detail in Box 7.3.3.[5]

Let's start by focusing on a single round of the purified BB'84 protocol. Imagine that Eve is trying to defeat the protocol: her task is to prepare a state ρ_{ABE} such that each of A and B is a single qubit, and E can be arbitrary. Moreover, she would like to achieve two properties. First of all, if Alice and Bob decide to measure their qubit in the same basis and compare their outcomes then they should be identical as often as possible (as otherwise the users will detect this and abort the protocol). Let's introduce random variables (Z_A, Z_B) to represent the result of measuring both qubits A and B in the standard basis, and (X_A, X_B) to represent the result of measuring both qubits in the Hadamard basis. These random variables are not all simultaneously well-defined because we can't measure the qubits in both bases, but

[5] In the literature on QKD, "making the i.i.d. assumption" is synonymous with "showing security under collective attacks."

each pair by itself is well-defined. Let's define

$$\delta = \frac{1}{2}\Pr(X_A \neq X_B) + \frac{1}{2}\Pr(Z_A \neq Z_B)\ , \tag{7.8}$$

so that δ is the probability to obtain different outcomes when the choice of the basis is made uniformly at random. Eve's first goal is to minimize this quantity. Second, Eve would like to be able to predict Alice's outcome, again in a random choice of basis (that she gets to learn). So, she would like to minimize the quantity

$$\frac{1}{2}H_{\min}(X_A|E) + \frac{1}{2}H_{\min}(Z_A|E)\ . \tag{7.9}$$

Furthermore, recall the interpretation of the conditional min-entropy as a guessing probability: this is precisely (minus the logarithm of) Eve's maximum chance of guessing X_A, when the chosen basis is the standard basis, or Z_A, when it is the Hadamard basis.

Wen need to show that Eve's goal is impossible: she can make one quantity or the other small, but not both at the same time. This is the same goal as in the previous section, except that in the previous section we combined both goals in a single one to obtain a simple game that we could analyze.

As we will see later, due to the i.i.d. assumption it turns out to be sufficient to consider analogues of (7.8) and (7.9) where uncertainty is measured using the von Neumann entropy, as opposed to the min-entropy. Concretely, we replace the success measure (7.8) by

$$\frac{1}{2}H(X_A|X_B) + \frac{1}{2}H(Z_A|Z_B)\ . \tag{7.10}$$

This measure is directly related to (7.8). Using a simple calculation based on the definition of the von Neumann entropy we can check that

$$\frac{1}{2}H(X_A|X_B) + \frac{1}{2}H(Z_A|Z_B) \ \leq\ h(\delta_{\max})\ , \tag{7.11}$$

where h is the binary entropy function and $1-\delta_{\max}$ is the probability of succeeding in the matching outcomes test using the state ρ_{AB}. This means that the users, who compute δ in the protocol and verify that $\delta \leq \delta_{\max}$, can thereby verify that (7.11) holds. Similarly, we replace (7.9) by

$$\frac{1}{2}H(X_A|E) + \frac{1}{2}H(Z_A|E)\ . \tag{7.12}$$

To bound the average of the two quantities (7.10) and (7.12) we use the following inequality which is an example of an *entropic uncertainty relation*. This relation states that for *any* state ρ_{ABE} such that each of A and B are a system of a single qubit we always have that

$$H(X_A|X_B) + H(Z_A|E) \geq 1 \qquad \text{and} \qquad H(Z_A|Z_B) + H(X_A|E) \geq 1\ . \tag{7.13}$$

In other words, it is impossible to have both $H(X_A|X_B)$ small *and* $H(Z_A|E)$ small! This is a manifestation of the *monogamy of entanglement*: If the state ρ_{ABE} is strongly correlated in the standard basis across A and B, then it must be unpredictable in the Hadamard basis across A and E – and vice-versa. Averaging these two inequalities we obtain

$$\Big(\frac{1}{2}H(X_A|X_B) + \frac{1}{2}H(Z_A|Z_B)\Big) + \Big(\frac{1}{2}H(X_A|E) + \frac{1}{2}H(Z_A|E)\Big) \geq 1\ . \tag{7.14}$$

Using (7.11) in (7.14) we get that

$$\frac{1}{2}H(X_A|E) + \frac{1}{2}H(Z_A|E) \geq 1 - h(\delta_{\max})\ . \tag{7.15}$$

This relation looks very similar to the equation (7.4) that we want to prove, but there are some differences. First of all, the measure of entropy is not exactly the same. Here, we have the conditional von Neumann

entropy, and in (7.4) we have the conditional min-entropy. Second, (7.2) applies simultaneously to all outcomes that were not measured, but (7.15) applies to a single round. Finally, in (7.4) the min-entropy is evaluated on the state conditioned on not aborting, whereas here we have not yet taken this conditioning into account.

The first difference is handled by making the i.i.d. assumption. The key leverage that we get from this assumption is that it allows us to use the *quantum asymptotic equipartition property*. This states that, when considering a large number of samples of a random variable X, the min-entropy converges to the von-Neumann entropy:

$$\frac{1}{n}\mathrm{H}^{\varepsilon}_{\min}(X_1 \cdots X_n) \approx_{n\to\infty} H(X)$$

for i.i.d. X, provided the smoothing parameter ε is chosen sufficiently large. (Informally, the smoothing parameter ε on the left-hand side means that when calculating the min-entropy we take the largest possible value among all distributions that are ε-close to $(X_1, \ldots, X_n)$ in total variation distance; see Box **??**.) Here, we can use this property to argue that $H_{\min}(X_A|E\Theta) \approx nH(X_{A,1}|E\Theta)$, where $X_{A,1}$ is the first bit of the n-bit string X_A.

The i.i.d. assumption lets us easily the second difference as well. Finally, the conditioning requires a little care, but we omit the technical details—let's just say that the result is similar to the dependence on $p(\Omega)$ which we already observed in the previous section.

To conclude we note that in general of course the i.i.d. assumption cannot be experimentally justified: in practice, the eavesdropper can do what they want. Hence security "under the i.i.d. assumption" is, arguably, not security at all, and this is the main limitation of our work in this section. In fact, in the most general case, without making the i.i.d. assumption, it is still possible to show (7.2) using a more involved uncertainty relation, that is shown directly for the min-entropy and therefore bypasses the need for the asymptotic equipartition property. If you are interested, we give a pointer in the chapter notes. i.i.d. assumption notwithstanding, as a result of all our hard work we managed to prove (7.2), with a coefficient almost 1 in front of the n! This is the best that we could hope for.

7.4 Correctness of BB'84 key distribution

We end by formally arguing correctness of the protocol. In case we make the i.i.d. assumption then correctness is very easy to show, as we already did informally: in that case the probability of succeeding in the matching outcomes test gives an estimate for the number of positions in which the strings x_{remain} and $\tilde{x}_{\mathrm{remain}}$ are expected to differ, and to get correctness it suffices to select an appropriately good information reconciliation protocol. Formally, we'd use a Chernoff bound — since we're about to give a more general argument we skip the details.

Showing correctness without making the i.i.d. assumption requires more work. The main issue is that in Protocol 3 the matching outcomes test in step 7 is performed on the rounds T selected for testing, but for the raw key we use the rounds in $R = S\backslash T$: how can we guarantee that information gathered through the tests performed on rounds in T has some implication for the rounds *not* in T? Intuitively this is because the tested rounds are chosen at random, and moreover they are chosen *after* the adversary has created the state ρ_{ABE}. So, even if ρ_{ABE} can be arbitrarily correlated there should be no way to arrange things so that tests in a randomly chosen subset of rounds pass and yet the untested rounds wouldn't have passed.

Let's explain how we can make this intuition precise. 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 $Z_j = 0$ indicates success 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 $\sum_{j \in T} Z_j \leq \delta |T|$. To select parameters for the information reconciliation protocol, however, they would like to have a bound on $\sum_{j \in S \setminus T} Z_j$ that they can be confident about. 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 available. 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 7.4.1 *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 $\delta, \nu > 0$,*

$$\Pr\Big(\sum_{j \in T} X_j \leq \delta k \ \ \wedge \sum_{j \in \{1,\ldots,m\} \setminus T} X_j \geq (\delta + \nu) n \Big) \leq e^{-2\nu^2 \frac{nk^2}{(n+k)(k+1)}}. \tag{7.16}$$

To see what the theorem says in our setting, set $m = |S| = 2n$ and $k = n$.[6] Let's also choose $\nu = \delta$ for convenience. Plugging in these parameters we get the bound

$$\Pr\Big(\sum_{j \in T} Z_j \leq \delta n \ \ \wedge \sum_{j \in S \setminus T} Z_j \geq 2\delta n \Big) \leq e^{-\delta^2 \frac{n^2}{n+1}}, \tag{7.17}$$

which is valid for any choice of $\delta > 0$. Eq. (7.17) implies that the probability that the test performed in step 8 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 8 of Protocol 3, we can use Bayes' rule to rewrite the bound above as

$$\Pr\Big(\sum_{j \in S \setminus T} Z_j \geq 2\delta n \ \Big| \ \neg \, \mathrm{abort} \Big) \leq \frac{e^{-\delta^2 \frac{n^2}{n+1}}}{\Pr(\neg \, \mathrm{abort})}. \tag{7.18}$$

Writing the bound in this way allows us to clarify our earlier discussion around the role of the probability of aborting. As you can see, the bound (7.18) is only good if $\Pr(\neg \, \mathrm{abort})$ is not too small; if this probability was extremely tiny, then the right-hand side of Eq. (7.18) 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.

To conclude, assuming that we choose the parameters of information reconciliation such that strings at a relative distance at most 2δ are corrected with probability at least ε_{IR} (where ε_{IR} is a correctness parameter for information reconciliation) it follows that the QKD protocol is $\varepsilon_c = eps_{IR} + \varepsilon_a$-secure, where $\varepsilon_a = e^{-\delta^2 n} / \Pr(\neg \, \mathrm{abort})$ is the right-hand side of (7.18). Note that here 2δ can be made arbitrarily close to δ by choosing as small a ν as we like (and paying a corresponding increase in the error term ε_a).

[6] There is a subtlety here, which is that only the *expected* size of T is n, but the size of T may vary from one execution of the protocol to another. We gloss over this issue here, and return to it in Section **??** in the next chapter.

Moreover, to be precise we should note that small additional error terms should be included to account for our assumption that the test set has size precisely $|T| = n$; we will see how to deal with this in the next chapter (spoiler: it is easier than what we just did, and leads to smaller errors, which justifies us neglecting this minor point so far).

7.5 Chapter notes

The origins of QKD can be traced back to ideas that Wiesner had in the 1970s [**?**]. The first concrete proposal for a QKD protocol is due to Bennett and Brassard [**?**]. Shortly after, Ekert discovered a different approach [**?**] which we review in the next chapter. Up to small variations these are the two main QKD protocols studied, and implemented, to date.

The uncertainty relation in Section 7.3.3 is from [**?**]. The quantum asymptotic equipartition property is shown in [**?**]. For a complete security proof based on the tripartite guessing game, see [**?**]. For a complete proof of security based on entropic uncertainty relations, see [**?**]. For the general non-i.i.d. case, one can use the uncertainty relations from [**?**], extended as in [**?**, Corollary 7.4]. Another method, less strong quantitatively but conceptually elegant, to reduce the analysis of a multi-round protocol to the i.i.d. case is to use "de Finetti reductions," see for example [**?**].

In this chapter the analysis remains high-level and focuses on the asymptotic setting, where we can assume that the number of rounds N of the protocol goes to infinity, in practice it is crucial to understand the error terms even for moderately small values of N. For this, see e.g. [**?**].