Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 146 (2005)

Quantum Cryptography: A Survey *

Dagmar Bruß,† Gábor Erdélyi,‡ Tim Meyer,† Tobias Riege,‡ and Jörg Rothe‡
Heinrich-Heine-Universität Düsseldorf
40225 Düsseldorf, Germany

March 13, 2006

Abstract

We survey some results in quantum cryptography. After a brief introduction to classical cryptography, we provide the quantum-mechanical background needed to present some fundamental protocols from quantum cryptography. In particular, we review quantum key distribution via the BB84 protocol and its security proof, as well as the related quantum bit commitment protocol and its proof of insecurity.

Categories and subject descriptors: E.3; E.4; F.1; F.2.2; J.2.
General Terms: Theory; Security; Algorithms; Experimentation.
Key words and phrases: quantum cryptography; quantum key distribution; quantum bit commitment.

1 Introduction

Cryptography is the science of keeping private information from unauthorized access. An algorithm, which is called a *cipher* in this context, scrambles the message via some rule such that restoring the original message is hard—if not impossible—without knowledge of the secret key. Cryptographic technology in use today relies on the hardness of certain mathematical problems. Classical cryptography faces the following two problems. First, the hardness of the problems on which the security of cryptosystems is based (e.g., integer factoring or the discrete logarithm problem) often is not a proven fact but rather a widely believed hypothesis. Second, the theory of quantum computation has yielded new methods to tackle these mathematical problems in a much more efficient way. Although there are still numerous challenges to overcome before a working quantum computer of sufficient power can be built, in theory all classical ciphers might be vulnerable to such a powerful machine. However, while quantum computation seems to be the coffin nail for classical cryptography in a possibly not so distant future, at the same time it offers new possibilities to build encryption methods that are safe even against attacks performed by means of a quantum computer. Quantum cryptography extends the power of classical cryptography by protecting the secrecy of messages using the physical laws of quantum mechanics.

*This work was supported in part by the DFG under Grants RO 1202/9-1 and RO 1202/9-3, and by the EU Integrated Project SECOQC. A preliminary version was presented at *A Magyar Tudomany Napja*, Eötvös József Föiskola, Baja, Hungary, in November 2005.

†Institut für Theoretische Physik. Email: {bruss, meyer}@thphy.uni-duesseldorf.de.
‡Institut für Informatik. Email: {erdelyi, riege, rothe}@cs.uni-duesseldorf.de.

ISSN 1433-8092

Looking back in the history of cryptography, one of the first encryption methods was the scytale. The first recorded use of the scytale dates back to the fifth century B.C. when the Spartans used it to exchange battle information between generals without revealing it to the enemy. To encrypt a message, which we call the *plaintext* in cryptography, a strip of leather or pergament was wrapped around a wooden cylinder, the scytale. The encrypted message, also called the *ciphertext*, was then written from left to right onto the leather, so that unravelling the strip would produce a meaningless alignment of seemingly random letters, see Figure 1 for the encryption of the plaintext "scytaleisatranspositioncipher" by "ssoicaspytihtrteaairlnoesnipc." The decryption of the ciphertext was achieved by using a scytale of the same diameter as the cylinder that was used for encryption.

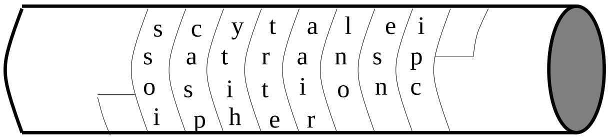

Figure 1: The Scytale

The scytale is a so-called transposition cipher, since only the order of the letters within the message is changed. Another type of encryption is the substitution cipher. Here, instead of swapping the positions of the letters, each plaintext letter is replaced by another letter according to some specific rule.

The method of encryption and decryption is called a *cryptosystem*, whereas the particular information used for encryption or decryption in an individual communication is called a *key*. In the case of the scytale, the diameter of the cylinder represents the secret key. Obviously, this ancient cryptosystem has a very low level of security. Once the method of encryption is known to the eavesdropper, he or she can simply try all possible diameters to reveal the original message. The fact that the cryptosystem is publicly known is not the reason for the insecurity of the communication, but rather the small number of possible keys that can be used for encryption. In the 19th century, Auguste Kerckhoffs stated the principle that the security of a cryptosystem must be based solely on the secrecy of the key itself. Therefore, when designing new ciphers, one should always treat the algorithm as if it were publicly known.

Over time, the amount of information that needed to be encrypted exploded, making it impossible to use simple and insecure procedures like the scytale. At first, mechanical devices were built to speed up the encryption and the decryption process, and to increase the complexity of the keys used to scramble the message. An infamous example of such a mechanical cryptosystem is the Enigma, which was used in World War II by the Germans to conceal their military communication. Not following Kerckhoffs' principle, the Germans considered the Enigma unbreakable, assuming that the mechanical device used for secure communication was not known to the enemy. However, allied cryptanalysts in Bletchley Park near London often were able to decrypt the German military messages during the war. One might argue that breaking the Enigma was one of the most crucial factors for the victory of the allied forces and for ending the war. After the war, it was the invention of the transistor that made the rise of the computer industry possible.

The huge speed-up in executing mathematical calculations resulted in the need to create much more secure cryptosystems, among them symmetric block ciphers such as the Data Encryption Standard (DES) and the Advanced Encryption Standard (AES) and public-key cryptosystems such as RSA and others, which are integrated in modern cryptographic applications currently in use. A nice and easy-to-read overview of the history of cryptography is given by Singh [Sin99].

With the currently emerging theory of quantum computation, we seem to be at the beginning of yet another era of cryptography.

2 Classical Cryptography

Overviews of classical cryptography can be found in various text books, see, e.g., [Rot05,Sti05]. Here, we present just the basic definition of a cryptosystem and give one example of a classical encryption method, the one-time pad.

Definition 1 *A* (symmetric) cryptosystem *is a five-tuple* $(\mathcal{P}, \mathcal{C}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ *satisfying the following conditions:*

1. $\mathcal{P}$ *is a finite set of possible* plaintexts*.*
2. $\mathcal{C}$ *is a finite set of possible* ciphertexts*.*
3. $\mathcal{K}$ *is a finite set of possible* keys*.*
4. *For each* $k \in \mathcal{K}$*, there are an* encryption rule $e_k \in \mathcal{E}$ *and a corresponding* decryption rule $d_k \in \mathcal{D}$*, where* $e_k : \mathcal{P} \to \mathcal{C}$ *and* $d_k : \mathcal{C} \to \mathcal{P}$ *are functions satisfying* $d_k(e_k(x)) = x$ *for each plaintext element* $x \in \mathcal{P}$*.*

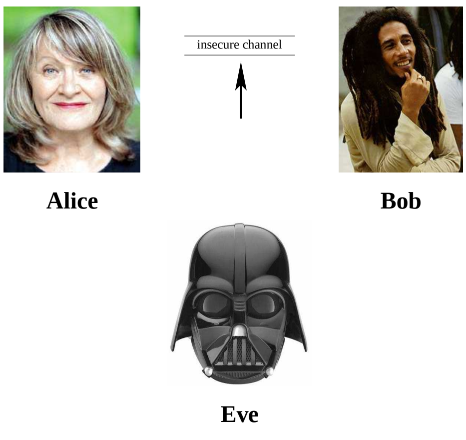

Figure 2: Communication between Alice and Bob, with Eve listening

In the basic scenario in cryptography, we have two parties who wish to communicate over an insecure channel, such as a phone line or a computer network. Usually, these parties are referred to as Alice and Bob. Since the communication channel is insecure, an eavesdropper, called Eve, may intercept the messages that are sent over this channel. By agreeing on a secret key k via a secure communication method, Alice and Bob can make use of a cryptosystem to keep their information secret, even when sent over the insecure channel. This situation is illustrated in Figure 2.

The method of encryption works as follows. For her secret message m, Alice uses the key k and the encryption rule e_k to obtain the ciphertext $c = e_k(m)$. She sends Bob the ciphertext c over the insecure channel. Knowing the key k, Bob can easily decrypt the ciphertext by the decryption rule d_k:

$$d_k(c) = d_k(e_k(m)) = m.$$

Knowing the ciphertext c but missing the key k, there is no easy way for Eve to determine the original message m.

There exist many cryptosystems in modern cryptography to transmit secret messages. An early well-known system is the *one-time pad*, which is also known as the *Vernam cipher*. The one-time pad is a substitution cipher. Despite its advantageous properties, which we will discuss later on, the one-time pad's drawback is the costly effort needed to transmit and store the secret keys.

Example 2 (One-Time Pad) *For plaintext elements in $\mathcal{P}$, we use capital letters and some punctuation marks, which we encode as numbers ranging from 0 to 29, see Figure 3. As is the case with most cryptosystems, the ciphertext space equals the plaintext space. Furthermore, the key space $\mathcal{K}$*

A	B	C	D	E	$\cdots$	X	Y	Z		!	-	.
00	01	02	03	04	$\cdots$	23	24	25	26	27	28	29

Figure 3: Letters and punctuation marks encoded by numbers from 0 to 29

also equals $\mathcal{P}$, and we have $\mathcal{P} = \mathcal{C} = \mathcal{K} = \{0, 1, \ldots, 29\}$.

*Next, we describe how Alice and Bob use the one-time pad to transmit their messages. One concrete example is shown in Figure 4. Let $m = m_1 m_2 \ldots m_n$ be a given message of length n, which Alice wishes to encrypt. For each plaintext element m_i, where $1 \leq i \leq n$, Alice randomly and uniformly chooses a key element $k_i \in \{0, 1, \ldots, 29\}$ and adds the plaintext numbers to the key numbers. The result is taken modulo 30. For example, the last letter of the plaintext from Figure 4, "*D*," is encoded by "*03*." The corresponding key is "*28*," so we have $c = 3 + 28 = 31$. Since $31 \equiv 1 \bmod 30$, our plaintext letter "*D*" is decrypted as "*B*." The encryption and decryption can be written as $c_i = (m_i + k_i) \bmod 30$ and $m_i = (c_i - k_i) \bmod 30$, respectively.*

To prove that the one-time pad achieves perfect secrecy, we need some elementary notions from probability theory.

Definition 3 *Let $\mathbf{X}$ be a discrete random variable that can take on values from a finite set $\mathcal{X}$ according to a given probability distribution on $\mathcal{X}$. We denote by $\Pr[\mathbf{X} = x]$ the probability that $\mathbf{X}$ takes on the value $x \in \mathcal{X}$. If $\mathbf{X}$ is clear from the context, we just write $\Pr[x]$. For all $x \in \mathcal{X}$, $\Pr[x] \geq 0$. Additionally, $\sum_{x \in \mathcal{X}} \Pr[x] = 1$. For another random variable $\mathbf{Y}$ defined on the finite set $\mathcal{Y}$, we define the conditional probability that $\mathbf{X}$ takes on the value $x \in \mathcal{X}$ given that $\mathbf{Y}$ takes on the value $y \in \mathcal{Y}$ by $\Pr[x|y]$.*

plaintext $p \in \mathcal{P}$	O	N	E	-	T	I	M	E		P	A	D
p encoded	14	13	04	28	19	08	12	04	26	15	00	03
key k	06	13	02	01	14	05	07	18	05	26	13	28
ciphertext $c \in \mathcal{C}$	20	26	06	29	03	13	19	22	01	11	13	01
c decoded	U		G	.	D	N	T	W	B	L	N	B

Figure 4: Encryption and decryption example for the one-time pad

Suppose that a probability distribution on the finite plaintext space $\mathcal{P}$ is given. Thus, the plaintext element defines a random variable, which we denote by $\mathbf{p}$. Similarly, the key chosen by Alice and Bob for their communication defines a random variable on the key space, denoted by $\mathbf{k}$. Both probability distributions, for $\mathbf{p}$ and $\mathbf{k}$, induce a probability distribution on the ciphertext space $\mathcal{C}$, which gives another random variable $\mathbf{c}$ for the ciphertext element. We now define the notion of perfect secrecy that was introduced by Shannon [Sha49].

Definition 4 *A cryptosystem is said to achieve* perfect secrecy *if and only if for each $p \in \mathcal{P}$ and for each $c \in \mathcal{C}$,*

$$\Pr[p|c] = \Pr[p].$$

That means that the event that some plaintext p was encrypted is independent of the ciphertext c being observed. In other words, knowing c yields no advantage when trying to retrieve the original plaintext p.

Shannon [Sha49] gave a characterization of when perfect secrecy can be achieved. Suppose that $(\mathcal{P}, \mathcal{C}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ is a cryptosystem with $||\mathcal{K}|| = ||\mathcal{C}||$ and such that every plaintext element will be encrypted with a positive probability. Then, this cryptosystem achieves perfect secrecy if and only if

1. the keys in $\mathcal{K}$ are uniformly distributed, and
2. for each $p \in \mathcal{P}$ and for each $c \in \mathcal{C}$, there exists a unique key k such that $e_k(p) = c$.

The proof of Shannon's theorem can be found in, e.g., [Sha49,Rot05,Sti05]. Using this theorem, it is easy to see that the one-time pad satisfies the property of perfect secrecy. Since a new key element is created for each single plaintext element randomly under the uniform distribution, knowing the ciphertext is no advantage for an eavesdropper who seeks to recover the original message.

Although it provides perfect secrecy, the one-time pad also has severe disadvantages that make it impractical to use. Recall that the key has to be as large as the message itself. Thus, the number of bits that need to be exchanged over a secure channel for obtaining a joint secret key increases with the amount of information that Alice and Bob wish to transmit secretly. In light of this fact, one might ask why they don't use the secure channel directly for their communication. Using the same key for encryption more than once is no alternative, as the one-time pad's perfect secrecy crucially depends on creating a new key for every single plaintext element.

The scytale and the one-time pad are two examples of a symmetric cryptosystem. That means that the same key is used for encryption and decryption (or, at least, that the decryption key can be easily determined from the encryption key). Thus, Alice and Bob have to agree on a joint secret key prior to their conversation via a secure channel. Secret-key agreement protocols were proposed by Diffie and Hellman [DH76], Rivest and Sherman (see [RS97,HR99,HPR01,HRS05]), and others. However, a major disadvantage of symmetric ciphers and the related issue of key distribution occurs when the communicating parties in a large communication network need to share joint secret keys. When n parties participate, $n(n+1)/2$ different keys have to be exchanged and stored securely.

Public-key cryptosystems, also called asymmetric cryptosystems, circumvent the key distribution and storage problem. Instead of having one key for every pair of parties, only one key per party is needed to communicate securely. In 1976, Whitfield Diffie and Martin Hellman [DH76] proposed the principle idea of public-key cryptography, namely to use two distinct keys, a public key for encryption and a private key for decryption.

The first public-key cryptosystem is the RSA system, named after its three inventors Ron Rivest, Adi Shamir, and Leonard Adleman [RSA78].[1] Up to date, RSA is still used in numerous cryptographic applications. Public-key cryptosystems are based on so-called *(trapdoor) one-way functions*, functions that are easy to compute but hard to invert (unless one possesses a certain "trapdoor" information required for authorized decryption).

To communicate via a public-key cryptosystem, Alice creates two keys, k_{public} and k_{private}. Her encryption key k_{public} is public, but Alice keeps her private decryption key k_{private} secret. Each time Bob wishes to communicate with Alice, he looks up her public key and uses it to encrypt his message. Since only Alice knows her private key, she alone can (efficiently) compute the original message, i.e., the inverse of the encryption function.

The key issue is to find one-way functions that are secure enough to use for public-key cryptography. The first one-way function designed for this purpose (i.e., the RSA encryption function) was based on the problem of factoring large integers. Up to now, no efficient algorithm for computing the prime factors of some given integer is known. Other public-key cryptosystems—such as the ElGamal system [ElG85]—are based on the presumed hardness of computing discrete logarithms. One disadvantage of such systems is that they often lack a proof of security. Another disadvantage is that the directory storing the public keys has to be protected against manipulation and unauthorized access. If eavesdropper Eve replaces Alice's public key with her own key, she can decrypt all messages sent to Alice.

Since the publication of Peter Shor about prime factorization and computing discrete logarithms with quantum computers [Sho97], all cryptosystems whose security is based on the hardness of solving these mathematical problems have become theoretically vulnerable. Although it will certainly take some time for the first practical quantum computers to emerge, it is advisable to look for alternative, new cryptosystems whose security is not based solely on the hardness of solving such mathematical problems with current computer technology. Quantum theory seems to be the perfect basis on which to build such a new cryptosystem that withstands even an attack by quantum computers.

[1] In 1997, the British Government Communications Headquarters revealed that its researchers James Ellis, William Cocks, and Malcolm Williamson had independently and even earlier discovered the principle idea of public-key cryptography, the cryptosystem now called RSA, and the secret-key agreement protocol now called Diffie–Hellman, see, e.g., the discussion in [Sin99,Rot05].

3 From Bits to Qubits

The most important unit of information in computer science is the *bit*. There are two possible values that can be stored by a bit: the bit is either equal to "0" or equal to "1." These two different states can be represented in various ways, for example by a simple switch or by a capacitor: if not charged, the capacitor holds the value zero; if charged, it holds the value one.

In general, a quantum state $|\psi\rangle$ is an element of a finite-dimensional complex vector space (or Hilbert space) H. We denote the scalar product of two states $|\psi\rangle$ and $|\phi\rangle$ by $\langle\psi|\phi\rangle$, where $\langle\psi| = \overline{|\psi\rangle}^T$ is the conjugate transpose of $|\psi\rangle$.[2] It is convenient to deal with normalized states, so we require $\langle\psi|\psi\rangle = 1$ for all states $|\psi\rangle$ that have a physical meaning.

The quantum analog of the bit is called *qubit*, which is derived from *qu*antum *bit*. A qubit $|\psi\rangle$ is an element of a two-dimensional Hilbert space, in which we can introduce an orthonormal basis, consisting of the two states $|0\rangle$ and $|1\rangle$. Unlike its classical counterpart, the quantum state can be in any *coherent superposition* of the basis states:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{3.1}$$

where α and β are, in general, complex coefficients. This is due to the fact that the quantum mechanical equation of motion, the Schrödinger equation, is linear: Any linear superposition of its solutions (the quantum states) is also a solution. Since we require quantum states to be normalized, we find that the coefficients in (3.1) have to fulfill $|\alpha|^2 + |\beta|^2 = 1$, where $|\cdot|$ denotes the absolute value.

There exist many possibilities to physically represent a qubit in practice, as every quantum system with at least two states can serve as a qubit. For example, the spin of an atom or the polarization[3] of a light particle can represent the state of a qubit. Even a cat with its two basic states "dead" and "alive," introduced by Schroedinger [Sch35] to visualize fundamental concepts of quantum mechanics, might serve as a representation. The cat's problem—or fortune from the animal's point of view—when being used as a quantum system is its sheer size compared to that of an atom or light particle. There is no way to protect such a big quantum instance from interaction with its environment, which in turn will result in decoherence of the superposition of the cat. For the rest of the chapter, we will leave the cat alone and use light particles as our preferred qubits.

The physical meaning of (3.1) can most easily be understood when we measure the quantum state $|\psi\rangle$. In quantum mechanics, this is achieved by a *positive operator valued measurement* (POVM), which is a set of positive-definite, hermitian operators $\mathcal{E} = \{E_x\}_{x\in\mathcal{X}}$ acting on the Hilbert space of the qubit. The elements of this set have to sum up to the identity, $\sum_{x\in\mathcal{X}} E_x = \mathbb{1}$. A simple, special case occurs when the E_x are orthogonal projectors, i.e., $E_x = |\phi_x\rangle\langle\phi_x|$ and $\langle\phi_x|\phi_y\rangle = \delta_{xy}$. This simple projection measurement is called *von Neumann measurement*. The result x of a von Neumann measurement will occur with probability $\Pr[x] = \langle\psi|E_x|\psi\rangle = |\langle\psi|\phi_x\rangle|^2$. Consider our qubit being represented by the polarization states of a photon. We denote horizontal polarization by

[2] Mathematically, $\langle\psi|$ is an element of the dual space H^*.

[3] Light particles, called photons, can be seen as electromagnetical waves. A specific property of them is their transversality, which means that the electric and the magnetic fields are orthogonal to each other and to the propagation direction. The inclination of the electric (or magnetic) field to the axis of the propagation is called polarization.

$|0\rangle$ and vertical polarization by $|1\rangle$. It is a physical property of the electromagnetic field that these two states are orthogonal,[4] i.e., $\langle 0|1\rangle = 0$, and thus form a basis in the two-dimensional Hilbert space. A simple measurement that tells us whether the qubit is in the state $|0\rangle$ or $|1\rangle$ is given by the projection set $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$. When performing this measurement, the qubit will be found in the state $|0\rangle$ with probability $|\alpha|^2$, and in the state $|1\rangle$ with probability $|\beta|^2$. We are free to choose a different basis in the Hilbert space; for instance, the one given by the two states

$$\begin{aligned} |0\rangle_\times &= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad \text{and} \\ |1\rangle_\times &= \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle). \end{aligned}$$

This is a rotated basis, and a photon in the state $|0\rangle_\times$ and $|1\rangle_\times$, respectively, has a polarization of $\pm 45°$ against the horizontal. If we measure in this basis by means of the projection measurement $\{|0\rangle_\times\langle 0|, |1\rangle_\times\langle 1|\}$, we find the qubit in the state $|0\rangle_\times$ with probability $1/2 + \Re(\alpha\bar{\beta})$, and in the state $|1\rangle_\times$ with probability $1/2 - \Re(\alpha\bar{\beta})$. Let us consider the special case where, for instance, $\beta = 0$: When we do the first measurement, we find the qubit in the state $|0\rangle$ with certainty. But when we apply the second measurement, the outcome will be completely random. This is an important property of the *conjugated bases* $\{|0\rangle, |1\rangle\}$ and $\{|0\rangle_\times, |1\rangle_\times\}$ with $|\langle i|j\rangle_\times| = 1/\sqrt{2}$ for all i and j, which will be exploited in many quantum key distribution protocols, as described below.

From POVMs it is just a small step to *observables*. Each measurable physical quantity is represented by a hermitian operator, called observable. When we write an observable A in its spectral decomposition, $A = \sum_i \lambda_i |i\rangle\langle i|$, where $\langle i|j\rangle = \delta_{ij}$, the corresponding POVM is given by the orthogonal projectors $\{|i\rangle\langle i|\}$. A measurement of A always yields one of the eigenvalues λ_i as a result, and the measured quantum state *collapses* onto the corresponding state $|i\rangle$.

An important concept in quantum mechanics is the *density matrix* or *density operator* ρ: The density matrix of a so-called *pure* state $|\psi\rangle$ is given by the projector $|\psi\rangle\langle\psi|$. In the case of a qubit, this is a complex-valued (2×2) matrix. The advantage of this representation is the possibility to describe systems with a statistical distribution of states. For instance, consider a system that is known to be in the state $|\psi_x\rangle$ with probability $\Pr[x]$, for $x \in \mathcal{X}$. Let $\mathcal{E} = \{E_y\}_{y\in\mathcal{Y}}$ be some POVM. Then the probability to get the result y *if* the system was known to be in the state $|\psi_x\rangle$ would be $\langle\psi_x|E_y|\psi_x\rangle$. But since we do not know, we have to average over all possible states, just as we would do if the system were prepared many times in one of the states $\{|\psi_x\rangle\}$ and we had repeated the measurement each time. The probability to measure y in the *ensemble* $\{|\psi_x\rangle, \Pr[x]\}$ is consequently

$$\Pr[y] = \sum_{x\in\mathcal{X}} \Pr[x]\langle\psi_x|E_y|\psi_x\rangle = \mathrm{tr}\left(E_y \sum_{x\in\mathcal{X}} \Pr[x]|\psi_x\rangle\langle\psi_x|\right), \tag{3.2}$$

where $\mathrm{tr}A$ denotes the trace of the matrix A, i.e., the sum of its diagonal elements. We can now introduce the density matrix $\rho = \sum_{x\in\mathcal{X}} \Pr[x]|\psi_x\rangle\langle\psi_x|$, such that (3.2) takes the simple form:

[4]This is by no means a consequence of the geometric relationship between "horizontal" and "vertical." For instance, the spin of a spin-1/2 particle like the electron can point "up" or "down," and the corresponding states $|\uparrow\rangle$ and $|\downarrow\rangle$ are orthogonal. However, the angle between the two spin settings is certainly not 90 degree.

$\Pr[y] = \mathrm{tr}(E_y \rho)$. From now on, we can concentrate on density matrices solely, since any *pure* state $|\psi\rangle$ is just a special case where one probability in the ensemble $\{|\psi_x\rangle, \Pr[x]\}$ is equal to one and all others vanish. In the general case, i.e., when at least two different states in the ensemble occur with nonvanishing probability, the system is said to be in a *mixed* state.

Once we consider composite quantum systems, the situation becomes more complicated—and more interesting. Let us consider Alice holding a state ρ_A, acting on a Hilbert space H_A, and Bob holding a state ρ_B acting on H_B. Both states are part of a total state ρ_{AB}, acting on the tensor product $H_A \otimes H_B$, and they are related by the partial trace: $\rho_A = \mathrm{tr}_B \rho_{AB}$ and $\rho_B = \mathrm{tr}_A \rho_{AB}$. This operation discards degrees of freedom in the respective subsystem. Composite states, such as ρ_{AB} can be divided into two classes: *separable* and *entangled* states. We first look at pure states, which means that ρ_{AB} is of the form $\rho_{AB} = |\psi_{AB}\rangle\langle\psi_{AB}|$. Separable pure states are *product states*:

$$|\psi_{AB}\rangle = |\psi_A\rangle \otimes |\psi_B\rangle \equiv |\psi_A\rangle|\psi_B\rangle \equiv |\psi_A \psi_B\rangle.$$

(The last three expressions are equivalent notations.) They are composed of two independent states of the two subsystems A and B. Pure states that cannot be written in this form are called entangled. A famous example of pure entangled states are the Bell states:

$$|\phi^\pm\rangle = \frac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle), \tag{3.3}$$

$$|\psi^\pm\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle). \tag{3.4}$$

These four states form a basis in the two-qubit Hilbert space. A mixed state is called *separable* if and only if it can be written as a convex sum of projectors onto product states [Wer89]:

$$\rho = \sum_{x \in \mathcal{X}} \Pr[x] |\psi_x^A \phi_x^B\rangle\langle\psi_x^A \phi_x^B| = \sum_{x \in \mathcal{X}} \Pr[x] |\psi_x^A\rangle\langle\psi_x^A| \otimes |\phi_x^B\rangle\langle\phi_x^B|, \tag{3.5}$$

with $\Pr[x] \geq 0$ for each $x \in \mathcal{X}$ and $\sum_{x \in \mathcal{X}} \Pr[x] = 1$. These states can be prepared locally in Alice's and Bob's laboratory only by means of *classical* communication, i.e., no quantum systems need to be sent. If a state cannot be written in the form (3.5), it is called *entangled*.

4 Quantum Key Distribution

Quantum cryptography exploits the quantum mechanical property that a qubit cannot be copied or amplified without disturbing its original state. This is the statement of the *No-Cloning Theorem* [WZ82], which is easily proven: Assume there exists a unitary transformation[5] U that can copy two states $|\psi_1\rangle$ and $|\psi_2\rangle$:

$$U|\psi_1\rangle|0\rangle = |\psi_1\rangle|\psi_1\rangle, \tag{4.6}$$

$$U|\psi_2\rangle|0\rangle = |\psi_2\rangle|\psi_2\rangle, \tag{4.7}$$

[5]The time-evolution of an isolated quantum system is described by a unitary transformation U: $|\psi\rangle \to U|\psi\rangle$.

where $|0\rangle$ is an arbitrary input state. If we equate the scalar products of the left-hand and right-hand sides, it follows by the unitarity of U that $\langle\psi_1|\psi_2\rangle = \langle\psi_1|\psi_2\rangle^2$, which implies that $\langle\psi_1|\psi_2\rangle$ equals 0 or 1. This means that we can copy only orthogonal or identical states. In contrast, arbitrary unknown states cannot be perfectly cloned. (Note that orthogonal or identical states are not viewed as "unknown" states, since we do know they are orthogonal, for example.)

The essence of this theorem is the main ingredient of quantum key distribution, where Alice and Bob use a quantum channel to exchange a sequence of qubits, which will then be used to create a key for the one-time pad in order to communicate over an insecure channel. Any disturbance of the qubits, for example caused by Eve trying to measure the qubits' state, can be detected with high probability.

In this section, we describe the BB84 protocol proposed by Charles Bennett and Gilles Brassard in 1984, see [BB84]. This is the first protocol designed to employ quantum mechanics for two parties to agree on a joint secret key.

4.1 The BB84 Protocol

In this protocol, Alice and Bob use a quantum channel to send qubits. They are also connected by a classical channel, which is insecure against an eavesdropper but unjammable. Alice and Bob use four possible quantum states in two conjugate bases (say, the rectilinear basis $+$ and the diagonal basis $\times$). We use $|0\rangle_+$ and $|0\rangle_\times = (|0\rangle_+ + |1\rangle_+)/\sqrt{2}$ for the classical signal "0," and we use $|1\rangle_+$ and $|1\rangle_\times = (|0\rangle_+ - |1\rangle_+)/\sqrt{2}$ for the classical signal "1." Note that the two bases are connected by the so-called Hadamard transformation

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \tag{4.8}$$

in the following way: Since $H^2 = \mathbb{1}$, we have $H|0\rangle_+ = |0\rangle_\times$ and $H|1\rangle_+ = |1\rangle_\times$, and vice versa.

The protocol works as follows (see also Table 1 for illustration):

1. Alice randomly prepares $2n$ qubits, each in one of the four states $|0\rangle_+$, $|0\rangle_\times$, $|1\rangle_+$, or $|1\rangle_\times$, and sends them to Bob.

2. For each qubit that Bob receives, he chooses at random one of the two bases ($+$ or $\times$) and measures the qubit with respect to that basis. If he chooses the same basis as Alice, his measurement result is the same as the classical bit that Alice prepared. If the bases differ, Bob's result is completely random.

3. Alice tells Bob via the classical channel which basis she used for each qubit. They keep the bits where Bob has used the same basis for his measurement as Alice. This happens in about half the cases, so they will have approximately n bits left. These are forming the so-called *sifted key*.

4. Alice and Bob choose a subset of the sifted key to estimate the error-rate. They do so by announcing publicly the bit values of the subset. If they differ in too many cases, they abort the protocol, since its security cannot be guaranteed.

5. Finally, Alice and Bob obtain a joint secret key from the remaining bits by performing error correction and privacy amplification.

Alice's string	1	1	0	1	0	0	1	0	1	1	1	1	0	0
Alice's basis	$+$	$+$	$+$	$\times$	$\times$	$+$	$\times$	$\times$	$\times$	$\times$	$+$	$+$	$+$	$+$
Bob's basis	$+$	$\times$	$+$	$+$	$\times$	$+$	$\times$	$+$	$\times$	$\times$	$+$	$+$	$+$	$+$
Bob's string	1	R	0	R	0	0	1	R	1	1	1	1	0	0
Same basis?	Y	N	Y	N	Y	Y	Y	N	Y	Y	Y	Y	Y	Y
Bits to keep	1		0		0	0	1		1	1	1	1	0	0
Test	Y		N		N	Y	N		N	N	N	Y	Y	N
Key			0		0		1		1	1	1			0

Table 1: The BB84 Key Distribution Protocol. Here, "Y" and "N" stand for "yes" and "no," respectively, and "R" means that Bob obtains a random result.

Which possibilities does Eve have to attack this protocol? And, consequently, what is the threshold of the error-rate, at which Alice and Bob should abort the protocol? To answer these questions, we look at a simple eavesdropping strategy, which is called "intercept and resend."

Eve's goal is to learn at least some part of the key. Thus, an obvious strategy for her is to intercept the qubits being transmitted from Alice to Bob. She cannot simply copy the qubits, since this would contradict the No-Cloning Theorem. In order to extract some information, she is forced to measure (and thus destroy) them. But since she does not know the basis in which they were prepared (Alice announces this information only after Bob received all signals), she can only guess or just flip a coin for the selection of the measurement basis. In about half the cases, she will happen to choose the same basis as Alice and get completely correlated bit values. In the other half, her results will be random and uncorrelated. Bob certainly expects to receive something from Alice, so Eve needs to send some qubits to him. However, she still has no idea which basis Alice used, so she prepares each qubit in the same basis as she measured it (or she chooses a basis at random). These newly created qubits again match Alice's bases in only half of the cases. After Bob receives Eve's qubits, he measures them, and Alice and Bob apply the sifting. Because of Eve's disturbance, about half of Bob's key was measured in a different basis than it was prepared by Alice. Since Bob's result is random in those cases, his sifted key will contain about 25% errors. In the error-estimation stage, if Alice and Bob obtain such a high error rate, it would be wise for them to abort the protocol.

If the error rate is below an agreed threshold value, Alice and Bob can eliminate errors with (classical) error correction. A simple method for error correction works as follows: Alice chooses two bits at random and tells Bob the XOR-value of the two bits. Bob tells Alice if he has the same value. In this case, they keep the first bit and discard the second bit. If their values differ, they discard both bits. The remaining bits form the key.

The last stage of the protocol is privacy amplification—a procedure in which Alice and Bob eliminate (or, at least, drastically reduce) Eve's knowledge about the key. They do so by choosing random pairs of bits of the sifted key and replacing them by their corresponding XOR-values. Thus, they halve the length of the key, in order to "amplify" their privacy. Note that Eve has less knowledge about the XOR-value, even if she knew the values of the single bits with high probability (but not with certainty).

There are even more sophisticated methods for error correction and privacy amplification. For more details on error-correcting codes and their usage in the physics of quantum information, we refer to Huffman and Pless [HP03] and Bouwmeester, Ekert, and Zeilinger [BEZ00].

4.2 Security of Quantum Key Distribution

Unlike many of the classical cryptosystems in use today, whose security often draws on unproven assumptions about the computational complexity of mathematical problems, the security of quantum cryptography is based on—and employs—the laws of physics. The term "unconditional security" is used to emphasize the fact that it does not rely on the presumed, yet unproven hardness of some mathematical problem. In this section, we present the proof of the unconditional security of the BB84 protocol, as devised by Peter Shor and John Preskill [SP00].

We divide the proof into three parts:

- In the first part, we present the so-called *entanglement-based* version of the BB84 protocol. In contrast, the scheme presented in the previous section is called a *prepare-and-measure scheme*, for obvious reasons. In the entanglement-based version, Alice and Bob's aim is to share a special entangled state that allows them to obtain perfectly correlated bits upon measuring their half of the state. We will see how they can construct such a state, how they can check whether they were successful, and how they can detect Eve's attempted attack.

- In the second part, we will show that the equivalent entanglement-based version is secure. More precisely, we will prove that Eve cannot deceive Alice and Bob into continuing the protocol, falsely believing that they are creating a secure key.

- In the third part, we show that the two schemes are equivalent indeed.

4.2.1 The Entanglement-Based Version of BB84

In this version of the protocol, Alice and Bob aim at creating a special entangled state, namely the Bell state

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle), \tag{4.9}$$

where Alice holds the first particle and Bob holds the second one. An important property of this state is that it is *maximally entangled*. This means that Alice's and Bob's measurement results are completely correlated whenever they measure the state $|\phi^+\rangle$ in the same basis. (Moreover, their results are random.) Since the state is pure, it cannot be entangled with anything else, in particular not with anything under Eve's control. Thus, whenever Alice and Bob are sure they share a $|\phi^+\rangle$ state, they know that (a) measuring in the same basis generates a shared random bit, and (b) Eve has no knowledge about this bit. To generate the whole key, Alice and Bob prepare a large number of these Bell states,

$$|\phi^+\rangle^{\otimes n} = |\phi^+\rangle \otimes \cdots \otimes |\phi^+\rangle,$$

and measure each qubit separately. We will now show how they can achieve this.

We need to take a brief detour to quantum error correction first. In contrast to a classical bit, a qubit can undergo three different errors: bit flips, phase errors, and combinations of these two:

- When a bit flip occurs, the state $|0\rangle$ becomes $|1\rangle$, and vice versa. This error is described by the Pauli matrix
$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$
- Phase errors transform the state $|1\rangle$ into $-|1\rangle$, but leave $|0\rangle$ unchanged. Such an error is described by the Pauli matrix
$$\sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}.$$
- Both these errors can also occur combined. For example, changing $|0\rangle$ to $-|1\rangle$ and $|1\rangle$ to $|0\rangle$ can be described by $\sigma_z\sigma_x = i\sigma_y$, where
$$\sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Let us now recall some elements of classical error correction. A (classical) linear $[n, k]$ code C that encodes k bits of information by an n bit string is a set of k codewords. Each codeword is an n-dimensional binary vector. The whole code can be described by an $(n \times k)$-dimensional generator matrix G that maps each message x to the encoded message Gx. Thus, the set of all possible codewords is the vector space that is spanned by the columns of G. We require those vectors to be linearly independent. Error correction for linear codes can be easily described by means of the parity check matrix H. This is an $((n - k) \times n)$ matrix with the property that $Hx = 0$ for all codewords x.

Suppose now that a message x is encoded as $y = Gx$. Due to an error e, one obtains $y' = y + e$. Since we have $Hy = 0$ for all codewords, it follows that $Hy' = He$, which is called the *(error) syndrome*. Thus, if the syndrome is 0, no error has occurred. Otherwise, H is constructed such that the syndrome contains information about the error that should make it possible to correct it. Finally, we introduce the concept of *duality*: Let C be a linear $[n, k]$ code with generator matrix G and parity check matrix H. Then we can define the dual code $C^{\perp}$ of C, which is the set of all codewords that are orthogonal to each codeword in C. The dual code $C^{\perp}$ is an $[n - k, n]$ code which is generated by H^T and has a parity check matrix G^T. Dual codes play an important role in the construction of CSS codes, as we explain below.

Definition 5 *Let C_1 and C_2 be classical linear $[n, k_1]$ and $[n, k_2]$ codes, respectively, such that $C_2 \subset C_1$. For each codeword $x \in C_1$, define the quantum state*
$$|x + C_2\rangle = \frac{1}{\sqrt{|C_2|}} \sum_{y \in C_2} |x + y\rangle. \tag{4.10}$$
The space spanned by $\{|x + C_2\rangle\}_{x \in C_1}$ defines an $[n, k_1 - k_2]$ quantum code, which is called the Calderbank–Shor–Steane code, $\mathrm{CSS}(C_1, C_2)$ *for short.*

Let x and x' in C_1 be codewords such that $x - x'$ is in C_2. Then one can show that

$$|x + C_2\rangle = |x' + C_2\rangle,$$

i.e., the state $|x + C_2\rangle$ depends only on C_1/C_2, that is, the coset to which x belongs.[6] It follows that if x and x' belong to different cosets, the states $|x + C_2\rangle$ and $|x' + C_2\rangle$ are orthogonal. As the number of cosets of C_2 in C_1 is $|C_1|/|C_2|$, the dimension of the space $\mathrm{CSS}(C_1, C_2)$ is $|C_1|/|C_2| = 2^{k_1-k_2}$, thus $m = k_1 - k_2$ qubits can be encoded.

Error correction with CSS codes works as follows. Suppose that C_1 and $C_2^{\perp}$ both can correct ℓ errors. Moreover, let H_1 be the parity check matrix for C_1, and let H_2 be that for $C_2^{\perp}$. Define

$$\sigma_\alpha^s = \sigma_\alpha^{s_1} \otimes \sigma_\alpha^{s_2} \otimes \cdots \otimes \sigma_\alpha^{s_n}, \tag{4.11}$$

where $\alpha \in \{x, y, z\}$, $\sigma_\alpha^0 = \mathbb{1}$, and $s = (s_1, s_2, \ldots, s_n)$ is an n bit vector. It can be shown that the syndrome for bit flip errors can be computed by measuring σ_z^r for each row vector r of H_1. Similarly, the syndrome for phase errors can be computed by measuring σ_x^t for each row vector t of H_2. In this way, ℓ bit flips and ℓ phase errors can be corrected.

We have now collected all the ingredients to describe the entanglement-based version of the BB84 protocol:

1. Alice creates $2n$ qubit pairs in the state $|\phi^+\rangle^{\otimes 2n}$.
2. She randomly selects n of those qubits which will later serve as check qubits.
3. Alice selects a random $2n$ bit string b and applies the Hadamard transformation (4.8) to her half of each qubit pair whenever the corresponding bit of b is "1."
4. She sends the other half of all qubit pairs to Bob.
5. Alice announces b and which qubits are to serve as check qubits.
6. Bob performs a Hadamard transformation on those of his qubits where b is "1."
7. Alice and Bob measure the check qubits in the $\{|0\rangle, |1\rangle\}$ basis to estimate the error rate. If more than ℓ results differ, they abort the protocol.
8. For the remaining qubits, Alice and Bob measure the syndromes for the codes C_1 and C_2, correct the errors, and obtain $|\phi^+\rangle^{\otimes m}$.
9. They measure this state in the $\{|0\rangle, |1\rangle\}$ basis to obtain a shared secret key.

The point of performing the Hadamard transformation on half of the qubits is that this operation effectively changes the basis, in which the qubits are prepared, from $\{|0\rangle_+, |1\rangle_+\}$ to $\{|0\rangle_\times, |1\rangle_\times\}$. This is necessary because if Eve knew the basis, she could launch the intercept-resend attack presented in the previous section and break the protocol.

[6] Let G and H be two groups with $G \subset H$. Then for any $h \in H$, we define the coset of G in H, determined by h, as $hG = \{h + g \mid g \in G\}$. The group H/G is the set of all elements of H that belong to different cosets.

4.2.2 Security of the Entanglement-Based Version

The goal of the entanglement-based version of the BB84 protocol is to provide Alice and Bob with a number of qubit pairs in the state $|\phi^+\rangle$, because measuring this state in a joint basis generates correlated, random bits. The distribution of this state is done by means of CSS codes, but this operation may not be 100% efficient. Moreover, the random sampling which Alice and Bob apply by randomly choosing a set of check bits may not provide a perfect sample. In this section, we deal with these issues.

We first need a technical lemma stating, roughly speaking, that a state that is "close" to the state $|\phi^+\rangle^{\otimes m}$ has a small entropy. The "distance" to a pure state is measured by means of the so-called *fidelity*, which is defined as $F(\rho, |\psi\rangle) = \langle\psi|\rho|\psi\rangle$. If $F = 1$, the two states are identical.

Lemma 6 *Let ρ be a density matrix and $s > 0$. If $F(\rho, |\phi^+\rangle^{\otimes m})^2 \geq 1 - 2^{-s}$, then*

$$S(\rho) \leq (2m + s + 1/\ln 2)2^{-s} + \mathcal{O}(2^{-2s}).$$

Here $S(\rho) = -\mathrm{tr}(\rho \ln \rho)$ denotes the von Neumann entropy. The proof of this lemma is simple and can be found in [NC00]. The amount of information that can be extracted from a quantum state is given by the *accessible information*. Holevo's bound [Hol73] shows that it can be upper-bounded by the von Neumann entropy $S(\rho)$. Now using the above lemma, we have shown that if the state shared by Alice and Bob is close to the state $|\phi^+\rangle^{\otimes m}$, then the information extractable by Eve is negligibly small.

It remains to show that by the random sampling that Alice and Bob apply, they can reliably estimate the fidelity of the remaining qubits. The main ingredient to prove this is again a lemma, which we state here without proof. (The proof is left to the reader, see Nielsen and Chuang [NC00].)

Lemma 7 *Let a random $2n$ bit string that might contain some errors, and a random subset of n check bits of that string be given. Then, for any two constants $\delta > 0$ and $\epsilon > 0$, the probability of finding less than δn errors on the check bits, and more than $(\delta + \epsilon)n$ errors on the remaining bits is less than $e^{-\mathcal{O}(\epsilon^2 n)}$, for sufficiently large n.*

Although this lemma is based on classical probability theory, we can give an argument for its validity in the quantum world: The observables that Alice and Bob measure on the check bits are both diagonal in the Bell basis, which means that the statistics of the results can be described purely classically. These measurements on $H_A \otimes H_B$ are given by the POVMs

$$\{P_{\mathrm{bf}} = |\psi^+\rangle\langle\psi^+| + |\psi^-\rangle\langle\psi^-|, \mathbb{1} - P_{\mathrm{bf}}\},$$

which are used to check for bit flips, and

$$\{P_{\mathrm{pe}} = |\phi^-\rangle\langle\phi^-| + |\psi^-\rangle\langle\psi^-|, \mathbb{1} - P_{\mathrm{pe}}\},$$

which are used to check for phase errors. Alice and Bob choose one of those measurements at random for each check qubit. In this way, they can calculate a lower bound for the fidelity of the remaining qubits.

Summarizing, these rough arguments show that by random sampling, the fidelity of the state shared by Alice and Bob can be lower-bounded, with an exponentially small probability of error. Moreover, with this bound, the information that Eve can obtain about this state (and consequently about the secret key) can be shown to be also exponentially small.

4.2.3 Equivalence of the Two Schemes

We prove the equivalence of the entanglement-based and prepare-and-measure versions of the BB84 protocol by successive simplifications. Each step is very simple, so it is easy to verify that the security of the protocol is not compromised.

A major simplification is that all measurements done by Alice after transmitting the particles can already be done at the very beginning: If Alice measures her part of the state $|\phi^+\rangle$, she obtains a random bit as a result, but on the other hand, Bob's part of the state collapses onto the correlated state $|0\rangle$ or $|1\rangle$. Thus, instead of sending entangled qubits for the check, Alice can as well prepare single qubits randomly in one of the states $|0\rangle$ and $|1\rangle$, and send those states to Bob. Of course, it is crucial for the security of the protocol that Eve does not know *a priori* which qubits will serve as check qubits and which as "key qubits"; otherwise, she could treat them differently and thus fudge the error estimation.

Another measurement Alice can do at the beginning is the measurement of her syndrome and her key qubits. This is not very obvious, so let us give some more detail: Given a CSS code $\mathrm{CSS}(C_1, C_2)$, we can define a family of equivalent codes $\mathrm{CSS}_{v,w}(C_1, C_2)$, in the sense that they have the same error correcting properties. The codewords of the code $\mathrm{CSS}_{v,w}(C_1, C_2)$ are given by

$$|x_k, v, w\rangle = \frac{1}{\sqrt{|C_2|}} \sum_{y \in C_2} (-1)^{v \cdot y} |x_k + y + w\rangle, \tag{4.12}$$

where x_k is one representative of one of the m cosets of C_2 in C_1, and v and w are arbitrary n bit strings. Since the $\{|x_k, v, w\rangle\}$ form a basis, we can rewrite

$$|\phi^+\rangle^{\otimes n} = \frac{1}{\sqrt{2^n}} \sum_{i=0}^{2^n-1} |i\rangle|i\rangle = \frac{1}{\sqrt{2^n}} \sum_{x_k, v, w} |x_k, v, w\rangle|x_k, v, w\rangle, \tag{4.13}$$

where i is in binary notation. If now Alice measures the error syndromes, namely σ_z^r for each row vector r of H_1 and σ_x^t for each row vector t of H_2, she obtains a random result for v and w. Finally, if she does a last measurement in the $\{|0\rangle, |1\rangle\}$ basis, she obtains a random codeword x_k. From (4.13), we see that Bob's state then collapses onto $|x_k, v, w\rangle$, which is a random qubit encoded in a random code.

As an intermediate result, we rephrase the entanglement-based protocol including all simplifications introduced so far:

1. Alice creates n random check qubits, each in the state $|0\rangle$ or $|1\rangle$, a random n bit string k, which will serve as the key, and two random n bit strings v and w. She prepares the state $|k\rangle$ and encodes it using $\mathrm{CSS}_{v,w}(C_1, C_2)$.

2. She randomly selects n positions for the check qubits and puts the encoded qubits in the remaining positions.

3. Alice selects a random $2n$ bit string b and applies the Hadamard transformation to her half of each qubit pair where b is "1."

4. She sends the other half of all qubit pairs to Bob.

5. Alice announces b, v, and w, and which qubits are to serve as check qubits.

6. Bob performs a Hadamard transformation on those of his qubits where b is "1."

7. Bob measures the check qubits in the $\{|0\rangle, |1\rangle\}$ basis. If he finds more than ℓ results that disagree with Alice's prepared states, they abort the protocol.

8. Bob decodes the key qubits from $\mathrm{CSS}_{v,w}(C_1, C_2)$ and obtains the state $|k\rangle$.

9. He measures $|k\rangle$ in the $\{|0\rangle, |1\rangle\}$ basis and obtains the key k as the result.

We will now simplify this protocol even further: Note that in the original version, Alice and Bob do not care whether they shared the state $|\phi^+\rangle$ or $|\phi^-\rangle = (|00\rangle - |11\rangle)/\sqrt{2}$, because measuring both states provides them with correlated, random bits; the relative phase is irrelevant. Thus, it is unnecessary to send the phase correction information v to Bob. This is why CSS codes are used: They decouple the bit flip error correction from the phase error correction. If now Bob were to measure his key qubits before the decoding, he would obtain $x_k + y + w + e$, where e denotes the bit errors that occurred during the transmission (or that were introduced by Eve). He can now classically decode this bit string by subtracting w, which was announced by Alice, and correct it to the codeword $x_k + y$, if e did not introduce too many errors. Bob finds the key by computing the coset to which $x_k + y$ belongs. But since Bob does not need v, why should Alice send it? If she never reveals that value, she effectively prepares a state that is a classical mixture of all possible values that v can take, weighted with the corresponding probabilities:

$$\rho_{x_k,w} = \frac{1}{2^n} \sum_v |x_k, v, w\rangle\langle x_k, v, w| = \frac{1}{|C_2|} \sum_{z \in C_2} |x_k + z + w\rangle\langle x_k + z + w|. \tag{4.14}$$

We see that this state can also be prepared by classically choosing a random codeword $z \in C_2$ and constructing $|x_k + z + w\rangle$. Thus, the preparation in Step 1 can be done equivalently in the following way: Alice creates n random check qubits, each in the state $|0\rangle$ or $|1\rangle$, a random n bit string w, a random string $x_k \in C_1/C_2$, and a random codeword $z \in C_2$. The n key qubits are prepared in the state $|x_k + w + z\rangle$, and the check qubits are placed at random positions.

Note that we can also remove the need for $z \in C_2$, if Alice instead of choosing $x_k \in C_1/C_2$ chooses $x_k \in C_1$. With this modification, Alice sends the state $|x_k + w\rangle$ as key qubits, which Bob then measures and corrects to $x_k + w$. Since $x_k + w$ is a completely random n bit string, Alice can as well just prepare $|y\rangle$, where y is a random n bit string. She sends it to Bob who measures it to obtain $y + e$, then Alice sends error correction information $y - x_k$, which Bob subtracts from $y + e$ to finally obtain $x_k + e$. He corrects it to x_k and calculates the key k as the coset to which x_k belongs.

What we have achieved is that now the check and the key qubits are just prepared randomly in one of the states $|0\rangle$ or $|1\rangle$. The whole protocol so far looks as follows:

1. Alice creates $2n$ random qubits, each in the state $|0\rangle$ or $|1\rangle$, and a random codeword $x_k \in C_1$.
2. She randomly selects n positions to be check qubits and the remaining n positions to be $|y\rangle$.
3. Alice selects a random $2n$ bit string b and applies the Hadamard transformation to her half of each qubit pair where b is "1."
4. She sends the other half of all qubit pairs to Bob.
5. Alice announces b and $y - x_k$, and which qubits are to serve as check qubits.
6. Bob performs a Hadamard transformation on those of his qubits where b is "1."
7. Bob measures the check qubits in the $\{|0\rangle, |1\rangle\}$ basis. If he finds more than ℓ results that disagree with Alice's prepared state, they abort the protocol.
8. Bob measures the key qubits to get $y + e$, subtracts $y - x_k$, and corrects $x_k + e$ to x_k.
9. He calculates the coset to which x_k belongs to get the key k.

Finally, we can remove the Hadamard transformation, and let Alice choose randomly one of the four states in $\{|0\rangle_+, |1\rangle_+, |0\rangle_\times, |1\rangle_\times\}$. Then Bob, instead of waiting for b to be announced, simply chooses one basis at random and measures the arriving qubits. As he will choose the wrong basis in roughly half the cases, Alice should double the number of input qubits to $4n$. After his measurement, Alice announces which basis she used and both discard all instances where they used a different basis. With this last modification, we finally arrived at the prepare-and-measure version of the BB84 protocol, only up to some small twists.

5 Quantum Bit Commitment

When talking about quantum cryptography, everyone is thinking about key distribution. There are, however, other cryptographic applications as well, such as bit commitment. In 1993, a bit commitment protocol based on quantum mechanics was introduced by Brassard et al. [BCJL93]. The unconditional security of the protocol (which means that the security of the protocol is independent of the computational resources, such as computing time, amount of memory used, and computer technology of the cheater) has been accepted without proof [Yao95]. Two years after it had been proposed, the protocol turned out to be insecure [May95].

A commitment protocol is a procedure in which one party, say Alice, deposits a message such that no one (and in particular not Alice) can read it nor change it. At some point in the future, Alice's message will be announced, and with high certainty it can be proven that the revealed message is the same as the one Alice had deposited originally. To illustrate this situation, suppose Bob wants to auction off a diamond ring, subject to the condition that each person wishing to participate in the auction can bid only one single amount of money. After each person has chosen a specific amount, the highest bidder gets the ring. So everyone writes their own bid on a piece of paper, puts it into a personal safe, which is then locked and given to Bob. Until all bids have been submitted to Bob, each bidder keeps the key matching the lock of his or her safe. In this way Bob cannot see any of the bids, which in turn cannot be changed once they have been submitted, since only Bob has access to the committed safes. All keys are handed over to Bob after he has received all safes from the people participating in the auction. The different offers are compared in public, so that everybody can be sure that only the highest bidder walks away with the diamond and an empty wallet.

We can describe this commitment protocol mathematically as follows: The protocol has two stages, the commit phase and the unveil phase. Alice commits herself to the data m by computing $c = f(m)$, and she sends c to Bob. Alice unveils the commitment by showing Bob the preimage m of c. In classical cryptography, and in particular in public-key cryptography, one-way functions are used for commitment. In quantum cryptography, we want to make use of the laws of quantum mechanics to create a fair protocol for both sides. Bit commitment is a special case of a commitment protocol, where the data m consists of only one single bit.

It is widely believed that it is impossible to create a perfectly secure classical bit commitment protocol. Regarding the extension to the quantum world, it was shown that unconditionally secure quantum bit commitment is also impossible [May97,LC97].

5.1 The BB84 Quantum Bit Commitment Protocol

The BB84 protocol was introduced in Section 4.1. A quantum bit commitment protocol can be created from the BB84 quantum key distribution protocol with a few minor changes [BB84]. Just as in the classical bit commitment protocol, the quantum protocol starts with the commit phase and ends with the unveil phase.

The commit procedure:

1. Alice chooses a bit $b \in \{0, 1\}$.

2. Alice creates a random binary string $w = w_1 \cdots w_n$ with n bits.

3. If Alice wants to commit to 0, she does a quantum encoding of each bit w_i in the two basis states of the rectilinear basis $+$. If she wants to commit to 1, she encodes the bits in the two basis states of the diagonal basis $\times$. Let θ_i denote the basis chosen for w_i.

4. Alice sends the sequence of n encoded quantum states to Bob.

5. Bob chooses a random measurement basis (rectilinear or diagonal) for each of the received quantum states, i.e., he chooses a string of random bases $\hat{\theta} = \hat{\theta}_1 \cdots \hat{\theta}_n \in \{+, \times\}^n$. He measures the ith state in the basis $\hat{\theta}_i$, and denotes the outcome by $\hat{w}_i$.

If we take a look at the two density matrices for the n states corresponding to $b = 0$ and $b = 1$, respectively, it is easy to see that they are the same, and equal to the identity matrix. Thus, Bob has no chance to get any information about the bit b.

The unveil procedure:

1. Alice publishes b (i.e., the basis that she used for encoding) and the string w.

2. For about half of the n states, Bob used the same basis for his measurement as Alice used for encoding. In these cases Bob can verify that Alice's revealed bits are matching his measurement results.

How could a dishonest party cheat in this protocol? For example, Alice could choose the bit $b = 1$ for the commit phase, so she encodes the states with the diagonal basis $\times$. Later during the unveil phase, she changes her mind and tells Bob that she committed to the bit $b = 0$, so Bob assumes that Alice has used the rectilinear basis $+$. In approximately $n/2$ cases, Bob measures the states with the rectilinear basis $+$, and in these cases Alice has to guess the bits Bob measured. Since Alice's success to make a right guess for one bit is $1/2$, her overall cheating will not be detected with a probability of $(1/2)^{n/2}$. Once n is chosen large enough, Alice has practically no chance to manipulate the protocol by this probabilistic method.

But what if Alice uses specially entangled states as in the entanglement-based version of the BB84 protocol (see Section 4.2.1, Equation (4.9))? Alice could create n pairs of entangled states and send one part of each pair to Bob. She doesn't have to commit to a bit in the beginning, because she can perform a measurement right before the unveil phase. If, for example, she chooses bit $b = 0$, she measures the states that she has kept in the rectilinear basis $+$. Bob's measurement results will be perfectly correlated, due to the shape of the entangled state in Equation (4.9). If Alice wants to choose bit $b = 1$ instead, she measures the states that she has kept in the diagonal basis $\times$. As the state from Equation (4.9) is form-invariant under a basis rotation by $45°$, Alice's announced encoded states will again match Bob's measurement results. Thus, Bob has no chance to notice the attack.

5.2 Impossibility of Unconditionally Secure Quantum Bit Commitment

As mentioned above, unconditionally secure quantum bit commitment is impossible. In this section we will review the main arguments to prove this statement. According to Lo and Chau [LC97], the ideas of all quantum bit commitment protocols proposed up to date can be roughly described by the following five steps:

1. Alice chooses a bit $b \in \{0, 1\}$ and prepares the state

$$|0\rangle = \sum_i \alpha_i |e_i^A\rangle \otimes |f_i^B\rangle$$

for $b = 0$, and the state

$$|1\rangle = \sum_j \beta_j |e_j'^A\rangle \otimes |f_j'^B\rangle$$

for $b = 1$, where $|e_i^A\rangle$ and $|e_j'^A\rangle$ are orthonormal bases of Alice's Hilbert space, i.e., $\langle e_i^A | e_k^A \rangle = \delta_{ik}$ and $\langle e_j'^A | e_l'^A \rangle = \delta_{jl}$. The states $|f_i^B\rangle$ and $|f_j'^B\rangle$ live in Bob's Hilbert space, and are not necessarily orthogonal to each other.

2. Now, Alice has to make a measurement on the first part of the above state, and will thus determine i or j, depending on her initial choice for b.

3. Alice sends the second part of the above state to Bob. This is the last step in the commit phase.

4. At the beginning of the unveil phase, Alice publicly announces i or j together with b.

5. Bob makes a measurement on his part of the state, in order to make sure that in Step 3, Alice committed to the same bit she has announced in Step 4.

To show that a cheating Alice cannot be detected, we distinguish two cases. We give only a sketch of the proof, for more details we refer to [May97,LC97].

In the first case, Bob cannot get any information about the bit b out of the state that Alice sent him. This means that his two reduced density matrices, corresponding to the two states $|0\rangle$ and $|1\rangle$, are the same, i.e., $\mathrm{tr}_A|0\rangle\langle 0| = \mathrm{tr}_A|1\rangle\langle 1|$. Thus, the fidelity[7] of the two states is $F = 1$. Now, we can write the Schmidt decomposition (i.e., a bi-orthogonal decomposition that can always be found, see, e.g., Nielsen and Chuang [NC00]) as

$$|0\rangle = \sum_k \sqrt{\lambda_k}|\hat{e}_k^A\rangle \otimes |\hat{f}_k^B\rangle$$

and

$$|1\rangle = \sum_k \sqrt{\lambda_k}|\hat{e}_k^{\prime A}\rangle \otimes |\hat{f}_k^B\rangle,$$

where $|\hat{e}_k^A\rangle$ and $|\hat{e}_k^{\prime A}\rangle$ are orthonormal bases of Alice's Hilbert space, and $|\hat{f}_k^B\rangle$ is an orthonormal basis of Bob's Hilbert space. The λ_k's are the eigenvalues of Bob's two reduced density matrices corresponding to $|0\rangle$ and $|1\rangle$ (which are identical). There always exists a unitary transformation U that maps an orthonormal basis $|\hat{e}_k^A\rangle$ of a Hilbert space to another orthonormal basis $|\hat{e}_k^{\prime A}\rangle$ of the same Hilbert space, and thus this *local* unitary transformation (a rotation on Alice's side only) can map $|0\rangle$ to $|1\rangle$.

Therefore, Alice can start her commit phase with the bit $b = 0$. She prepares the state $|0\rangle$, skips the measurement (delays until Step 4) and sends Bob's part of the state $|0\rangle$ directly to Bob. At the beginning of the unveil phase, Alice has to choose the value b. If she chooses $b = 0$, she can proceed with the original protocol honestly. If she chooses $b = 1$, she can execute the unitary transformation U, and switch $|0\rangle$ to $|1\rangle$. Because $F = 1$, Bob has no chance to detect the cheating.

In the second case, let $F \neq 1$, i.e., the two reduced density matrices of Bob, corresponding to the two states $|0\rangle$ and $|1\rangle$, are not the same. The fidelity must be close to 1; otherwise, Bob could easily distinguish between the bits 0 and 1, and so he could cheat. Alice can again use her cheating strategy from above. Mayers [May97] has shown that with a cheating Alice, the probability of Bob being able to distinguish between 0 and 1 will not be larger. Thus Alice can cheat again with a probability close to 1.

As we can see, a dishonest party uses the same algorithm as an honest party. Hence it is impossible for the honest party to detect the cheater, and thus secure quantum bit commitment is not possible.

[7] Note that the fidelity F for two mixed states ρ_1 and ρ_2 is defined as $F = \mathrm{tr}\sqrt{\sqrt{\rho_1}\rho_2\sqrt{\rho_1}}$.

6 Outlook and Conclusions

The security of quantum key distribution relies on the inviolable laws of quantum mechanics: nonorthogonal quantum states are used as signal states in the BB84 protocol. The impossibility of perfect cloning of nonorthogonal states implies the security of this protocol.

In the security proof for the BB84 protocol, we have employed an equivalent entanglement-based protocol. The main idea is that local measurements on a maximally entangled state, shared by Alice and Bob, have perfectly correlated outcomes that can be used as the key. A maximally entangled state is necessarily pure, and a pure state cannot be entangled with an eavesdropper's state—thus Eve cannot learn anything about the key. The idea for quantum cryptography with entangled states goes back to Artur Ekert [Eke97], who suggested to confirm the existence of quantum correlations in the state of Alice and Bob by a Bell inequality test.

6.1 Other Quantum Key Distribution Protocols

A variety of quantum key distribution protocols can be found in the literature. All known prepare-and-measure schemes can be seen as variations of the BB84 protocol, which are obtained by changing the number and/or dimension of the quantum states.

In 1992, Charles Bennett [Ben92] proposed a protocol—which now is named after him the B92 protocol—in which only two nonorthogonal states are used. In the so-called six-state protocol [Bru98,BPG99], the six eigenstates of the three Pauli operators are used. This protocol has a lower efficiency compared with BB84, as in only one third of the cases Alice and Bob use the same basis, but it is more difficult for Eve to retrieve any information, thus the security is enhanced.

In this paper, we have always considered qubits, i.e., two-level systems as information carriers. What happens if one considers higher-dimensional systems, such as qutrits (three-level systems)? Intuitively, one would expect that the increased number of degrees of freedom makes it more difficult for Eve to extract information on the key. As proven in [BM02], higher-dimensional systems indeed offer increased security.

A recently suggested protocol [SARG04] introduces a new sifting method: rather than announcing the basis, Alice gives Bob a list of two nonorthogonal states from which the signal state was taken. This protocol has certain security advantages that are connected with experimental implementations of quantum cryptography.

6.2 Experimental Status

So far, we have presented quantum key distribution in a rather theoretical, abstract manner. What is the experimental situation—can the ideas of quantum cryptography be made reality? In recent years, much effort has been devoted to experiments on quantum cryptography, and much progress has been made. In most experiments, polarized photons are representing the qubits: photons are polarized if their electromagnetic field oscillates in a fixed direction of space (which has to be orthogonal to the direction of flight). The two degrees of freedom for a photonic qubit can be, e.g., horizontal and vertical polarization (the rectilinear basis in the BB84 protocol), or polarization rotated by 45° with respect to the horizontal/vertical direction—this corresponds to the diagonal basis in BB84. The experimentalist "only" has to produce single polarized photons on demand.

This, however, is one of the main experimental challenges: an attenuated laser pulse consists of Poisson-distributed number states, i.e., with a certain probability more than one photon will be emitted. These events with more than one photon allow for a dangerous eavesdropping strategy, the so-called photon-number splitting attack, where Eve splits off a photon and receives full information about the key. Apart from experimental progress towards true single-photon sources, new algorithms that can cope with this sort of attack have been developed. One example, the protocol by Scarani et al. [SARG04], has already been mentioned above.

The long-term goal in experimental quantum key distribution is to reach high key rates over large distances. For the transmission of photons, two possibilities exist: either transmission via optical fibers, or transmission in free space. Rather than trying to summarize all existing experiments, let us mention just two examples. A very stable, robust system with optical fiber transmission has been developed by Gisin and Zbinden at the University of Geneva, see [GRTZ02]. They were able to transmit a secret key from Geneva to Lausanne (i.e., over a distance of about 67 km), with a rate of 130 bit/s. Regarding free space quantum cryptography, Weinfurter from LMU Munich [KZH$^+$02] recently demonstrated secret key exchange over about 23.4 km (in the Alps, from Zugspitze to Karwendelspitze), with a rate of about 1000 bit/s.

Long-term goals of quantum key distribution are the realistic implementation via fibers, e.g., for different buildings of a bank or company (with a relatively small distance), and free space key exchange via satellites. Future practical developments will have to prove which one of the described protocols will turn out to be successful. At the moment, demonstrators for long-range quantum key distribution are being built within the EU project SECOQC (for further information, see www.secoqc.net). Quantum cryptography already provides the most advanced technology of quantum information science, and is on the way to achieve the (quantum) jump from university laboratories to the real world.

References

[BB84] C. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems, and Signal Processing*, pages 175–179. IEEE Press, 1984.

[BCJL93] G. Brassard, C. Crépeau, R. Jozsa, and D. Langlois. A quantum bit commitment scheme provably unbreakable by both parties. In *Proceedings of the 34th IEEE Symposium on Foundations of Computer Science*, pages 362–371. IEEE Computer Society Press, 1993.

[Ben92] C. Bennett. Quantum cryptography using any two nonorthogonal states. *Physical Review Letters*, 68:3121–3124, 1992.

[BEZ00] D. Bouwmeester, A. Ekert, and A. Zeilinger. *The Physics of Quantum Information*. Springer-Verlag, 2000.

[BM02] D. Bruß and C. Macchiavello. Optimal eavesdropping in cryptography with three-dimensional quantum states. *Physical Review Letters*, 88:127901(1)–127901(4), 2002.

[BPG99] H. Bechmann-Pasquinucci and N. Gisin. Incoherent and coherent eavesdropping in the six-state protocol of quantum cryptography. *Physical Review A*, 59:4238–4248, 1999.

[Bru98] D. Bruß. Optimal eavesdropping in quantum cryptography with six states. *Physical Review Letters*, 81:3018–3021, 1998.

[DH76] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, 1976.

[Eke97] A. Ekert. Quantum cryptography based on Bell's theorem. *Physical Review Letters*, 67:661–663, 1997.

[ElG85] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Transactions on Information Theory*, IT-31(4):469–472, 1985.

[GRTZ02] N. Gisin, G. Ribordy, W. Tittel, and Hugo Zbinden. Quantum cryptography. *Rev. Mod. Phys.*, 74:145–195, 2002.

[Hol73] A. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problems of Information Transmission (Problemy Peredachi Informatsii)*, 9:177–183, 1973.

[HP03] W. Huffman and V. Pless. *Fundamentals of Error-Correcting Codes*. Cambridge University Press, 2003.

[HPR01] L. Hemaspaandra, K. Pasanen, and J. Rothe. If P $\neq$ NP then some strongly noninvertible functions are invertible. In *Proceedings of the 13th International Symposium on Fundamentals of Computation Theory*, pages 162–171. Springer-Verlag *Lecture Notes in Computer Science #2138*, August 2001.

[HR99] L. Hemaspaandra and J. Rothe. Creating strong, total, commutative, associative one-way functions from any one-way function in complexity theory. *Journal of Computer and System Sciences*, 58(3):648–659, June 1999.

[HRS05] L. Hemaspaandra, J. Rothe, and A. Saxena. Enforcing and defying associativity, commutativity, totality, and strong noninvertibility for one-way functions in complexity theory. In *Proceedings of the Ninth Italian Conference on Theoretical Computer Science*, pages 265–279. Springer-Verlag *Lecture Notes in Computer Science #3701*, October 2005.

[KZH$^+$02] C. Kurtsiefer, P. Zarda, M. Halder, H. Weinfurter, P. Gorman, P. Tapster, and J. Rarity. A step towards global key distribution. *Nature*, 419:450, 2002.

[LC97] H. Lo and H. Chau. Is quantum bit commitment really possible? *Physical Review Letters*, 78:3410–3413, 1997.

[May95] D. Mayers. The trouble with quantum bit commitment. Technical Report quant-ph/9603015v3, Computing Research Repository (CoRR), 1995. Available on-line at http://arxiv.org/pdf/quant-ph/9603015.

[May97] D. Mayers. Unconditionally secure quantum bit commitment is impossible. *Physical Review Letters*, 78:3414–3417, 1997.

[NC00] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[Rot05] J. Rothe. *Complexity Theory and Cryptology. An Introduction to Cryptocomplexity*. EATCS Texts in Theoretical Computer Science. Springer-Verlag, Berlin, Heidelberg, New York, 2005.

[RS97] M. Rabi and A. Sherman. An observation on associative one-way functions in complexity theory. *Information Processing Letters*, 64(5):239–244, 1997.

[RSA78] R. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signature and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.

[SARG04] V. Scarani, A. Acín, G. Ribordy, and N. Gisin. Quantum cryptography protocols robust against photon number splitting attacks for weak laser pulse implementations. *Physical Review Letters*, 92:057901(1)–057901(4), 2004.

[Sch35] E. Schrödinger. Die gegenwärtige Situation in der Quantenmechanik. *Die Naturwissenschaften*, 23:807–812, 823–828, 844–849, 1935.

[Sha49] C. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, 28(4):657–715, 1949.

[Sho97] P. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[Sin99] S. Singh. *The Code Book. The Science of Secrecy from Ancient Egypt to Quantum Cryptography*. Fourth Estate, London, 1999.

[SP00] P. Shor and J. Preskill. Simple proof of security of the BB84 quantum key distribution protocol. *Physical Review Letters*, 85:441–444, 2000.

[Sti05] D. Stinson. *Cryptography: Theory and Practice*. CRC Press, Boca Raton, third edition, 2005.

[Wer89] R. Werner. Quantum states with Einstein–Podolsky–Rosen correlations admitting a hidden-variable model. *Physical Review A*, 40(8):4277–4281, October 1989.

[WZ82] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299:802–803, 1982.

[Yao95] A. Yao. Security of quantum protocols against coherent measurements. In *Proceedings of the 27th ACM Symposium on Theory of Computing*, pages 67–75. ACM, New York, 1995.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/