QUANTUM INFORMATION & COMPUTATION

Nilanjana Datta, DAMTP Cambridge

1 Quantum cryptography: BB84 quantum key distribution

Introduction

The use of general quantum states to represent information and encode messages would appear, from what we have seen, to have some significant drawbacks compared to the use of classical states! Firstly (a): a received unknown quantum state cannot be reliably identified (unless it has been promised to belong to a specified orthonormal set) so the receiver cannot reliably read the message. Secondly (b): any attempt to read the message (in the context of general signal states) results in only partial information and is always accompanied by irrevocable corruption ("measurement collapse") of the quantum state and correspondingly, part of the sent message is necessarily irretrievably destroyed.

Remarkably these seemingly negative features can be used to positive effect to provide valuable benefits for some cryptographic and information security issues, which in some cases turn out to be impossible to achieve with classical signals. For example, intuitively here, in communication between distant parties, (b) implies that any attempted eavesdropping on the message en route must leave a mark on the signal which can then in principle be detected by actions of receiver in (public) discussion with the sender. It turns out that this can be used to provide communication that is provably secure against eavesdropping. Classical messages on the other hand can always be read en route and sent on to the receiver perfectly intact. Also it turns out (cf more below) that the effects of (a) for the communicators can be circumvented by a suitably clever (non-obvious) protocol involving further (public) discussion between them.

It is now known that quantum effects can provide benefits for a wide variety of cryptographic tasks beyond just secure communication. The associated subject of *quantum cryptography* is currently a highly active and flourishing area of scientific research worldwide, with evident huge practical and theoretical significance for modern society which is becoming increasingly reliant on secure information technology. In this course we will consider only one cryptographic issue, but perhaps the most fundamental of all – that of secure communication between two spacially separated parties traditionally named Alice (A) and Bob (B) with eavesdropper Eve (E). We will discuss the Bennett-Brassard protocol (the so-called BB84 protocol) for *quantum key distribution* which provides a means of communication that is provably secure against eavesdropping. The protocol itself is relatively simple to describe but the proof of unconditional security against general attacks is very involved and technical (beyond the scope of this course). Below we will describe the protocol and be content with making some remarks about its security in

some restricted situations.

How can we communicate securely?
The issue of secure communication has a long history going back thousands of years. Circa 100 BC Julius Caesar used a cipher in which the letters of the text were simply shifted forward by three places in the alphabet. A more elaborate version of this kind of encryption method (subsequently historically used in a variety of contexts) is to apply some more general chosen fixed permutation of the alphabet, known securely only to the sender and receiver. However such schemes are insecure (against suitably intelligent adversaries) for example by compiling a table of symbols with their occurring frequencies, and comparing this to a similar table derived from typical texts in the language.

With the development of mathematics (particularly number theory, abstract algebra, group theory, coding theory, computational complexity theory etc.) a variety of more sophisticated (classical) schemes for secure communication were invented but none of these apart from the *one time pad* (which in turn requires a method of *key distribution*) is provably secure. We will discuss the one time pad below as it is also an underpinning ingredient for quantum key distribution (QKD) schemes such as BB84. QKD schemes will be able to circumvent shortcomings of classical key distribution schemes which render them unsuitable in many common situations (cf below).

Remark (on public key cryptosystems).
Amongst more sophisticated schemes in common use today are the so-called public key crypto systems (Diffie-Hellman scheme, RSA, elliptic curve cryptography). The security of these schemes is not absolute but relies on (unproven but widely believed) computational hardness assumptions i.e. a belief that certain computational tasks while computable in principle require so much computing time that they are effectively uncomputable in practice. For example given two large prime numbers p and q (of say hundreds of digits each) it is easy to compute their product (e.g. using a very large sheet of paper and careful long multiplication we could even probably do it by hand over a rainy weekend). But conversely, given a composite number N (similarly having hundreds of digits) there is no known "fast algorithm" to factorise it. In fact for N having just several hundreds of digits, and using our best known classical factoring algorithms with all the classical computing power on earth today, it would generally take longer than the age of the universe to complete the task! These kinds of issues are the subject of *computational complexity theory* which we will see more of in the second half of the course (quantum computation and quantum algorithms). More precisely here, the computational task of multiplication of n-digit integers can be completed in a number of steps growing polynomially (quadratically) with n, whereas our best classical factoring algorithms for n digit integer factorisation require an exponential (super-polynomial) number of steps, exponential in the cube root of n, and this exponential versus polynomial growth in number of digits makes that latter task effectively uncomputable in practice for modest sizes of n. Public key cryptosystems exploit such asymmetric (assumed) computational hardness properties of various tasks to provide security. They also have the remarkable very useful feature that the communicating parties do not need any prior secret shared information known only to them (such as knowledge of the permutation in a Caesar-like cypher) so, for example, they need never have previously met. However there are significant draw-

backs:
(a) it has not been proven that faster classical algorithms for the tasks may be discovered in the future e.g. a factoring algorithm that requires only a polynomially growing number of steps to complete.
(b) *quantum* computation provides entirely new (non-classical) modes of computing consistent with the laws of physics (as we will see in the second half of the course). These modes lead to new kinds of algorithms which can be used to solve some computational problems exponentially faster than any known classical method. And coincidentally, known tasks of this kind include those on which public key cryptosystems are based. So public key crypto systems in common use today could be readily broken if we had a working quantum computer! The most famous such algorithm is Shor's quantum algorithm for integer factorisation (and also computation of discrete logarithms in number theory) discovered by Peter Shor in 1994, which achieves these tasks in a number of (quantum-) computational steps growing only polynomially with n i.e rendering them feasible in practice.

Thus on the one hand quantum physics (via Shor's and other quantum algorithms) allows the breaking of some classical cryptosystems that are not known to be classically breakable, while on the other hand, via quantum key distribution protocols, it offers a method for provably secure communication.

The one time pad

We assume that our message is a bit string M of length n (without loss of generality e.g. we could represent letters of the alphabet and some punctuation symbols as distinct 5-bit strings).

For the one time pad Alice and Bob need to share a *secret private key* K which is a uniformly random bit string of the same length n as the message, and which is known only to them.

Alice encrypts her message by adding K to M. Here addition is addition mod 2 and it is carried out separately at each bit position of the strings. This produces the cryptotext $C = M \oplus K$ which she sends to Bob (over a public classical channel).

Bob receives C and computes $C \oplus K = M \oplus K \oplus K = M$ (with the last equality since $0 \oplus 0 = 1 \oplus 1 = 0$) thus decrypting the message.

Features and remarks

If K is uniformly distributed amongst bit strings then so is C. Thus any potential eavesdropper Eve can learn nothing about M (apart from its length) by looking at C. Hence this scheme cannot be broken, a feature that can be proven more formally in the context of classical information theory, introduced by Claude Shannon in the 1940s.

It is important for security that the key K is used only once (hence the name "one time pad") e.g. if it would be used twice to generate $C_1 = M_1 \oplus K$ and $C_2 = M_2 \oplus K$ then $C_1 \oplus C_2 = M_1 \oplus M_2$ which would generally contain information about M_1 and M_2 and (with C_1 and C_2 available) about K too, if Alice were to use K again.

Thus the scheme is rather inefficient in its ongoing needed secret resource, with Alice and Bob needing fresh secret key of length equal to that of each subsequent message. But given that, all seems fine and the key question is: how can Alice and Bob acquire their secret key? It is impossible for two parties to classically generate a secure private key over a public channel. Thus they would need to meet (and carry away e.g. a private one time pad book, for later use) or else use a trusted intermediary to distribute the key. Each of these has evident limitations and potential significant security risks. Quantum key distribution (QKD) provides a method for Alice and Bob to generate a shared secret key over public classical and quantum channels without the need to meet or to use a trusted intermediary.

In QKD schemes the quantum signals are used to generate the shared secret key rather than to encode the message itself, which is subsequently communicated using the classical one time pad scheme. A variety of quantum key distribution schemes have been proposed including:
(1) BB84 (C. Bennett and G. Brassard 1984), uses four qubit signal states that include non-orthogonal pairs;
(2) B92 (C. Bennett 1992), uses only two (non-orthogonal) quantum signal states;
(3) E91 (A. Ekert 1991), uses (one qubt of) an entangled pair of qubits in place of the signal states of BB84 which are later created using local measurements by Alice and Bob; and others. We will discuss only BB84.

The BB84 quantum key distribution protocol

We assume that Alice and Bob can communicate over a public classical channel and they can also send qubits over a quantum channel. Eve also has access to these channels and she wants to acquire information without being detected, about the secret key that Alice and Bob will generate. The bottom line will be that this will be impossible for Eve to achieve by any means whatsoever consistent with the laws of physics.

For quantum transmissions we will use the following four qubit states

$$\begin{aligned}
|\psi_{00}\rangle &= |0\rangle \\
|\psi_{10}\rangle &= |1\rangle \\
|\psi_{01}\rangle &= |+\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \\
|\psi_{11}\rangle &= |-\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle)
\end{aligned}$$

making up two orthonormal qubit bases viz. $\mathcal{B}_0 = \{|\psi_{00}\rangle, |\psi_{10}\rangle\}$ and $\mathcal{B}_1 = \{|\psi_{01}\rangle, |\psi_{11}\rangle\}$. These are the computational basis (or Pauli Z eigenbasis), and the diagonal (or conjugate) basis (or Pauli X eigenbasis). These bases are called *mutually unbiased* since if any state of one basis is measured in the other basis, the outcomes are always equally likely.

We now give the BB84 protocol as a series of steps:

BB84 Step1:

Alice generates two uniformly random binary strings $\mathbf{x} = x_1 x_2 \dots x_m$ and $\mathbf{y} = y_1 y_2 \dots y_m$, $x_i, y_i \in \{0, 1\}$. Then she prepares m qubits in the states

$$|\psi_{x_1 y_1}\rangle \, |\psi_{x_2 y_2}\rangle \dots |\psi_{x_m y_m}\rangle$$

and sends these m qubits over to Bob.
Here x_i will represent the bit value she is trying to send and y_i is her choice of quantum encoding (choice of basis) for that bit. Using such a random choice of mutually unbiased bases for encoding each bit value is sometimes called conjugate coding.

BB84 Step 2:

When Bob receives the m qubits they may no longer be in the states $|\psi_{x_i y_i}\rangle$ that Alice sent, since the quantum channel may have been noisy or eavesdropping may have occurred. To understand how the protocol works let us imagine first that there is no eavesdropping and that the channel is perfectly noiseless i.e. Bob receives precisely the states $|\psi_{x_i y_i}\rangle$ that Alice sent.

Bob chooses a uniformly random bit string $\mathbf{y}' = y'_1 y'_2 \dots y'_m$ and measures the i^{th} received qubit in basis $\mathcal{B}_{y'_i}$ to get a result x'_i i.e. y'_i is Bob's guess at Alice's choice of encoding basis and x'_i is his guess at Alice's bit value x_i. Let $\mathbf{x}' = x'_1 x'_2 \dots x'_m$ be the string of Bob's measurement outcomes. Note that if $y'_i = y_i$ (i.e. Bob correctly guessed Alice's encoding basis) then $x'_i = x_i$ and he will learn her message bit correctly with certainty. But if $y'_i \neq y_i$ then x'_i is completely uncorrelated with x_i (recalling the mutually unbiased relationship between the bases).

BB84 Step 3:

After the completion of Step 2 Alice and Bob publicly reveal and compare their choice of bases i.e. their strings y and y' (but they do not reveal the strings $\mathbf{x}$ and $\mathbf{x}'$!). They discard all bits x_i and x'_i for which $y_i \neq y'_i$ leaving shorter strings of expected length $m/2$. Call these strings $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}'}$. Under our assumptions of no noise and no eavesdropping in the quantum channel, these bit strings would provide the desired outcome of a shared secret key.

[An example of Steps 1 to 3 given in the lecture.]

In reality there will always be some noise in transmissions and we need to deal with possible eavesdropping too. To address these issues the BB84 protocol concludes with the following steps 4 and 5, which we discuss further below. These are entirely issues and techniques from classical cryptography.

BB84 Step 4 (information reconciliation):

Alice and Bob want next to estimate the *bit error rate (BER) i.e. the proportion of bits in* $\tilde{\mathbf{x}'}$ *that are not equal to those in* $\tilde{\mathbf{x}}$. To do this they publicly compare a random sample of their strings (say half of their bits chosen at random positions), and then discard all the announced bits. They assume that the remaining bits have about the same proportion of errors as those checked. Next they want to correct these remaining errors (albeit at unknown positions) to obtain two strings that agree in a high percentage of positions with high probability. Remarkably this can be done (at the expense of sacrificing some more bits) without giving everything away, if the bit error rate is not too large.

BB84 Step 5 (privacy amplification):

From the estimated bit error rate Alice and Bob can estimate the maximum amount of information that an eavesdropper is likely to have obtained about the remaining bits. From this information estimate they use techniques of so-called privacy amplification from classical cryptography to replace their strings by even shorter strings about which the eavesdropper can have practically no knowledge whatever (with high probability). This concludes the BB84 quantum key distribution protocol.

Further remarks about information reconciliation and privacy amplification

A rigorous treatment of the details of Steps 4 and 5 requires much further technical development from classical information theory, the theory of error correcting codes and classical cryptography. A full treatment is beyond the scope of this course and here we will only draw attention to some of the essential ideas.

In Step 5 the bit error rate provides an upper bound on the amount of information that an eavesdropper can have gained because (as we have previously discussed) non-orthogonal states cannot be reliably distinguished and any attempt to acquire information about the state identity certainly involves irreversible state disturbance. In the full theory one can prove information disturbance tradeoff relations that quantify the intuition that more information gain is necessarily accompanied by more disturbance. As a consequence of this fundamental property of quantum information, the amount of Eve's acquired information is reflected in the bit error rate. Of course noise in the channel also generates bit errors but Alice and Bob can reliably upper bound Eve's information by simply assuming that the whole error rate arose from eavesdropping.

There are many ways in which Eve could attempt to acquire information, such as:
(a) the intercept-resend attack: Eve can intercept each transmitted qubit separately, measure it in some chosen basis to acquire some information about it, and then send on the post-measurement state to Bob.
(b) general coherent attack: (could be much more general and complicated) Eve can introduce an auxiliary (possibly very large) probe quantum system E of her own and unitarily interact E with many of the passing qubits. Finally she can measure E to acquire information, which now can be joint information about many of the qubits. Her measurement here can even be postponed until after she overhears Alice and Bob's public discussions in Steps 2,3,4,5 and chosen in response to what she hears.
(c) Note that the standard classical strategy for eavesdropping on classical bits viz. reading them and retaining a copy, and then sending them on perfectly intact, is not available for our quantum state bit encodings because of the no-cloning theorem and the use of non-orthogonal states in the set of encoding states!

Remarkably the privacy amplification techniques of classical information theory can be shown to provide security against any possible eavesdropping strategy that is consistent with the laws of physics.

Example. (an intercept-resend attack)
Assume that the quantum channel is noiseless but Eve intercepts each passing qubit and

measures it in the so-called Breidbart basis:

$$\begin{aligned} |\alpha_0\rangle &= \cos\tfrac{\pi}{8}\,|0\rangle + \sin\tfrac{\pi}{8}\,|1\rangle \\ |\alpha_1\rangle &= -\sin\tfrac{\pi}{8}\,|0\rangle + \cos\tfrac{\pi}{8}\,|1\rangle\,. \end{aligned}$$

This is a good choice of basis as it lies "midway" between the two BB84 encoding bases: the squared overlaps of $|\alpha_0\rangle$ with the two states $|0\rangle$ and $|+\rangle$ used to encode bit value 0 are equal (being $\cos^2\pi/8$) and similarly for $|\alpha_1\rangle$ with $|1\rangle$ and $|-\rangle$. For any other choice of basis one of these four overlaps will be smaller and intuitively the $|\alpha_i\rangle$'s thus provide the best (most parallel) simultaneous approximations to the two non-orthogonal states used to encode each bit value i; and Eve will learn each bit of $\tilde{X}$ with probability $\cos^2\pi/8 \approx 0.85$. *Detailed discussion in class.*

Let us compute the *bit error rate* (BER) in the strings $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{x}'}$ arising from Eve's intervention (supposing that Eve measured the intercepted state relative to the Breidbart basis. For each of the four encoding states (used with probability 1/4) we compute $\mathrm{Prob}(x \neq x')$. Let us denote measurement outcomes by the corresponding basis states and write for example, P(B infers 1 | he received $|\phi_0\rangle$) to denote the probability that Bob's measurement result is 1 given that he received a qubit in state $|\phi_0\rangle$ etc.
For state $|0\rangle$ sent by Alice in basis $\mathcal{B}_0$ to encode bit value 0, we know that Bob will measure in basis $\mathcal{B}_0$ (i.e. same bases for the tilde strings) but in between Eve will have measured in the Breidbart basis. Thus for this case

$$\begin{aligned} \mathrm{Prob}(x' \neq x) &= \mathrm{P}(\text{B infers 1} \mid \text{A sent } |0\rangle) \\ &= \mathrm{P}(\text{E sends } |\phi_0\rangle \mid \text{A sent } |0\rangle)\cdot\ \mathrm{P}(\text{B infers 1} \mid \text{he received } |\phi_0\rangle) \\ &\quad + \mathrm{P}(\text{E sends } |\phi_1\rangle \mid \text{A sent } |0\rangle)\cdot\ \mathrm{P}(\text{B infers 1} \mid \text{he received } |\phi_1\rangle) \\ &= |\langle\phi_0|0\rangle|^2 |\langle 1|\phi_0\rangle|^2 + |\langle\phi_1|0\rangle|^2 |\langle 1|\phi_1\rangle|^2 \\ &= \cos^2\pi/8 \sin^2\pi/8 + \sin^2\pi/8 \cos^2\pi/8 \\ &= \tfrac{1}{4}. \end{aligned}$$

The other three encoding states similarly give the same result. Thus the eavesdropping will result in a disturbance amounting to a bit error rate of 25%. This is in fact the minimal value over all choices of Eve's basis (cf Exercise Sheet 2). □

Having estimated the bit error rate in step 4, Alice and Bob now perform information reconciliation to correct the errors (albeit in unknown positions!) in their remaining strings. This can be achieved using techniques from the theory of error correcting codes (that we will not discuss in this course) or other methods from classical cryptography.

Information reconciliation leaks further limited information to Eve and the final result is a pair of shorter strings that are guaranteed to be equal in each position with high probability, but about which Eve may still have information. Finally privacy amplification is performed to produce two final strings of still shorter length, about which Eve is guaranteed with high probability to have no significant information at all.

Eve may have different kinds of information about the string. For example she may know some specific bits, or parities of some subsets of bits, or some other Boolean function of bits, or perhaps probabilistic information e.g. that a particular bit or Boolean function has probability 2/3 of being 0 etc. It is thus prima facie remarkable that privacy amplification can be done at all, without publicly revealing the whole string. Here's an example to illustrate how it can be achieved in a very simple case.

We have here really just drawn attention to the ideas of information reconciliation and privacy amplification, and their significance for the BB84 QKD protocol. Readers interested in more details should consult the (extensive) classical cryptography literature.

(Q) Why should Eve choose the Breidbart basis?

Discussion in class.