Found Phys (2010) 40: 494–531
DOI 10.1007/s10701-010-9408-4

Quantum Cryptography

Serge Fehr

Received: 6 March 2009 / Accepted: 7 January 2010 / Published online: 20 January 2010

Abstract Quantum cryptography makes use of the quantum-mechanical behavior of nature for the design and analysis of cryptographic schemes. Optimally (but not always), quantum cryptography allows for the design of cryptographic schemes whose security is guaranteed solely by the laws of nature. This is in sharp contrast to standard cryptographic schemes, which can be broken in principle, i.e., when given sufficient computing power. From a theory point of view, quantum cryptography offers a beautiful interplay between the mathematics of adversarial behavior and quantum information theory. In this review article, we discuss the traditional application of quantum cryptography, *quantum key distribution* (QKD), from a modern perspective, and we discuss some recent developments in the context of quantum *two-party cooperation* (2PC). QKD allows two distant parties to communicate in a provably-secure way in the presence of an outside eavesdropper, whereas 2PC is concerned with protecting information against possibly malicious insiders. We show the basic idea of constructing quantum cryptographic schemes, but we also show some connections to quantum information theory as needed for the rigorous security analyses, and we discuss some of the relevant quantum-information-theoretic results.

Keywords Quantum cryptography · Quantum information theory · Hilbert space formalism · Key distribution · Secure cooperation

1 Introduction

CRYPTOGRAPHY aims at providing tools for securing private information and preventing critical information-processing operations from adversarially provoked malfunction. These are very crucial objectives in today's society where information plays

S. Fehr (✉)
Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
e-mail: Serge.Fehr@cwi.nl

a fundamental role. As such, great effort is put into designing and implementing cryptographic schemes that offer privacy-protecting solutions for various tasks. Whereas traditional cryptography is concerned with *secure communication*, i.e., the transmission of private messages over a (potentially) insecure communication channel, with the advent of widespread electronic communication, new cryptographic tasks have become increasingly important. We would like to be able to do e-voting, on-line auctions, Internet gambling, data-mining etc. in a secure way. These tasks involve parties with different and possibly conflicting interests, and we want that the correctness of the outcome is guaranteed while at the same time the privacy of individual users remains protected.

However, the security of most of the cryptographic schemes currently used relies on *unproven* computational complexity assumptions (like the assumed hardness of factoring large numbers), combined with an assumed bound on a potential attacker's computing power. This complexity-theoretic approach of designing cryptographic schemes leads to very practical solutions but obviously has its downside: one cannot be fully certain about the security of the scheme! Indeed, the underlying computational complexity assumption might be broken from one day to another (e.g. by an efficient factoring algorithm being discovered) since complexity theory is still far from being able to prove some computational problem to be "hard" in the sense as needed. Furthermore, it is known that the standard complexity assumptions used in practice (factoring and computing discrete-logs) break down as soon as a quantum computer can be built. Finally, even if it *is* computationally infeasible for an attacker to extract sensitive data from the information available to him *at the time the cryptographic scheme is used*, the attacker can still store, say, an intercepted ciphertext and wait until computer technology has advanced enough so that he eventually *can* recover the data that was to be protected. This clearly poses a serious threat to long-term highly-sensitive data.

QUANTUM CRYPTOGRAPHY offers a beautiful approach to overcome the above drawbacks. The idea behind quantum cryptography is to make use of the quantum-mechanical behavior of nature for the design and analysis of cryptographic schemes that do not have to rely on unproven complexity assumptions. This adventurous approach goes back to ideas by Wiesner from the late sixties, but they were unnoticed for about a decade. Optimally, but not always, quantum cryptography allows for the design of cryptographic schemes that can be proven secure under the sole assumption that the laws of quantum mechanics are correct—or that they at least describe sufficiently well the behavior of certain particles like photons or spin-$\frac{1}{2}$ particles, which would be used to implement the quantum-cryptographic schemes.

However, quantum cryptography not merely uses the theory of quantum mechanics as a tool box, but rather there is a fruitful interplay between the design and analysis of quantum cryptographic schemes and the development of the information-theoretical understanding of quantum mechanics. For instance, the search for a rigorous analysis of one of the first quantum cryptographic schemes led to important insights into quantum information theory, which in turn proved to be useful for the design of new quantum cryptographic schemes.

Let us give some intuition why quantum mechanical effects could indeed prove useful for designing cryptographic schemes. Consider two parties, called *Alice* and

Bob, who can produce and exchange quantum-mechanical particles, for instance single polarized photons or spin-$\frac{1}{2}$ particles. However, we assume that this quantum communication is under the control of an attacker *Eve*. Thus, Eve has full access to the exchanged particles. Nevertheless, the laws of quantum mechanics restrict the information accessible to Eve! Indeed, Heisenberg's uncertainty principle, and its extension by Robertson, guarantees that Eve *cannot* obtain full information on the state of the transmitted particles: if Eve measures the linear polarization of a photon then its circular polarization becomes unpredictable and vice versa, and, similarly, if Eve measures the spin of a spin-$\frac{1}{2}$ particle along one axis then its spin along either of the other two axes becomes unpredictable. This not only means that Eve may get at most limited information on the state of any exchanged particle, but even more importantly, any attempt to obtain information inevitably causes the state of the particle to *change*; if cleverly set-up, this can be detected by Alice and Bob, so that they can abort before any harm is done. This is in sharp contrast to classical means of communication (e.g. over the phone or the Internet) where in principle an eavesdropper can listen into the conversation without actively affecting it, and thus without any chance of being detected.

IN THIS ARTICLE, one the one hand, we would like to give the basic intuition behind the design and the security of quantum cryptographic schemes. As the reader will see, the quantum cryptographic schemes we show and the intuitive reasoning why they should be secure are rather simple and can be appreciated even by laymen with a very limited (and possibly wrong) understanding of quantum mechanics.

On the other hand, we also want to present quantum cryptography as an exact mathematical science that combines elements from classical cryptography, information theory and quantum mechanics. Therefore, besides the quantum-cryptographic schemes we show, we also discuss the theoretical foundations needed to rigorously understand and prove their security. These are quantum-information-theoretic results, specifically developed for the analysis of quantum-cryptographic schemes, but can be appreciated in their own right as providing interesting insight into the theory of quantum information. For instance, we show a meaningful way to measure the uncertainty that some piece of classical (meaning non-quantum) data contains when given a correlated quantum state, and we show that this measure determines the number of nearly-random-and-independent bits that can be extracted from the classical data. Also, we show a variant of the uncertainty principle that expresses the amount of uncertainty in terms of the above measure.

As of specific quantum cryptographic results, we focus in this article on the question of tackling classical (i.e. non-quantum) cryptographic tasks by quantum-cryptographic means, like how to securely communicate a classical private message by using a quantum channel. Specifically, we focus on *quantum-key distribution* (QKD), which is the traditional application of quantum cryptography, and on recent new developments in the context of quantum *two-party cooperation* (2PC).

QKD allows two parties, Alice and Bob, to agree on a secret key K by *public communication*, i.e., even if an attacker Eve can access the complete conversation between Alice and Bob. By the laws of quantum mechanics, it is guaranteed that the agreed-upon secret key K is (close to) random-and-independent of Eve's (quantum) view. As such, K can then be safely used for instance as encryption key for a (possibly perfectly-secure) encryption scheme to securely communicate a private message via the public communication channel.

2PC, on the other hand, is concerned with protecting information against *inside* attackers. Unfortunately, quantum cryptographic 2PC schemes whose security is guaranteed by the laws of quantum mechanics *alone* do not exist (unless one settles for a very low level of security), but in addition, some "technological restriction" needs to be assumed about the attacker: for instance, that he *cannot* reliably store arbitrarily many, say, photons without affecting their polarization. While the theory of quantum physics permits to store quantum states, doing so in the form of photons, for instance, is technically very challenging and essentially impossible with current technology. It is thus reasonable to base security upon it.

Another direction of quantum cryptography, which is not covered here, is to specify and study *quantum*-cryptographic tasks, like how to encrypt or authenticate a quantum state; this direction leads to questions and results that are interesting from a theoretical point of view but so far lack a practical significance. On the other hand, there is promising progress in the development of the technology needed to actually implement the quantum cryptographic schemes discussed in this article, with actual devices already being sold on the market. Nevertheless, this article is of theoretical nature and does not discuss implementational issues; for a more practical-oriented treatment of the topic, we refer to the excellent review article by Gisin, Ribordy, Tittel and Zbinden [31].

THE STRUCTURE of the article is as follows. The upcoming Sect. 2 provides some information on the history of quantum cryptography as of interest for the topics covered in this article, and in Sect. 3 we introduce the notation that we use throughout. In Sect. 4 we discuss and construct schemes for QKD, and in Sect. 5 we develop the tools we then use to rigorously prove security of the QKD schemes in Sect. 6. Finally, in Sect. 7 we discuss the recent developments of quantum cryptography in the context of 2PC, and we conclude in Sect. 8.

2 A Brief History of Quantum Cryptography

The history of quantum cryptography starts off in 1970 when Stephen Wiesner wrote *Conjugate Coding*. In this highly innovative article, he explains how in principle the laws of quantum mechanics can be used to produce bank notes that would be impossible to counterfeit, and how to implement a *multiplexing channel*, a notion that was re-invented more than 10 years later under the name of *oblivious transfer* [30, 48]. However, Wiesner's manuscript was not accepted for publication. Fortunately, Wiesner knew Charles H. Bennett quite well and told him about his work; otherwise his pioneering ideas might have been lost forever. In the subsequent years, Bennett mentioned Wiesner's work to various people, but without raising anyone's interest.

Quantum cryptography was revived in 1979 when Bennett approached Gilles Brassard and explained to him Wiesner's approach to use quantum mechanics in order to design unforgeable banknotes. Brassard was very excited about such an approach, and they combined the (at this time new) concept of public-key cryptography with Wiesner's quantum approach,[1] resulting in a *Crypto 82* paper by Bennett, Brassard, Breidbart and Wiesner [8], which coined the term *quantum cryptography*. This also brought Wiesner's manuscript back to life, and it was subsequently published in *Sigact News* [60].

At this point in time, quantum cryptography was considered pure science fiction, because the technology required to implement the suggested schemes was (and actually still is) out of reach. For instance the proposed unforgeable bank notes require to store a single polarized photon or spin-$\frac{1}{2}$ particle for days without significant absorption or loss of polarization. As such, quantum cryptography was considered to be doomed from the start as being unrealistic.

This changed when, as Brassard has expressed it in [17], "we [Bennett and Brassard] realized [...] that God had meant photons to travel rather than to stay put!" Although it has to be said that already Wiesner's multiplexing scheme was based on traveling photons, with no need for storing them. Driven by this motivation, Bennett and Brassard started to look for quantum cryptographic schemes that were based on the *transmission* of quantum states via a quantum channel. They first came up with a one-time-pad-like quantum encryption scheme that allows the key to be *re-used* [7]; the scheme, however, was still not very practical. They submitted this result to several major theoretical-computer-science conferences, but failed to get it accepted.

In 1983, Bennett and Brassard abandoned their quantum encryption scheme when they realized that it would be much simpler to use the quantum channel to securely transmit a *random key*, rather than the actual message to be securely communicated. And once the key is securely transmitted, it can then be used to one-time-pad encrypt and securely communicate the actual message in a standard way. *Quantum key distribution* (QKD) was born! Their new finding got accepted to an information-theory conference in 1983 [5]; however, this conference only published one-page abstracts. Shortly after, Brassard was invited to present and publish a paper on a topic of his choice at the 1984 *IEEE International Conference on Computers, Systems, and Signal Processing*, which took place in India. Having experienced how hard it was at that time to get these kinds of results published, Brassard took the opportunity to publish the full description of their QKD scheme [6], which then became known as the *BB84* QKD scheme.

We note that at this point in time, the BB84 QKD scheme could at best be proven secure against "feasible" *individual* attacks, where the attacker, Eve is assumed to interact with each communicated photon individually, but BB84 was *conjectured* to be secure against *general* attacks that are only restricted by the laws of quantum physics, and where for instance Eve may interact with the communicated photons *collectively*.

Quantum cryptography was also picked up by other researchers, for instance by Claude Crépeau, and a lot of effort was put into designing quantum-cryptographic

[1] From today's perspective, it looks odd to mesh public-key cryptography, which inherently can only be *computationally* secure, with quantum cryptography, whose goal is to obtain security guaranteed by the laws of nature.

schemes for other cryptographic tasks. In particular, a lot of effort was put into trying to design schemes for *bit-commitment* (BC) and for *oblivious transfer* (OT), two important building blocks for secure *2-party cooperation* (2PC) [10, 16, 18]. Also for those schemes, security could be argued only for individual attacks, and security against general attacks was typically conjectured.

However, in the following years, all the proposed schemes for BC and OT got eventually broken by sophisticated quantum attacks. And in 1996, it was then proven by Mayers and independently by Lo and Chau that bit commitment, and essentially any interesting 2PC (including OT), *cannot* be implemented by means of a quantum-cryptographic scheme with security only relying on the correctness of quantum mechanics [41, 42, 46]. This negative result came as a shock for the quantum cryptography community. Not only was the belief shattered that quantum cryptography could provide unconditionally-secure solutions for any reasonable cryptographic problem, but since the BB84 scheme was still not rigorously proven to withstand sophisticated quantum attacks, also the confidence in QKD was undermined. As a result, in the subsequent years, little work was done in the context of secure 2PC,[2] and a lot of effort was put into proving QKD unconditionally secure.

In the meantime, some variants of the original BB84 QKD scheme had been proposed. Most notably is the scheme by Ekert [29], which is based on entangled particles (like so-called EPR pairs [28]) and on Bell's theorem [3], and its modification due to Bennett, Brassard and Mermin, which avoids the use of Bell's theorem and was shown to be equivalent to the original BB84 QKD scheme from a security point of view. Although technically more challenging to implement, entanglement-based QKD schemes play an important role because they provide a convenient handle for proving QKD schemes secure against general attacks.

The very first QKD security proofs (for BB84) against general attacks were given by Mayers [45, 47] and, subsequently, by Biham, Boyer, Boykin, Mor and Roychowdhury [14]. However, their security proofs were very complicated and have only been reluctantly accepted. Lo and Chau proposed a security proof that was easier to understand, but was for a new entanglement-based QKD scheme that required the honest participants of the scheme to have quantum computers. It was then up to the seminal work of Shor and Preskill in 2000 [57], more than 15 years after the invention of QKD by Bennett and Brassard, to give a fully-satisfactory security proof against general attacks for the original BB84 QKD scheme.[3]

[2] There was some work on BC and OT secure against *computationally-bounded* quantum attacks; however, this approach goes somewhat against the main motivation of quantum cryptography, which is to avoid relying on the assumed hardness of some computational problem.

[3] Another reason why, from today's perspective, the early security proofs by Mayers and Biham et al. are not fully satisfactory, is that they implicitly assume the adversary Eve to *measure* all her information at the end of the execution of the QKD scheme. This was later realized (see e.g. [39]) to cause a lack of composability, meaning that even though the QKD scheme is secure when executed in isolation, it may actually become insecure as soon as the key is used in another application (and what's the point in producing a key when it cannot be used?). Although the original proof by Shor and Preskill also makes this implicit assumption, it does not crucially rely on it, and a security proof that does imply composability can be obtained by obvious modifications.

In the recent years, our understanding of the security of BB84 and other QKD schemes has significantly increased, mainly due to new insights into *quantum information theory*, put forward to a great deal by the work of Renner [21, 49–51]. In a sequence of works, he showed that for typical QKD schemes, security against general attacks follows "for free" from security against individual attacks, which is much easier to prove.

In the meantime, the problem of designing quantum-cryptographic schemes for 2PC tasks was picked up again. Clearly, there was no hope to construct fully-fledged unconditionally-secure quantum-cryptographic scheme: the impossibility result by Mayers and Lo and Chau implies that for every candidate scheme there exists an attack that breaks the aspired security of the scheme. However, even though such attacks exist *in principle*, they would be hard to execute *in practice* as they typically involve perfect storage of large quantum states. The technological hardness of launching these attacks had been folklore knowledge for some time, but no one had given it much attention, until 2005 when Damgård, Fehr, Salvail and Schaffner realized its potential. They were able to design quantum-cryptographic schemes for certain 2PC tasks for which they could prove that *any* attack that would break security necessarily must involve large quantum-storage capacities. In a sequence of works [23–26], they showed the existence of practical quantum-cryptographic schemes for a variety of 2PC tasks, provably secure in the above sense in the *bounded-quantum-storage model*.

The success of this approach gave new life to the problem of designing quantum-cryptographic schemes for 2PC tasks, after the set-back in the late nineties, and it motivated other researchers to look for extensions and alternatives, like the *noisy*-quantum-storage model, where the dishonest participant has the ability to store all the communicated photons, but the storage is assumed to be noisy [59].

3 Notation and Basic Concepts

We assume the reader to be familiar with the basic concepts of quantum mechanics and with its Hilbert-space formalism. For completeness, and since the view we take and the terminology and notation we use might be slightly different from what the reader is used to, we briefly recall the basic concepts of quantum mechanics as understood from a quantum-information-processing point of view.

3.1 Dirac's Bra-ket Notation

Let $\mathcal{H}$ be a complex Hilbert space. We use *Dirac's bra-ket* notation as commonly used in quantum physics. This means vectors in $\mathcal{H}$ are denoted as *ket*'s $|\cdot\rangle$, and for any $|\varphi\rangle \in \mathcal{H}$, the corresponding *bra*-vector is defined as the linear functional $\langle\varphi| : \mathcal{H}\to\mathbb{C}$ that maps $|\psi\rangle \in \mathcal{H}$ to the *inner product* of $|\varphi\rangle$ and $|\psi\rangle$, which is denoted as $\langle\varphi|\psi\rangle$; hence, by definition, $\langle\varphi||\psi\rangle = \langle\varphi|\psi\rangle$. Furthermore, for $|\varphi\rangle, |\psi\rangle \in \mathcal{H}$, the *outer product* of $|\varphi\rangle$ and $|\psi\rangle$ is defined as the linear function $|\varphi\rangle\langle\psi| : \mathcal{H}\to\mathcal{H}$ that maps $|\eta\rangle \in \mathcal{H}$ to $|\varphi\rangle\langle\psi|\eta\rangle$; hence, by definition, $|\varphi\rangle\langle\psi||\eta\rangle = |\varphi\rangle\langle\psi|\eta\rangle$.

Throughout, we only consider *finite-dimensional* Hilbert spaces, so that we always may assume that $\mathcal{H} = \mathbb{C}^d$ for some (finite) dimension d, and any operator in $\mathrm{End}(\mathcal{H})$

is bounded and can be thought of as a $(d \times d)$-matrix with entries in $\mathbb{C}$. Also, a vector $|\varphi\rangle \in \mathcal{H}$ can be thought of as a *column* vector with entries $a_1, a_2, \ldots, a_d \in \mathbb{C}$ and $\langle\varphi|$ as the corresponding transpose complex-conjugate *row* vector $|\varphi\rangle^\dagger = (\bar{a}_1, \ldots, \bar{a}_d)$, and $\langle\varphi||\psi\rangle$, $|\varphi\rangle\langle\psi|$ etc. can be understood as matrix multiplication.

Finally, for two (and similarly for more) vectors $|\varphi\rangle \in \mathcal{H}$ and $|\psi\rangle \in \mathcal{H}'$, we often write $|\varphi\rangle|\psi\rangle$ as well as $|\varphi, \psi\rangle$ (or even $|\varphi\,\psi\rangle$) as a short hand for the tensor product $|\varphi\rangle \otimes |\psi\rangle \in \mathcal{H}\otimes\mathcal{H}'$ of $|\varphi\rangle$ and $|\psi\rangle$.

3.2 Quantum Systems

A (*quantum*) *system* is associated with a complex Hilbert space $\mathcal{H}$, the *state space* of the system, and the *state* of the system is described by a positive semi-definite operator $\rho \in \mathrm{End}(\mathcal{H})$ with trace $\mathrm{tr}(\rho) = 1$. Such an operator is called *density* operator (or matrix). We write $\mathcal{D}(\mathcal{H})$ for the set of all density operators $\rho \in \mathrm{End}(\mathcal{H})$, and we write $\rho \geq 0$ to express that the operator ρ is positive semi-definite. We typically identify a quantum system by an abstract name, e.g. A, and then by default denote the state space of A by $\mathcal{H}_A$ and the density matrix describing the state of A by ρ_A.

A quantum state is *pure* if its density matrix $\rho \in \mathcal{D}(\mathcal{H})$ has rank 1, which is equivalent to saying that there exists $|\varphi\rangle \in \mathcal{H}$ such that $\rho = |\varphi\rangle\langle\varphi|$, where the trace condition on ρ implies that $|\varphi\rangle$ is *normalized*, i.e., $\||\varphi\rangle\|^2 = \langle\varphi|\varphi\rangle = 1$. In case of a pure state $\rho = |\varphi\rangle\langle\varphi|$, we may also use the *state vector* $|\varphi\rangle$ to describe the state.

From a geometric point of view, the pure states are given by the extremal points of the convex set $\mathcal{D}(\mathcal{H})$, in particular, any $\rho \in \mathcal{D}(\mathcal{H})$ can be written as a convex-linear-combination $\rho = \sum_{\ell=1}^{L} \varepsilon_\ell |\varphi_\ell\rangle\langle\varphi_\ell|$ (i.e. $\varepsilon_1, \ldots, \varepsilon_L \geq 0$ and $\sum_\ell \varepsilon_\ell = 1$) of pure states. Such a system can alternatively be understood to be in pure state $|\varphi_\ell\rangle$ with probability ε_ℓ.

The state space of the *joint* quantum systems AB, which consist of two (or more) subsystems A and B, is given by the tensor product $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$ of $\mathcal{H}_A$ and $\mathcal{H}_B$. If the state of the joint system is given by ρ_{AB}, then the state of the subsystem A when considered as a "stand alone" system is given by the reduced density matrix $\rho_A = \mathrm{tr}_B(\rho_{AB}) \in \mathcal{D}(\mathcal{H}_A)$, where the *partial trace* $\mathrm{tr}_B : \mathrm{End}(\mathcal{H}_A \otimes \mathcal{H}_B) \to \mathrm{End}(\mathcal{H}_B)$ is the (well-defined) linear operator with $\mathrm{tr}_B(|\varphi_A\rangle\langle\psi_A| \otimes |\varphi_B\rangle\langle\psi_B|) = |\varphi_A\rangle\langle\psi_A|\,\mathrm{tr}(|\psi_B\rangle\langle\varphi_B|) = |\varphi_A\rangle\langle\psi_A|\langle\psi_B|\varphi_B\rangle$ for all $|\varphi_A\rangle, |\psi_A\rangle \in \mathcal{H}_A$ and $|\varphi_B\rangle$, $|\psi_B\rangle \in \mathcal{H}_B$. Similarly, the state of B is given by $\rho_B = \mathrm{tr}_A(\rho_{AB})$.

Here, as is common in quantum information processing, we consider the quantum state of a system to be *static*, meaning that it does not change over time, unless it is actively operated on. A quantum system A can be operated on by means of applying a *unitary* transformation $U \in \mathrm{End}(\mathcal{H}_A)$; as a result, the state $\rho_A \in \mathcal{D}(\mathcal{H}_A)$ of A evolves to the new state $\rho'_A = U\rho_A U^\dagger$. We write $\mathcal{U}(\mathcal{H})$ for the set of all unitary operators $U \in \mathrm{End}(\mathcal{H})$. In case of a pure state described by its state vector $|\varphi_A\rangle \in \mathcal{H}_A$, the state evolves as $|\varphi'_A\rangle = U|\varphi_A\rangle$.

The only way to gain information on the state of a quantum system A is by means of a *measurement*. A measurement is described by an *observable*, which is given by a (finite) collection $\{\Pi_i\}_{i\in I}$ of *orthogonal projections* $\Pi_i \in \mathrm{End}(\mathcal{H}_A)$ that satisfy the condition $\sum_i \Pi_i = \mathbb{I}_A$, where $\mathbb{I}_A$ denotes the identity in $\mathrm{End}(\mathcal{H}_A)$.[4,5] For quantum system A in state $\rho_A \in \mathcal{D}(\mathcal{H}_A)$, measuring A with respect to $\{\Pi_i\}_{i\in I}$ has the following effect. (1) An *outcome* $i \in I$ is observed, with the probability that a specific $i \in I$ is observed given by $p_i = \mathrm{tr}(\Pi_i \rho_A)$; and (2) after the measurement, the state ρ has *collapsed* to $\rho'_A = \Pi_i \rho_A \Pi_i{}^\dagger / p_i$ where i is the outcome observed. If A is part of a joint system AB, then measuring A with respect to $\{\Pi_i\}_{i\in I} \subset \mathrm{End}(\mathcal{H}_A)$ acts as measuring AB with respect to $\{\Pi_i \otimes \mathbb{I}_B\}_{i\in I}$.

We often consider measurements where the Π_i's are projections onto an *orthonormal basis* $\{|i\rangle\}_{i\in I}$ of $\mathcal{H}_A$: $\Pi_i = |i\rangle\langle i|$.[6] In this case, we say that A is measured *in basis* $\{|i\rangle\}_{i\in I}$. If the state of A is pure, given by state vector

$$|\varphi\rangle = \sum_{i\in I} \alpha_i |i\rangle$$

(where by the normalization condition $\sum_i |\alpha_i|^2 = 1$), then it follows from the above that measuring A in basis $\{|i\rangle\}_{i\in I}$ has the effect that $i \in I$ is observed with probability

$$p_i = |\alpha_i|^2,$$

and the state collapses to $|i\rangle$. Furthermore, if the state of a joint system AB is pure, given by state vector $|\varphi\rangle = \sum_i \alpha_i |i\rangle|\psi_i\rangle$ with normalized $|\psi_i\rangle \in \mathcal{H}_B$, then measuring A in basis $\{|i\rangle\}_{i\in I}$ has the effect that $i \in I$ is observed with probability $p_i = |\alpha_i|^2$, and the state collapses to $|i\rangle|\psi_i\rangle$.

To simplify the language, we will sometimes be somewhat sloppy in distinguishing between a quantum system, its state, and the density matrix or state vector describing the state. For instance, we may speak of "measuring a state ρ" when we actually mean that a system A whose state is given by the density matrix ρ is measured.

A *qubit* is a quantum system with state space $\mathcal{H} = \mathbb{C}^2$. $\{|0\rangle, |1\rangle\}$ denotes the *computational* basis $|0\rangle = \left[\begin{smallmatrix}1\\0\end{smallmatrix}\right]$ and $|1\rangle = \left[\begin{smallmatrix}0\\1\end{smallmatrix}\right]$ of $\mathbb{C}^2$ and $\{|+\rangle, |-\rangle\}$ the *Hadamard* basis

$$|+\rangle = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\1\end{bmatrix} = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big) \quad \text{and} \quad |-\rangle = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\-1\end{bmatrix} = \frac{1}{\sqrt{2}}\big(|0\rangle - |1\rangle\big).$$

Note that one can write $|+\rangle = H|0\rangle$ and $|-\rangle = H|1\rangle$, where H is the *Hadamard* transform $H = \left(\begin{smallmatrix}1 & 1\\ 1 & -1\end{smallmatrix}\right)/\sqrt{2}$. Thus, $H^b\{|0\rangle, |1\rangle\} = \{H^b|0\rangle, H^b|1\rangle\}$ denotes the computational basis if $b = 0$ and the Hadamard basis if $b = 1$. An n-qubit systems consists of n qubits, i.e., is a quantum system whose state space is the n-fold tensor product $(\mathbb{C}^2)^{\otimes n} = \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2$.

[4] Equivalently, an observable can be given by a Hermitian operator O in $\mathrm{End}(\mathcal{H}_A)$, such that the Π_i's are the projections into the eigenspaces and the i's (encoded as real numbers) the corresponding eigenvalues: $O = \sum_i i\Pi_i$.

[5] There actually exists a more general notion of measurements, so-called *POVM*'s; however, the *Von Neumann* (also known as *projective*) measurements considered here are sufficient for our purposes.

[6] Note that we are using the indices $i \in I$ as the "names" of the basis vectors; indeed we will often name basis vectors by numbers, like $\{|0\rangle, |1\rangle\}$, but the index set I may just as well consists of other "symbols".

The *trace distance* of two density operators $\rho, \sigma \in \mathcal{D}(\mathcal{H})$ is defined as $\delta(\rho, \sigma) := \frac{1}{2} \operatorname{tr}|\rho - \sigma|$, where $|\rho - \sigma|$ is the unique positive semi-definite square root of $(\rho - \sigma)(\rho - \sigma)^\dagger$. In other words, $\delta(\rho, \sigma) = \frac{1}{2}\sum_i |\lambda_i|$, where the λ_i's are the (not necessarily distinct) real eigenvalues of $\rho - \sigma$. One can show that for any physical processing, the two states ρ and σ behave in an indistinguishable way except with probability at most $\delta(\rho, \sigma)$. Thus, informally, if $\delta(\rho, \sigma)$ is very small then, without making a significant error, the quantum state ρ can be considered to be equal to σ.

3.3 Hybrid Systems: Combining Classical and Quantum Information

Consider a situation where the state of a quantum system E is *randomized*: with probability $P_X(x)$ system E is in state $\rho_{E|X=x} \in \mathcal{D}(\mathcal{H}_E)$, where X is a random variable with finite range $\mathcal{X}$ and P_X is its probability distribution (i.e. $P_X(x) = P[X{=}x]$ for any $x \in \mathcal{X}$). Such a situation occurs naturally when subsystem A of a joint system AE is measured in a basis $\{|x\rangle\}_{x\in\mathcal{X}} \subset \mathcal{H}_A$, where the random variable X then captures the observed value and $\rho_{E|X=x}$ denotes the state E collapses to when x is observed. Or, it occurs when an "experimenter" tosses some coins to determine x and then prepares system E to be in a state that depends on his choice x.

For an observer that only has access to system E but is ignorant of the value of the index x, the state of E is given by

$$\rho_E = \sum_x P_X(x)\rho_{E|X=x}.$$

By "encoding" the choice of x into a quantum state $|x\rangle$, where $\{|x\rangle\}_{x\in\mathcal{X}}$ is a fixed orthonormal basis of $\mathcal{H}_X = \mathbb{C}^{|\mathcal{X}|}$ (typically the canonical basis), and where "decoding" works by measuring in basis $\{|x\rangle\}_{x\in\mathcal{X}}$, we may understand the *hybrid* system XE, consisting of the random variable X and the quantum system E, as a joint *quantum* system XE whose state is given by the density matrix

$$\rho_{XE} = \sum_x P_X(x)|x\rangle\langle x| \otimes \rho_{E|X=x}.$$

We also say that the state ρ_{XE} has a *classical* X (with respect to $\{|x\rangle\}_{x\in\mathcal{X}}$). Finally, we write $\rho_X = \operatorname{tr}_E(\rho_{XE}) = \sum_x P_X(x)|x\rangle\langle x|$ for the random variable X understood as a quantum state. This formalism naturally extends to states that depend on several, possibly dependent, random variables X, Y etc. To simplify notation, we often write ρ_E^x instead of $\rho_{E|X=x}$.

The random variable X is *independent* of the quantum system E in that $\rho_E^x = \rho_E^{x'}$ (and thus $= \rho_E$) for all $x, x' \in \mathcal{X}$, if and only if

$$\rho_{XE} = \rho_X \otimes \rho_E.$$

This in particular implies that no information on X can be obtained from having access to the quantum system *E*. Similarly, X is *random-and-independent* of the state of *E* if and only if

$$\rho_{XE} = \mu_X \otimes \rho_E$$

where μ_X denotes the *completely mixed state* $\mu_X = \frac{1}{|\mathcal{X}|}\sum_x |x\rangle\langle x| = \frac{1}{|\mathcal{X}|}\mathbb{I}_X$ in $\mathcal{D}(\mathcal{H}_X)$. This is the situation we aim for in cryptography, where X is intended to be used as a cryptographic key and E collects the information the attacker has. Typically, one needs to allow a small "error probability" and has to settle for $\delta(\rho_{XE}, \mu_X \otimes \rho_E)$ being sufficiently small. By the properties of the trace distance, this then implies that no matter how X is used, it behaves like being perfectly random-and-independent of E except with small probability.

4 Quantum Key Distribution (QKD)

4.1 Problem Description

The classical problem in cryptography concerns *secure communication*. Consider two parties, named *Alice* and *Bob*, who are geographically separated but can communicate over a given communication channel. However, the communication channel is *public* in the sense that an *attacker*, named *Eve*, can read the complete communication that takes place over the channel.[7] How can Alice still communicate a message M to Bob in such a way that only Bob learns M but not Eve? How can Alice "scramble" M so that it looks like nonsense to Eve yet Bob can recover M?

Technically, this is done by means of an encryption function *enc*, which takes as input a *key* K and the message M, and which outputs a *ciphertext*: $C = enc(K, M)$. Not knowing the key K (but possibly *enc*), it should be impossible for Eve to obtain M (or any partial information on it) from C; on the other hand, Bob, who knows K, should be able to recover M from C by means of a suitable decryption function: $M = dec(K, C)$.

However, encryption does not fully solve the problem, it only *reduces* it; namely to the problem of Alice and Bob establishing a common key K that is secret to Eve. For instance the so-called *one-time-pad* encryption scheme[8] enjoys *perfect security* in the sense that C is statistically independent of M, but requires a fresh random key K, known to Alice and Bob but secret to Eve, for every new message M to be encrypted.

One approach to establish K would be to have Alice produce K and try to communicate it securely to Bob over the public channel, but then we are obviously back to our initial problem. Another approach is to look for a "physical" solution: for instance Alice and Bob could meet at a safe place to agree on K, or use a trusted courier to securely transfer K. However, these kinds of solutions are typically very inconvenient and not acceptable in many cases. It would be much more convenient if Alice and Bob could generate a common secret key K "on the fly" simply by communicating over the public channel. But can this be possible at all? Can Alice and Bob agree on a secret key when Eve can follow the whole conversation?

[7] It is irrelevant if the communication can easily be read by any outsider (like for radio broadcast), or if Eve needs to—but indeed does—possess sophisticated eavesdropping devices that allow her to listen into the conversation (like for e-mail).

[8] The one-time pad encrypts message $M \in \{0,1\}^\ell$ as $C = M \oplus K$ where K is a secret key in $\{0,1\}^\ell$ and $\oplus$ denotes bit-wise addition modulo 2.

If we relax the requirement that Eve should have no information on K and "only" require that it is *computationally infeasible* (but possible in principle) for her to compute (any information on) K, then, under certain unproven computational complexity assumptions, this can be done by means of *public-key cryptography* techniques, one of the greatest inventions of modern cryptography. However, without any breakthrough result in computational complexity theory and in particular without solving the famous $\mathbf{P} \overset{?}{=} \mathbf{NP}$ problem, such an approach needs to rely on an unproven computational complexity assumption (like the conjectured hardness of factoring large integers), which we want to avoid.

What if we do not want to rely on unproven computational complexity assumptions and want K to be secret in an information-theoretic sense? Is it still possible for Alice and Bob to agree on a secret key K by public communication? Surprisingly, the answer is still yes: by means of a so-called *quantum key distribution* (QKD) scheme. Such a scheme makes clever use of the quantum-mechanical behavior of some particles, like photons, in order to allow Alice and Bob to jointly produce a secret key K by public communication. The secrecy of K solely relies on the correctness of the laws of quantum mechanics, and not even infinite computing power allows Eve to obtain any information on K: as long as Eve is constrained by the laws of quantum mechanics, the key K is provably secret.[9]

In Sect. 4.2 below, we give some ideas on how quantum mechanics could be useful in order to allow Alice and Bob to agree on a secret key by public communication. These ideas will then be worked out to fully-fledged QKD schemes in the subsequent sections. But first, we formally specify the communication infrastructure, which, together with the problem description, is depicted in Fig. 1.

We assume that Alice and Bob can communicate via a classical communication channel, which allows them to send bit strings to each other. This channel is public in the sense that the attacker Eve may read all communication over it; however, we assume that she *cannot* insert or modify messages sent over the channel (as indicated by the one-way arrow from the channel to Eve in Fig. 1). If this is not *per se* guaranteed, then it can be achieved by means of information-theoretic message authentication [58].[10] Note that without (implicit or explicit) authentication, there is no way to prevent Eve from simply impersonating Bob, so that Alice unwittingly shares her key with Eve.

[9] However, from a practical point of view, one has to be aware that such a security proof is always with respect to a mathematical model that is assumed to capture reality, and as such security is only guaranteed if the model correctly captures reality. For instance in the security proof for QKD we assume that the devices Alice and Bob use to produce and measure the particles work according to their description. Obviously, a security proof is meaningless if, say, Alice's computer is infected by a virus that sends K to Eve by some hidden means. Thus, even a provably secure cryptographic scheme should not be trusted blindly, and one has to be aware of the possible failures.

[10] This comes at the price of requiring Alice and Bob to share a short secret one-time authentication key, so that, at first glance, we again seem to run into a circularity: in order to produce a secret key, Alice and Bob need a secret key to start with. However, for authentication, a relatively short secret key is sufficient, even for large messages. Thus, it suffices for Alice and Bob have a *short* secret key to start with in order to produce a *much larger* secret key. Of this larger key, a small part can then be used as authentication key for the next round etc.

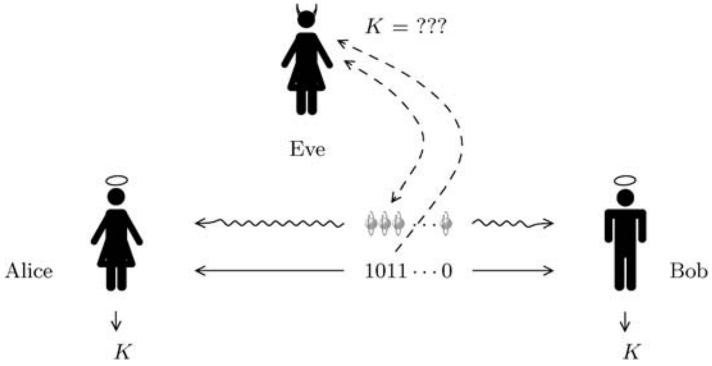

Fig. 1 Quantum key distribution by public communication

In addition to the classical communication channel, we assume that Alice and Bob are connected by a *quantum channel*, which allows Alice to send qubit systems to Bob. Also this channel is accessible to Eve; in fact, we allow Eve to have complete control over it. This means that when Alice sends qubit systems $A_1, \dots, A_n$ to Bob, then Eve can intercept $A_1, \dots, A_n$, apply an arbitrary unitary transformation $U \in \mathcal{U}(\mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2 \otimes \mathcal{H}_E)$ to $A_1, \dots, A_n$ and E, where E is a quantum system controlled by Eve (in some default initial state), and forward the transformed $A_1, \dots, A_n$ to Bob while keeping system E.

We would like to point out that Eve can in particular "block" the quantum communication (by forwarding default qubits to Bob). In this case, the quantum channel is useless for Alice and Bob, and one can show that it is impossible for Alice and Bob to produce a common secret key. Thus, the best we can hope for, is that Alice and Bob do agree on a common key if no Eve is present, and that if Eve is present and Alice and Bob manage to agree on a common key, then this key is secret to Eve.

4.2 Towards QKD

Consider a pure state $|\varphi\rangle = \sum_i \alpha_i |i\rangle$. Measuring $|\varphi\rangle$ in basis $\{|i\rangle\}_{i \in I}$ has the effect that i is observed with probability $p_i = |\alpha_i|^2$. Furthermore, this randomness is *fresh* and as such the outcome i of the measurement is *secret*: anyone who has not observed the outcome of the measurement has no information on which i was observed (beyond knowing its probability distribution). We stress that for this to hold, it is crucial that the initial state $|\varphi\rangle$ is *pure*.

As an example, measuring the qubit state $|+\rangle = H|0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ (or similarly $|-\rangle = H|1\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$) in the computational basis $\{|0\rangle, |1\rangle\}$ has the effect that 0 is observed with probability $\frac{1}{2}$ and 1 is observed with probability $\frac{1}{2}$. Thus, a secret random bit is obtained. Equivalently, measuring $|0\rangle$ (or $|1\rangle$) in the Hadamard basis $\{|+\rangle, |-\rangle\} = H\{|0\rangle, |1\rangle\}$ has the effect that "+" and "−" are observed each with probability $\frac{1}{2}$. Thus, identifying "+" with 0 and "−" with 1, which we do from now on, again a secret random bit is obtained. Repeating this procedure, i.e. measuring $|0\rangle \cdots |0\rangle$ qubit-wise in the Hadamard basis, can be used to produce a secret random bit *string*.

However, this does not really address our problem yet: it shows e.g. how Alice can produce a key K about which Eve has no information, but also Bob will have no information about K! In order to obtain a procedure that allows Alice and Bob to obtain a *common* secret key, consider now the 2-qubit state

$$|\Phi\rangle = \frac{1}{\sqrt{2}}\big(|00\rangle + |11\rangle\big) \in \mathbb{C}^2 \otimes \mathbb{C}^2,$$

called an *EPR pair* [28]. Recall that $|00\rangle$ is short for $|0\rangle \otimes |0\rangle$ etc. Measuring $|\Phi\rangle$ in the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ has the effect that 00 is observed with probability $\frac{1}{2}$ and 11 is observed with probability $\frac{1}{2}$. Thus, yet again, a secret random bit is produced. However, measuring in the product basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ can also be understood as measuring both qubits individually (and in whatever order) in the computational basis $\{|0\rangle, |1\rangle\}$. Thus, if the first qubit subsystem of $|\Phi\rangle$ is under Alice's control and the second under Bob's, and both measure their respective qubit in the computational basis, then they both observe *the same* random bit. And again, it is guaranteed that this bit is *secret*: any third party who has not observed the outcome of Alice or Bob's measurement has no information on the bit obtained. From the equality $|00\rangle + |11\rangle = |{+}{+}\rangle + |{-}{-}\rangle$, which is straightforward to verify, it follows immediately that the same also holds when Alice and Bob measure their respective qubit subsystems of $|\Phi\rangle$ both in the *Hadamard* basis; this will be important later on. Thus, when given n EPR pairs, Alice and Bob can obtain a random common secret key $K \in \{0,1\}^n$ by measuring within each EPR pair the two respective qubits in the same basis (computational or Hadamard).

What remains to be solved is: where do the EPR pairs come from, and in particular how can it be ensured that these indeed are EPR pairs? Note that if instead of an EPR pair Alice and Bob use for instance the first two qubits of the 3-qubit state $(|000\rangle + |111\rangle)/\sqrt{2} \in \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$ and measure them in the computational basis, where the third qubit is controlled by Eve, then Eve will also learn the random bit, simply by measuring her qubit. Thus, for the secrecy of K it is crucial that the ought-to-be EPR pair is really (close to) an EPR pair.

4.3 A QKD Scheme

Alice and Bob can try to obtain a list of shared EPR pairs as follows. Alice locally prepares n EPR pairs, i.e. 2-qubit quantum systems $A_i B_i$ that are in state $|\Phi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$, and sends the second qubit, B_i, of each pair to Bob. However, since Eve has full control over the quantum channel, there is no guarantee that the common state is not disturbed by Eve.

For instance, Eve could apply a so-called *controlled-NOT* [11] to B_i and E_i, where qubit E_i is in default state $|0\rangle$, such that the state of $A_i B_i E_i$ evolves from $|\Phi\rangle \otimes |0\rangle$ to $(|000\rangle + |111\rangle)/\sqrt{2}$. As discussed in the previous section, if Alice and Bob now decide to measure A_i and B_i in the *computational* basis, then the resulting common

[11] The controlled-NOT is given by the unitary transformation *CNOT* defined as $\mathit{CNOT}|b\rangle|c\rangle = |b\rangle|b \oplus c\rangle$ for $b, c \in \{0,1\}$, with $\oplus$ denoting addition modulo 2.

bit is completely insecure (Eve can learn it by measuring E_i in the computational basis). On the other hand, it is not hard to see that if, instead, Alice and Bob decide to measure A_i and B_i in the *Hadamard* basis, then they observe two random and *independent* bits, so that with probability $\frac{1}{2}$ the two bits are *distinct*. Recall from the previous section that when Alice and Bob hold a correct EPR pair, then they obtain the *same* random bit also when they measure in the Hadamard basis. Similarly, by applying a different suitable operation, Eve can enforce that Alice and Bob observe the same, yet insecure, bit when they use the *Hadamard* basis; but in this case, Alice and Bob observe independent bits when they use the *computational* basis instead.

This suggests the following procedure. For each qubit-pair $A_i B_i$, Alice and Bob decide at *random* and after Bob has received B_i, whether they should both use the computational or both the Hadamard basis to obtain the presumably common secure bit. Then, Alice and Bob compare the two n-bit strings X and Y that they respectively obtain at a randomly chosen subset of positions. If there are too many errors, meaning that X and Y differ at too many positions (within the chosen subset of positions), then Alice and Bob conclude that Eve has been heavily interacting with the quantum communication and they abort. Otherwise, if there are only a few errors, they proceed (see the preparation and error-estimation phase in Fig. 2).[12]

Intuitively, this seems to guarantee that Eve cannot have too much information on X, and similarly on Y (or else Alice and Bob abort). Indeed, at the time she can interact with B_i, she does not know yet the basis Alice and Bob will use, and therefore if she tries to entangle herself as in the above example to try to learn the key bit Alice and Bob obtain, she is likely to introduce an error. Thus, the number of errors between X and Y should indicate the amount of information Eve may have,

EPR-QKD:

Preparation: Alice creates n EPR pairs, and sends the second half of each pair to Bob, who confirms the receipt of the qubits. Then, Alice picks a random $\Theta \in \{0,1\}^n$ and sends it to Bob. For $j = 1, \dots, n$, Alice and Bob measure their respective parts of the j-th EPR pair in basis $H^{\Theta_j}\{|0\rangle, |1\rangle\}$ to obtain X_j on Alice's side and Y_j on Bob's side. (We expect $X_j = Y_j$ for all j.)

Error estimation: Alice chooses a random subset $\mathit{Test} \subset \{1, \dots, n\}$ of linear size and sends it to Bob. Then, Alice and Bob exchange and compare $X_{\mathit{Test}} = (X_i)_{i \in \mathit{Test}}$ and $Y_{\mathit{Test}} = (Y_i)_{i \in \mathit{Test}}$. If they differ at too many positions, Alice and Bob abort.

Error correction: Alice sends suitable error correcting information U to Bob that allows him to correct the remaining errors in Y and thus to recover X.

Key extraction: Alice and Bob apply a suitable function, chosen by Alice and announced to Bob, to X to obtain their common key K.

Fig. 2 An EPR-based QKD scheme

[12] The reason why Alice and Bob allow *some* errors is to tolerate a certain amount of noise in the quantum communication, which is inherent to current technology.

and it follows from sampling theory that the errors in a random subset reflects the number of errors on the whole. We stress, however, that up to this point, this is only intuition and no proof, as Eve may use a different strategy to attack the scheme than the very specific one considered here. Indeed, Eve may arbitrarily interact with the qubits communicated from Alice to Bob, and it is not clear that the above is her best strategy (and actually it is not). We will discuss later how to analyze this rigorously.

After the checking, and if Alice and Bob decide to proceed, they are still confronted with two problems. First, the two strings X and Y that Alice and Bob hold may still contain some limited number of errors, and, second, Eve may still have some limited amount of information. To correct the errors between X and Y without leaking too much information to Eve, a standard technique can be used: Alice chooses a random codeword $C \in \{0,1\}^n$ from a suitable error correcting code and sends $U := C \oplus X$ to Bob, and Bob decodes $C' := U \oplus Y$ to the closest codeword, $\hat{C}$, within the code and computes $\hat{X} := \hat{C} \oplus U$ as his guess for X. It is easy to see that if X and Y differ in only a small number of positions, then this also holds for C and C', and thus the error-correcting code guarantees that $\hat{C} = C$ from which $\hat{X} = X$ follows.

Taking care of the problem that Eve may have some limited information on X is done by means of *privacy amplification*. The purpose of privacy amplification is to transform a weakly-secret key X, by applying a suitably chosen function, into a fully-secure key K about which Eve has essentially no information. More details on how privacy amplification works is given in Sect. 5.3. The resulting EPR-based QKD scheme EPR-QKD is summarized in Fig. 2 above.

4.4 The BB84 QKD Scheme

The above QKD scheme requires Alice to produce EPR pairs, and it requires Bob to have *quantum memory* in order to store his parts of the EPR pairs until he learns Θ. Producing EPR pairs is feasible with current technology but more involved than producing single (unentangled) qubits. However, storing quantum states, e.g. in the form of polarized photons, turns out to be technically extremely difficult, such that even though scheme EPR-QKD can be implemented in theory, it is, to the best of our knowledge, not possible using current technology.

Here, we briefly show how to modify scheme EPR-QKD—without weakening its security—so that no quantum memory is needed and no EPR pairs have to be produced; Alice and Bob only need to prepare, send and measure-upon-arrival qubits. These tasks *can* be implemented using current technology. The resulting scheme coincides (up to some details) with Bennett and Brassard's original BB84 scheme.

The first modification we apply to scheme EPR-QKD is as follows. First of all, we denote the number of EPR pairs transmitted by N rather than by n. Furthermore, instead of using Alice's choice for Θ, Bob chooses "his own" $\Theta' \in \{0,1\}^N$ at random and measures the j-th qubit in basis $H^{\Theta'_j}\{|0\rangle, |1\rangle\}$ to obtain Y_j, and then Alice and Bob exchange their respective choices Θ and Θ', and only keep the positions j with $\Theta_j = \Theta'_j$. We argue that this modification does not weaken security. Indeed, if we blind out the transmissions of the qubits corresponding to the positions j with $\Theta_j \neq \Theta'_j$, then the modified scheme coincides with the original scheme EPR-QKD, with

$n = |\{j : \Theta_j = \Theta'_j\}| \approx N/2$. As such, if Eve could break the modified scheme, then she could also break the original scheme EPR-QKD.

As a next modification, we let Alice and Bob choose Θ and Θ' respectively, and let them do their measurements *as early as possible*. This means, Bob measures his qubits *upon arrival*, and Alice measures each of her qubits as soon as she has prepared the corresponding EPR pair. Changing the points in time where Alice and Bob do their local measurements does not change the outcome nor Eve's view of the scheme, and as such has no influence on its security. The resulting scheme coincides in spirit with the scheme by Bennett, Brassard and Mermin [11], which is a modification of Ekert's original EPR-based scheme [29], and does not require any quantum memory anymore.

To avoid the usage of EPR pairs, note now that measuring the first qubit of the j-th EPR pair in basis $H^{\Theta_j}\{|0\rangle, |1\rangle\}$ has the effect that Alice observes a random bit X_j and the qubit to be sent to Bob collapses to $H^{\Theta_j}|X_j\rangle$. Therefore, Alice could just as well choose $X_j \in \{0, 1\}$ at random and prepare and send qubit $H^{\Theta_j}|X_j\rangle$ to Bob. This then result in scheme BB84-QKD, summarized in Fig. 3, which is at least as secure as scheme EPR-QKD, but requires Alice and Bob to only prepare, send and measure-upon-arrival single qubits.

We would like to point out that from an intuitive point of view, BB84-QKD can also be appreciated directly, without the detour via EPR-QKD. Indeed, if Eve tries to obtain information on the transmitted qubits $H^{\Theta_j}|X_j\rangle$ by measuring (some of) them, then, because she does not know the "right" basis, she inevitably disturbs some of the qubits, which will be detected by Alice and Bob. The more information she tries to obtain the more qubits she disturbs, so that either Alice and Bob abort because they observed too many errors, or then Eve has gained only little information (which is taken care of by privacy amplification). This intuitive reasoning falls short of providing a rigorous security proof because it assumes Eve to treat the transmitted qubits individually, whereas quantum mechanics allows Eve to act on all of them collectively, as explained at the end of Sect. 4.1.

In Sect. 6, based on some tools developed in Sect. 5, we show how to rigorously analyze the EPR-pair-based scheme EPR-QKD. The provable security of the easier-to-implement scheme BB84-QKD then follows automatically.

BB84-QKD:

Preparation: Alice chooses random strings $X, \Theta \in \{0, 1\}^N$ and sends the qubits $H^{\Theta_1}|X_1\rangle \cdots H^{\Theta_N}|X_N\rangle$ to Bob. At the same time, Bob chooses a random $\Theta' \in \{0, 1\}^N$ and for $j = 1, \ldots, n$ measures the j-th qubit upon arrival in basis $H^{\Theta'_j}\{|0\rangle, |1\rangle\}$ to obtain Y_j, and he confirms the receipt of the qubits. Alice and Bob exchange Θ and Θ', and they update X and Y, respectively, by restricting them to the coordinates in $J = \{j : \Theta_j = \Theta'_j\}$.

Error estimation etc. as in EPR-QKD, with $n = |J|$.

Fig. 3 The BB84 QKD scheme

4.5 The Tolerable Noise

As briefly mentioned in footnote 12, current technology does not offer noise-free quantum communication. This means that in EPR-QKD, as well as in BB84-QKD for $j \in J$, even if Eve is not interacting with the communicated qubits and thus X_j is supposed to be equal to Y_j, it happens that $X_j \neq Y_j$ with some positive probability $\beta_\circ < \frac{1}{2}$. Then, in the error-estimation phase, Alice and Bob need to accept (slightly more than) a $\beta_\circ$-fraction of errors; otherwise, the scheme is aborted and thus no secret key is produced even when no Eve is attacking. As we will see in Sect. 6, allowing a $\beta_\circ$-fraction of errors implies that Eve has potentially $\mathrm{h}(\beta_\circ)n$ bits of information (in a well-defined sense) on the n-bit string X. Here and throughout the article, h is the *binary entropy function* $\mathrm{h}(p) = -(p \cdot \log(p) + (1-p) \cdot \log(1-p))$ for $0 < p < 1$, and $\mathrm{h}(p) = 0$ for $p = 0$ or 1, and log denotes the binary logarithm. Furthermore, it follows from coding theory that the error correction step leaks essentially another $\mathrm{h}(\beta_\circ)n$ bits of information on X, so that Eve may possibly have up to $2\,\mathrm{h}(\beta_\circ)n$ bits of information on the n-bit string X. Thus, if $2\,\mathrm{h}(\beta_\circ) \geq 1$ then Eve potentially knows all of X and as such it is not possible anymore to extract a strongly-secret key K from X. Therefore, for protocols EPR-QKD and BB84-QKD to work, it is needed that the error probability of the quantum communication satisfies $\mathrm{h}(\beta_\circ) < \frac{1}{2}$, which evaluates to $\beta_\circ \lesssim 11\%$.

4.6 Other Variants

Since the introduction of QKD with the BB84 scheme, a large variety of alternative QKD schemes has been proposed. Some of them offer a better *secret-key rate*, i.e., the number of key bits that can be generated per quantum-channel use, others tolerate a larger amount of noise than the 11% BB84-QKD can cope with, or are by some other means better suited for implementations. These schemes typically still follow the original construction design of BB84-QKD (or EPR-QKD in their respective EPR versions), but incorporate some modification. We give here a few examples (which can also be combined with each other) without trying to be exhaustive. The first one we mention is obtained by using a different set of designated states (rather than the four states induced by the computational and the Hadamard bases). For instance the *six-state* scheme [19] uses a set of three mutually-unbiased bases (resulting in six designated state vectors), or the B92 scheme [4] uses just two but non-orthogonal states. Another variant of BB84-QKD is obtained by having Alice and Bob choose each Θ_i and Θ'_i, respectively, in $\{0, 1\}$ not uniformly at random, but biased towards, say, 1. This increases the probability that $\Theta_i = \Theta'_i$ and thus the number of positions Alice and Bob can keep [43]. Finally, one can add an interactive so-called *advantage-distillation* step right after the preparation phase, which increases Bob's reliability in X without increasing Eve's. This then leads to a larger amount of noise that can be tolerated [32].

5 Some Quantum-Information-Theoretic Tools

5.1 Subset Sampling—Classical and Quantum

The *Hamming weight* $W(X)$ of a bit-string $X = (X_1, \dots, X_m) \in \{0,1\}^m$ is defined to be the number of 1's occurring within X. Similarly, the *relative* Hamming weight $\omega(X)$ of X is given by its Hamming weight divided by its bit-length m: $\omega(X) = W(X)/m$. We say that the relative Hamming weight of X is *ε-close* to β, denoted as $\omega(X) \approx_\varepsilon \beta$, if $|\omega(X) - \beta| \leq \varepsilon$. For any subset $T \subseteq \{1, \dots, m\}$ of size k, we write X_T for the restriction of X to the positions in T: $X_T = (X_i)_{i \in T} \in \{0,1\}^k$.

Consider the following problem: we want to *estimate* the (relative) Hamming weight of an unknown but fixed string $X \in \{0,1\}^m$ (of known bit-length m) by only looking at a small number of positions in X. A canonical way to do so is as follows: choose at random a sample subset $T \subset \{1, \dots, m\}$ of linear size (i.e. size αm for some constant $0 < \alpha < 1$), and take $\omega(X_T)$ as estimate for $\omega(X)$. Very generally, we allow the following kinds of estimation strategies: choose a sample subset $T \subset \{1, \dots, m\}$ according to *some* fixed probability distribution P_T, and compute the estimate for $\omega(X)$ as *some* (possibly randomized) function $\mathit{estim}(X_T)$ of X_T.

We want to measure the reliability of such a general estimation strategy, i.e., how well it predicts the (relative) Hamming weight of the string X. Actually, for technical reasons (and because the positions within the sample subset T are anyway revealed), we want to measure how well such a general strategy predicts the (relative) Hamming weight of $X_{\bar{T}} \in \{0,1\}^n$ (where $n = m - |T|$), i.e., of X restricted to the positions $\bar{T} = \{1, \dots, m\} \setminus T$ outside of the sample T (see Fig. 4, top). Therefore, for any $\varepsilon > 0$, we introduce the *error probability*

$$\mathit{err}_\varepsilon(m) := \max_{x \in \{0,1\}^m} P\big[\omega(x_{\bar{T}}) \not\approx_\varepsilon \mathit{estim}(x_T)\big]$$

where the probability is over the choice of T according to P_T.[13] By definition, for any choice of $X \in \{0,1\}^m$: $\omega(X_{\bar{T}}) \approx_\varepsilon \mathit{estim}(X_T)$ except with probability at most $\mathit{err}_\varepsilon(m)$. Using classical sampling theory (see e.g. [35]), one can e.g. show that for the above

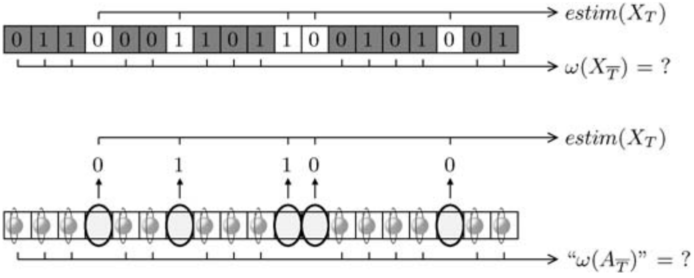

Fig. 4 Estimating ω for a string (*top*) and a quantum state (*bottom*)

[13] If the computation of *estim* is randomized, then the probability is also over this randomness.

canonical example with a random T of size αn with $\alpha \leq \frac{1}{2}$: $err_\varepsilon(m) \leq 2e^{-\varepsilon^2 \alpha m/2}$, and as such is exponentially small in m for fixed $\varepsilon, \alpha > 0$.

We now turn to the corresponding quantum problem. Consider a n-qubit system $A = A_1 \cdots A_n$, possibly entangled with another system E, and we want to estimate in a similar way as above how far away the state of A is from the all-zero state $|0 \cdots 0\rangle$, by *measuring* only a small number of the qubits. An estimation strategy as used above in the classical setting, defined by the distribution P_T and the function *estim*, can also be applied here: choose $T \subset \{1, \dots, m\}$ according to P_T, measure the qubits $A_T = (A_i)_{i \in T}$ qubit-wise in the computational basis to obtain X_T, and compute $estim(X_T)$ as estimate for the "relative Hamming weight" of the remaining system $A_{\bar{T}}$ (see Fig. 4, bottom). It remains to discuss what it should mean to have an estimation of the Hamming weight of a multi-qubit quantum system, and how to define and compute the reliability of a general strategy in this context. Formally, this is given by the following proposition; its proof is given in [15].[14]

Proposition 1 *Let AE be in pure state $|\varphi_{AE}\rangle$, and let $\rho_{TAE} = \rho_T \otimes \rho_{AE} = \sum_t P_T(t)|t\rangle\langle t| \otimes |\varphi_{AE}\rangle\langle\varphi_{AE}|$ be the hybrid state obtained by including the (independent) choice T. Then for any $\varepsilon > 0$ there exists $\tilde{\rho}_{TAE} = \sum_t P_T(t)|t\rangle\langle t| \otimes |\tilde{\varphi}^t_{AE}\rangle\langle\tilde{\varphi}^t_{AE}|$ with*

$$\delta(\rho_{TAE}, \tilde{\rho}_{TAE}) \leq \sqrt{err_\varepsilon(m)}$$

and such that for any $t \subset \{1, \dots, m\}$

$$\left|\tilde{\varphi}^t_{AE}\right\rangle \in \operatorname{span}\left\{|x\rangle \in \mathcal{H}_A : \omega(x_{\bar{t}}) \approx_\varepsilon estim(x_t)\right\} \otimes \mathcal{H}_E.$$

Informally, this guarantees that if the estimation strategy behaves well in the classical setting, in that $err_\varepsilon(m)$ is small, then it also behaves well in the quantum setting, in that we are guaranteed to be close to an *ideal* situation where the estimate $\beta = estim(X_T)$ predicts *with certainty* the approximate Hamming weight of the remaining qubits $A_{\bar{T}}$. The latter is to be understood in that after the measurement of A_T, the state of $A_{\bar{T}}E$ is of the form $|\varphi_{A_{\bar{T}}E}\rangle = \sum_y \alpha_y |y\rangle \otimes |\varphi^y_E\rangle$ where the sum is over all $y \in \{0,1\}^n$ with $\omega(y) \approx \beta$. What will be important for us is that β allows us to bound the number of y's occurring in the sum. Indeed, it is known that for any $\beta \leq \frac{1}{2}$, the number of $y \in \{0,1\}^n$ with $\omega(y) \approx_\varepsilon \beta$ is upper bounded by $2^{\mathrm{h}(\beta+\varepsilon)n}$, where h is the binary entropy function as introduced in Sect. 4.5.

5.2 (Conditional) Min-Entropy

The objective of information theory is to be able to quantify *information*, or the lack thereof, called *uncertainty* or *entropy*. A notion that proved to be useful in that it characterizes important operational quantities (for instance by how much data can be

[14] For simplicity, we implicitly assume here in Proposition 1 the function *estim* to be *deterministic*, but the corresponding also holds in case it is randomized (by also including the randomness used to compute *estim* into the hybrid state).

compressed, or how much information can be reliably sent over a noisy communication channel), is the Shannon entropy $\mathrm{H}(P_X) = -\sum_x P_X(x) \log P_X(x)$ (and the resulting Shannon information) [56], and its quantum information-theoretic analogue: the *von Neumann* entropy $\mathrm{S}(\rho) = -\operatorname{tr}(\rho \log \rho)$.

For cryptographic purposes, however, the Shannon entropy (respectively von Neumann entropy) is usually a too weak uncertainty measure, and in the context of classical cryptography, the stronger[15] notion of *min-entropy*

$$\mathrm{H}_\infty(P_X) = -\log\Big(\max_x P_X(x)\Big),$$

and its conditional version $\mathrm{H}_\infty(P_{XY}|Y) = -\log(\sum_y P_Y(y) \max_x P_{X|Y}(x|y))$,[16] turned out to be a much more useful measure for uncertainty. The (conditional) min-entropy captures how hard it is to *guess* the value described by the random variable X. Indeed, the min-entropy $\mathrm{H}_\infty(P_X)$ of a random variable X equals the negative-log of the success probability of predicting X from scratch when using an optimal strategy (which obviously predicts the most likely outcome). Similarly, the conditional min-entropy $\mathrm{H}_\infty(P_{XY}|Y)$ equals the negative-log of the success probability of predicting X from Y.

In quantum cryptography, we typically need to measure or lower-bound the uncertainty the attacker, holding a quantum system E, has on some classical piece of data X, held by the "good guy". We thus need a corresponding notion of min-entropy of a random variable *conditioned on a quantum system*: $\mathrm{H}_\infty(\rho_{XE}|E)$ for a hybrid state $\rho_{XE} = \sum_x P_X(x)|x\rangle\langle x| \otimes \rho_E^x \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_E)$ with classical X; such a notion was introduced by Renner [49].[17] The formal definition, given by

$$\mathrm{H}_\infty(\rho_{XE}|E) := \sup_{\sigma_E} \max\{h \in \mathbb{R} : 2^{-h} \cdot \mathbb{I}_X \otimes \sigma_E - \rho_{XE} \geq 0\}$$

where the supremum is over all density matrices $\sigma_E \in \mathcal{D}(\mathcal{H}_E)$, is not very relevant to us; we merely rely on some elementary properties, which we will mention below. When P_X (respectively ρ_{XE}) is clear from the context, we may speak of the min-entropy of X (conditioned on E), and we may write $\mathrm{H}_\infty(X)$ instead of $\mathrm{H}_\infty(P_X)$ (respectively $\mathrm{H}_\infty(X|E)$ instead of $\mathrm{H}_\infty(\rho_{XE}|E)$).

Similarly to the min-entropy $\mathrm{H}_\infty(P_{XY}|Y)$ conditioned on a random variable, $\mathrm{H}_\infty(\rho_{XE}|E)$ equals the negative-log of the success probability of predicting X when using an optimal strategy and having access to the quantum system E [40]. This in particular implies (but can also be directly deduced from the above definition of $\mathrm{H}_\infty(\rho_{XE}|E)$, see [49]) that $0 \leq \mathrm{H}_\infty(X|E) \leq \log|\mathcal{X}|$, with equality on the left if and only if there exists a measurement of E that allows X to be predicted with certainty, and with equality on the right if and only if X is random-and-independent of E (i.e. $\rho_{XE} = \mu_X \otimes \rho_E$).

[15] Stronger in the sense that $\mathrm{H}(P_X) \geq \mathrm{H}_\infty(P_X)$.

[16] In some literature one may also find a different definition for the conditional min-entropy: $\mathrm{H}_\infty(P_{XY}|Y) = -\sum_y P_Y(y) \log(\max_x P_{X|Y}(x|y))$; this version, though, turned out to be less useful.

[17] Actually, Renner's definition also allows X to be quantum.

In the following section, we argue that the (conditional) min-entropy also captures how many nearly-random bits can be extracted from X; this was actually the original motivation behind the definition proposed by Renner in [49].

But first, we state the following two useful properties (proven in [49]). The first one is the so called *chain rule*, which guarantees that getting additionally access to a t-qubit quantum system E reduces the min-entropy by at most t:

Lemma 1 *For any $\rho_{XEF} \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_E \otimes \mathcal{H}_F)$ with classical X:*

$$\mathrm{H}_\infty(\rho_{XEF}|EF) \geq \mathrm{H}_\infty(\rho_{XF}|F) - \log(\dim(\mathcal{H}_E)).$$

The second one relates the (conditional) min-entropy of a superposition with the (conditional) min-entropy of the corresponding mixture:

Lemma 2 *Let $|\varphi_{AE}\rangle \in \mathcal{H}_A \otimes \mathcal{H}_E$ and $\tilde{\rho}_{AE} \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_E)$ be quantum states of the form*

$$|\varphi_{AE}\rangle = \sum_{i \in J} \alpha_i |i\rangle |\varphi_E^i\rangle \quad \textit{and} \quad \tilde{\rho}_{AE} = \sum_{i \in J} |\alpha_i|^2 |i\rangle\langle i| \otimes |\varphi_E^i\rangle\langle\varphi_E^i|,$$

respectively, where $\{|i\rangle\}_{i\in I}$ is a basis of $\mathcal{H}_A$ and $J \subseteq I$. Furthermore, let ρ_{XE} (respectively $\tilde{\rho}_{XE}$) be the state obtained by measuring A of $|\varphi_{AE}\rangle$ (respectively of $\tilde{\rho}_{AE}$) in some basis $\{|x\rangle\}_{x\in\mathcal{X}}$, where the random variable X describes the outcome of the measurement. Then,

$$\mathrm{H}_\infty(\rho_{XE}|E) \geq \mathrm{H}_\infty(\tilde{\rho}_{XE}|E) - \log|J|.$$

Informally, this means that if X is obtained by measuring a state that can be written as a "small" *superposition* of orthogonal states, then the corresponding *mixture* behaves similarly, in that the entropy of X is close to the entropy obtained by measuring the mixture instead.

5.3 Privacy Amplification

Consider a common situation in cryptography where the "good guys" hold some (classical) information, say a uniformly distributed n-bit string, given by a random variable X, and the attacker controls some quantum system E, which may contain some information on X. However, we assume that the amount of information E contains on X is limited. Using the above introduced entropy measure, this is formalized by requiring that $\mathrm{H}_\infty(X|E)$ is lower bounded by some value $t > 0$. For instance, E may actually be classical and consist of $n - t$ arbitrary positions of the bit-string X, or, more generally, of an arbitrary $(n-t)$-bit-output function applied to X. Or, E may be quantum and arbitrarily depend on X but $\log(\dim(\mathcal{H}_E)) \leq n - t$. However, in its full generality of the problem, we do not put any restriction on *how* E's information on X is limited, only *that* it is and by *how much* it is.

The goal in such a situation is to transform the good guys' *weak key* X into a *secure key* K that is close to *random and independent* of E, i.e., from the adversary's point of view. Such a process is called *privacy amplification*. We stress once more that the resulting secure key K should be secure no matter how X and E are correlated, as long as $\mathrm{H}_\infty(X|E) \geq t$ for some (known) bound $t > 0$.

It is easy to see that privacy amplification, as described above, cannot be done by means of a *deterministic* transformation $K = f(X)$, even if f extracts just one bit. Indeed, for any candidate function f with a 1-bit output, the (limited) information E the adversary is allowed to have on X may actually consists of $f(X)$, so that $\mathrm{H}_\infty(X|E) \geq t := \mathrm{H}_\infty(X) - 1$, but nevertheless the extracted key $K = f(X)$ is far from random-and-independent of E. However, as we see below, privacy amplification is possible by means of a *randomized* transformation; this means K is computed from X as $K = f(S, X)$, where S, called the *seed*, is randomly chosen from some finite domain $\mathcal{S}$. For such a randomized transformation, we will require K to be close to random-and-independent of S *and* E: $\rho_{KSE} \approx \mu_K \otimes \rho_{SE}$; in other words, we allow the adversary to learn S. Otherwise, the problem would be trivial (simply by using the function $f(S, K) = S$) and thus not interesting,[18] and would not capture the typical situation where the adversary indeed does learn S.

As we argue below, privacy amplification can be done by means of universal hash functions. A function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{K}$ is called *universal* if

$$P\big[f(S,x) = f(S,x')\big] \leq \frac{1}{|\mathcal{K}|}$$

for all $x \neq x' \in \mathcal{X}$, where the seed S is uniformly distributed over $\mathcal{S}$.[19] Since for the following result to be meaningful it is needed that $|\mathcal{K}| < |\mathcal{X}|$, i.e., the function (for any fixed seed) is compressing, one also speaks of a universal *hash* function. An example of a universal hash function from $\mathcal{X} = \{0,1\}^n$ to $\mathcal{K} = \{0,1\}^\ell$ is given by $f(A,x) = Ax$, where the seed A is a uniformly random $(\ell \times n)$-matrix $A \in \{0,1\}^{\ell \times n}$, and where the computations are done modulo 2. More efficient examples (in terms of the size of the seed, e.g. with a seed of bit-size n instead of ℓn) exist.

Universal hash functions were originally introduced by Carter and Wegman [20] (though they use the terminology "universal$_2$"). The initial application was to storage-and-retrieval problems, but universal hash functions turned out to be useful for various other problems, like authentication or, as we discuss here, privacy amplification.

The possibility of doing privacy amplification in the classical setting, and of doing it by means of universal hashing, was pioneered by Bennett, Brassard and Robert in [12] and further worked out in [9, 13, 33, 36]. The generalization to the quantum setting (Theorem 1) is due to Renner and König [49, 52].

Theorem 1 *Let X be a random variable and E a quantum system with joint state $\rho_{XE} \in \mathcal{D}(\mathcal{H}_X \otimes \mathcal{H}_E)$. Let $f : \mathcal{S} \times \mathcal{X} \to \{0,1\}^\ell$ be a universal hash function, and let S be uniformly distributed over $\mathcal{S}$, independent of X and E. Then $K = f(S, X)$*

[18] The problem becomes non-trivial again by requiring K to be larger in bit-size than S; solutions to this problem are called *extractors*, which play a very important role in theoretical computer science.

[19] Equivalently, instead of a fixed function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{K}$ that takes a random seed $s \in \mathcal{S}$ as additional input, one can also think of a *family* $\{f_s = f(s,\cdot) : s \in \mathcal{S}\}$ of functions from which one member is then chosen at random.

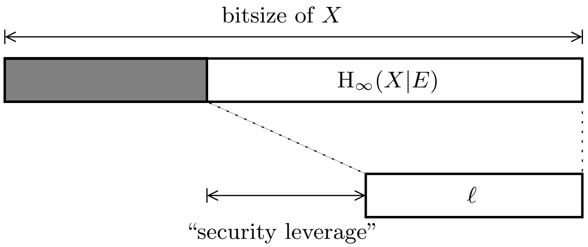

Fig. 5 Size of the key extracted by means of privacy amplification

satisfies

$$\delta\big(\rho_{KSE}, \mu_K \otimes \rho_{SE}\big) \leq \frac{1}{2} \cdot 2^{-\frac{1}{2}(H_\infty(X|E)-\ell)},$$

where μ_K is the completely mixed state $\mu_K = \frac{1}{|\mathcal{K}|}\mathbb{I}_K$.

This means that for any $0 < \ell < H_\infty(X|E)$, an ℓ-bit key $K = f(S, X)$ can be extracted, such that K is ε-close to random and independent of S and E where ε is exponentially small in $H_\infty(X|E) - \ell$ (see Fig. 5).

6 QKD Security Proof: Putting Things Together

With the above tools, we are now well equipped to do a rigorous security proof for the QKD scheme `EPR-QKD` in Fig. 2, which then also implies security for `BB84-QKD`. The proof we present here follows the approach of [15]; it uses elements from the quantum-information-theoretic approach proposed by Renner, but instead of reducing the security against general collective attacks to security against individual attacks (which is then much easier to analyze) as e.g. in [21, 49, 51], security against general attacks is proven directly with the help of Proposition 1.[20]

Formally, what we need to prove is the following. Given that the extracted key K has bit-length ℓ, the state $\rho_{K\mathit{Comm}E}$ is close (in terms of the trace distance δ) to $2^{-\ell}\mathbb{I} \otimes \rho_{\mathit{Comm}E}$, where $\mathit{Comm} = (\Theta, \mathit{Test}, \ldots)$ consists of the classical communication between Alice and Bob (which is potentially known to Eve), and E denotes Eve's quantum state at the end of the execution of `EPR-QKD`. We would like to point out that the value of ℓ depends on the course of the protocol and may be influenced by and known to Eve.[21] What the above guarantees is that if $\ell > 0$ then ℓ is (essentially) the *only* information Eve can have on the key K. Furthermore, by the properties of the trace distance, it follows that the real key K behaves (up to a small error probability)

[20] The downside of this approach is that it seems to be somewhat tailored to BB84-like QKD schemes. For instance, it does also apply to the scheme by Lo, Chau and Ardehali [43], which works similar to standard BB84 except that Alice and Bob choose their bases (computational or Hadamard) with some bias, but we do not know yet if this approach can be used to prove the six-state scheme [19] secure.

[21] For instance by fully *blocking* the quantum communication between Alice and Bob, Eve can always enforce $\ell = 0$; thus, there is no way to ensure that Alice and Bob obtain a key of positive length when under attack.

like an ideal key of the same bit length ℓ that is (fully) random-and-independent of Eve's view, no matter how K is used (e.g. as an encryption key); thus, *composability* as discussed in footnote 3 is guaranteed.

First we would like to analyze the error-estimation step in EPR-QKD. This step can be understood in that Alice and Bob try to estimate how far away their common state is from n copies of an EPR pair $|\Phi\rangle$. By a straightforward generalization of the observations from Sect. 5.1 on estimating the (relative) Hamming weight, where we replace the qubits by 2-qubit systems and the computational basis by the so-called *Bell basis*, this can be done by measuring a subset of the ought-to-be EPR pairs in the Bell basis. However, this is *not* what Alice and Bob do in scheme EPR-QKD, and this with good reasons: Alice and Bob are geographically separated and thus can only act locally on their qubits. Thus, the observations from Sect. 5.1 cannot be directly applied to analyze the error-estimation step in scheme EPR-QKD, but some twist is needed.

Let $|\psi_{ABE}\rangle$ be the state, shared between Alice, Bob and Eve, after the quantum communication and before Alice and Bob measure their parts. Note that we may indeed assume this state to be *pure*, since if not, it can be purified by increasing the dimension of E sufficiently. Furthermore, let $\rho_{\Theta XYE}$ be the hybrid state, consisting of Alice's choice Θ, of Alice and Bob's respective measurement outcomes X and Y, and of Eve's quantum system E. Recall that X_i and Y_i are obtained by measuring A_i and B_i, respectively, in basis H^{Θ_i}, as illustrated in Fig. 6 (left). By introducing the additional random variables $S = (S_1, \dots, S_n)$ and $W = (W_1, \dots, W_n)$, defined as

$$S_i = X_i \oplus Y_i \quad \text{and} \quad W_i = \begin{cases} X_i & \text{if } \Theta_i = 0, \\ Y_i & \text{if } \Theta_i = 1, \end{cases} \tag{1}$$

where $\oplus$ denotes addition modulo 2, we obtain the hybrid state $\rho_{\Theta XYSWE}$. Below, we show how to obtain the very same hybrid state by a different "experiment", which will be more convenient to analyze.

Consider the unitary transformation $U \in \mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ defined by

$$U|b\rangle|c\rangle = H|b\rangle|b \oplus c\rangle$$

for $b, c \in \{0, 1\}$; in other words, U is a controlled-NOT followed by a Hadamard transform on the first (i.e. the control) qubit. It is straightforward to verify that, simi-

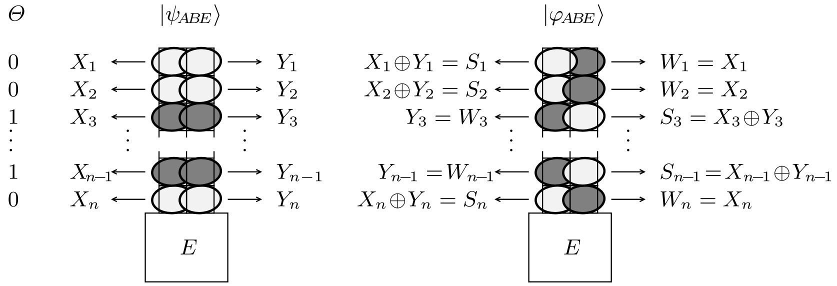

Θ
$|\psi_{ABE}\rangle$
$|\varphi_{ABE}\rangle$
0 X_1 Y_1
0 X_2 Y_2
1 X_3 Y_3
1 X_{n-1} Y_{n-1}
0 X_n Y_n
E
$X_1 \oplus Y_1 = S_1$ $W_1 = X_1$
$X_2 \oplus Y_2 = S_2$ $W_2 = X_2$
$Y_3 = W_3$ $S_3 = X_3 \oplus Y_3$
$Y_{n-1} = W_{n-1}$ $S_{n-1} = X_{n-1} \oplus Y_{n-1}$
$X_n \oplus Y_n = S_n$ $W_n = X_n$
E

Fig. 6 Obtaining X and Y from $|\psi_{ABE}\rangle$ (*left*), and from $|\varphi_{ABE}\rangle$ (*right*). White and grey ellipses correspond to measurements in the computational and the Hadamard basis, respectively

larly, $U(H|b\rangle H|c\rangle) = |b \oplus c\rangle H|c\rangle$. Let $|\varphi_{ABE}\rangle$ be obtained from $|\psi_{ABE}\rangle$ by applying U to each pair $A_i B_i$, and let $\sigma_{\Theta SWE}$ be the hybrid state obtained by choosing Θ at random in $\{0,1\}^n$, and obtaining S and W as follows. For every $i \in \{1,\dots,n\}$, if $\Theta_i = 0$ then S_i is obtained from measuring B_i in the computational basis and W_i from measuring A_i in the Hadamard basis, and if $\Theta_i = 1$ then S_i is obtained from measuring A_i in the computational basis and W_i from measuring B_i in the Hadamard basis, as illustrated in Fig. 6 (right). Defining X and Y such that (1) holds then determines the hybrid state $\sigma_{\Theta XYSWE}$. From the properties of U, and from the way $\sigma_{\Theta XYSWE}$ is obtained, it is not hard to see that $\sigma_{\Theta XYSWE} = \rho_{\Theta XYSWE}$.

This means that when analyzing the error-estimation step in scheme EPR-QKD in Fig. 2, we may assume X, Y etc. to be obtained as described in the above experiment, by measuring $|\varphi_{ABE}\rangle$ appropriately. Furthermore, we may also assume that first the qubits that produce the S_i's are measured (in the computational basis), and only afterwards, the qubits that produce the W_i's are measured (in the Hadamard basis). Recall that in the error-estimation step in scheme EPR-QKD in Fig. 2, Alice and Bob choose a small subset $\mathit{Test} \subset \{1,\dots,n\}$ and count the (relative) number of errors between X_{Test} and Y_{Test}; obviously, this equals the (relative) Hamming weight of S_{Test}. But now, this fits perfectly into the setting of Sect. 5.1: decide for each i, by choosing $\Theta_i \in \{0,1\}$ at random, whether to measure A_i or B_i to observe S_i—this determines the sample subset T using the notation/terminology from Sect. 5.1—and then compute $\mathit{estim}(S)$ as $\omega(S_{\mathit{Test}})$ for a randomly chosen subset $\mathit{Test} \subset \{1,\dots,n\}$. Note that here the function *estim* is randomized with *Test* as randomness.[22] It therefore follows from Sect. 5.1, that we are close to an ideal situation where *with certainty* the state of the remaining n qubits has relative Hamming weight approximately β (in the sense as discussed in Sect. 5.1), where $\beta = \mathit{estim}(S) = \omega(S_{\mathit{Test}})$ equals the error-rate Alice and Bob compute; furthermore, closeness is given by (the square-root of) the error probability of this estimation strategy in the classical case. From classical sampling theory, it follows that this error probability is exponentially small in n (for a subset *Test* of linear size $|\mathit{Test}| = \alpha n$).

Thus, we make only a small error when we conclude the following: for any possible choices of Θ and *Test* and for every possible outcome S when measuring the corresponding qubits (as illustrated in Fig. 6, right), the resulting state of the n qubits from which W is obtained—by measuring them in the *Hadamard* basis—and of E is of the form $\sum_z \alpha_z |z\rangle \otimes |\varphi_E^z\rangle$, where the sum is over all $z \in \{0,1\}^n$ with $\omega(z) \approx_\varepsilon \beta$.

Now we are in good shape. From Lemma 2 it follows that $\mathrm{H}_\infty(W|ES\Theta\mathit{Test})$ can be lower bounded with the help of the entropy of the measurement outcome when measuring the corresponding mixture.[23] Furthermore, measuring the corresponding mixture in the Hadamard basis produces a fully *random and independent* outcome. From the observation that the number of $z \in \{0,1\}^n$ with $\omega(z) \approx_\varepsilon \beta$ is upper bounded by $2^{\mathrm{h}(\beta+\varepsilon)n}$, it thus follows that

$$\mathrm{H}_\infty(W|ES\Theta\mathit{Test}) \geq n - \mathrm{h}(\beta + \varepsilon)n.$$

[22] It is important here that we consider *all* the positions that are measured to obtain an S_i as the measured subset T, and not only the ones that are used to actually compute the estimate.

[23] The additional conditioning on S etc. follows since the claim holds for any fixed choices for S etc.

Since X is uniquely determined when given Θ and S, the same bound also holds for the entropy $\mathrm{H}_\infty(X|ES\Theta \mathit{Test})$. Furthermore, by using the chain rule (Lemma 1), additionally conditioning on X_{Test} reduces the entropy by at most $|\mathit{Test}|$ (and conditioning on Y_{Test} comes "for free" as it is determined by S and X_{Test}):

$$\mathrm{H}_\infty(X|ES\Theta \mathit{Test} X_{\mathit{Test}} Y_{\mathit{Test}}) \geq n - \mathrm{h}(\beta+\varepsilon)n - |\mathit{Test}|.$$

Furthermore, one can show that the error-correcting information U decreases the entropy by at most $n-k$, where k is the log of the number of codewords in the code. This is easy to see in case of a *linear* code: instead of sending $U = C \oplus X$ as described at the end of Sect. 4, Alice could equivalently send the *syndrome* of X. The claim then follows by noting that the bit size of the syndrome is given by $n-k$, and applying the chain rule. Thus, it follows that

$$\mathrm{H}_\infty(X|ES\ldots Y_{\mathit{Test}}U) \geq n - \mathrm{h}(\beta+\varepsilon)n - |\mathit{Test}| - (n-k) \approx (1-2\mathrm{h}(\beta))n,$$

where the approximation follows by choosing $\varepsilon > 0$ as well as $\alpha > 0$ with $|\mathit{Test}| = \alpha \cdot n$ small enough, and by using an error-correcting code that approaches the bound $n - k \geq \mathrm{h}(\beta) \cdot n$ for correcting a β-fraction of errors. It now follows from privacy amplification that an ℓ-bit key K can be extracted from X for $\ell = \lambda n$ with $0 \leq \lambda < (1-2\mathrm{h}(\beta))$, such that $\rho_{K\mathit{Comm}E}$ and $2^{-\ell}\mathbb{I} \otimes \rho_{\mathit{Comm}E}$ are close, as required.[24]

Thus, we can indeed conclude that for a small enough error rate β observed by Alice and Bob in the error-estimation phase (specifically $\beta \lesssim 11\%$ such that $1 - 2\mathrm{h}(\beta) > 0$), a secure key of positive bit-length is obtained in the key-extraction phase of `EPR-QKD` (whereas otherwise Alice and Bob abort). Since `BB84-QKD` is at least as secure as `EPR-QKD`, as demonstrated in Sect. 4.4, we can also conclude security of `BB84-QKD`, with a secure key of positive bit-length under the same condition on β.

7 Secure Cooperation in the Bounded-Quantum-Storage Model

7.1 Beyond Secure Communication: Secure *Cooperation*

Whereas secure communication provides the means to control the information flow to a potential "outside" attacker Eve, when interacting with not necessarily trustworthy parties, we might also want to control the information flow to (possibly dishonest) "insiders". This is what secure *cooperation* tries to achieve. In the case of *two* mutually distrustful parties, Alice and Bob, this is referred to as secure *2-party cooperation* (2PC).

In a very general form, the 2PC problem is as follows. Given that Alice holds X and Bob holds Y, how can they jointly compute $f(X,Y)$ and $g(X,Y)$ for some fixed (possibly randomized and possibly identical) known functions f and g, simply by communicating with each other and doing local computations, in such a way that Alice learns $f(X,Y)$ and Bob learns $g(X,Y)$ *but neither party learns anything beyond.*

[24] It is straightforward, but somewhat tedious and not of importance for us, to compute the exact "error".

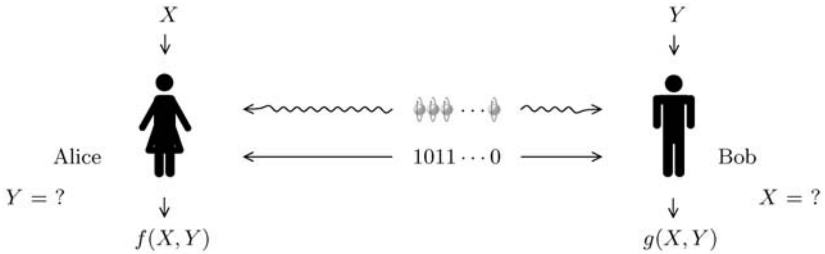

Fig. 7 Secure two-party cooperation

In particular Alice does not learn anything on Y beyond what she can deduce from X and $f(X,Y)$, and similarly for Bob; see Fig. 7. Furthermore, Alice should not be able to influence Bob's outcome $g(X,Y)$ beyond choosing her input X, and similarly for Bob. These properties should hold even if one of the parties, say Bob, is dishonest and actively attacks the scheme by arbitrarily deviating from the prescribed behavior in order to try to gain additional information on Alice's input X.

For instance, in the so-called millionaires' problem (as initially proposed by Yao [61]), Alice and Bob want to find out which of the two has more money, but neither of them is willing to reveal his/her actual wealth to the other party. In another example, Alice and Bob represent two companies that want to find out if it makes sense for them to merge, and as such they would like to find out how big the overlap between their respective customer sets is, but neither of them is willing to reveal his/her list of customers to the other company.

Is it possible, like for QKD, to design provably-secure quantum-cryptographic schemes for 2PC problems, whose security relies on the correctness of quantum mechanics alone? Unfortunately, here the answer is "no!" as follows from work by Mayers, Lo and Chau [41, 42, 46]. Nevertheless, as we argue in the sections below, it is still possible to design quantum-cryptographic 2PC schemes with interesting and strong security guarantees.

It is interesting to point out that even though quantum mechanics alone does not allow for (fully) secure 2PC schemes, it does allow for (certain) 2PC schemes, like Ambainis' oblivious transfer or coin flipping schemes [1, 2], whose security properties are stronger than what can be achieved in the classical (non-quantum) setting, but which are still (by far) too weak for typical applications.

7.2 The Bounded-Quantum-Storage Model

From the above no-go result we know that against any possible quantum-cryptographic scheme trying to solve a typical 2PC problem there exists a successful attack. However, it turns out that for cleverly designed schemes, the attacks implied by the no-go result do exist *in principle* but are hard to launch *in practice* because they require the attacker to reliably store a large number of quantum systems without affecting their respective states. Such an attack is indeed hard to launch since with current technology, we do not know how to store, say, photons without affecting their polarization. This motivates the study of whether it is possible to design quantum-cryptographic 2PC schemes that can be broken in principle (this is unavoidable), but *cannot* be broken in practice since breaking them would necessarily require large quantum-storage capacities. In other words, is it possible to design quantum-cryptographic schemes for 2PC problems, provably-secure under the assumption that an attacker can only store a certain number of photons (and that quantum mechanics is correct)? We show here that the answer is indeed yes!

We note that the *bounded-quantum-storage model*, as the above approach is called, departs somewhat from the original motivation for studying quantum cryptography, which was to obtain cryptographic schemes whose security relies on the correctness of the laws of quantum mechanics *alone*. Nevertheless, when compared with the complexity-theoretic approach discussed in the introduction, security in the bounded-quantum-storage model still offers several advantages. First of all, it still avoids to rely on an unproven complexity assumptions (whereas the assumed bound on the attacker's computing power is traded by a bound on his quantum storage capacity). Furthermore, in contrast to the complexity-theoretic approach, secure schemes in the bounded-quantum-storage model *cannot* be broken "in retrospect": if the attacker fails to store a large enough quantum state *during* the execution of the scheme, then the information needed to break the scheme is lost forever, and cannot be recovered at a later point in time even with unbounded quantum memory. This property is sometimes referred to as *everlasting security*. Finally, in order to get a very high level of security, one can combine the bounded-quantum-storage with the computational-complexity-theoretic approach and design cryptographic schemes that can be broken only if the attacker can efficiently solve some computational problem *and* has large quantum memory (as e.g. in [22]).

Formally, we specify the bounded-quantum-storage model as follows (see also Fig. 7). Like in the QKD setting, Alice and Bob can communicate via a quantum and via a classical communication channel.[25] In contrast to the QKD setting, we assume now that Alice and Bob have limited quantum storage capacities. In particular, we allow a dishonest party attacking the scheme to only store a certain number, q, of qubits—actually, we will only need this for dishonest *Bob*. We will express q as $q = \gamma n$, where n denotes the number of qubits transmitted in the scheme. The 2PC scheme we discuss here can be proven secure for any constant $\gamma < \frac{1}{4}$; this means that for the scheme to be secure it suffices that the number of transmitted qubits is larger than four times the number of qubits a dishonest Bob can store. Other schemes for other 2PC tasks may provide provable security for other values of γ.

When designing quantum cryptographic schemes in the bounded-quantum-storage model, we aim for schemes for which the faithful execution by the honest parties requires *no* quantum storage at all. In contrast to QKD where this feature is "only" aimed for because of practicality reasons (as explained in Sect. 4.4), here it is necessary for the model to make sense to have an as strong as possible separation between the quantum storage capacity of dishonest Bob and of the honest participants, since we must expect the dishonest party to be (much) better equipped. The scheme we discuss here and other schemes proposed in the literature (see Sect. 7.8) indeed require no quantum storage capacities from the honest parties, they merely need to prepare, communicate, and measure-upon-arrival qubits.

[25] Since here we deal with a malicious Alice or Bob but not with an outside attacker, there is no issue regarding security or authenticity of these channels.

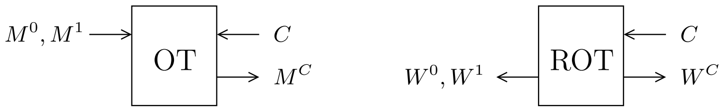

Fig. 8 Oblivious transfer (*left*) and randomized oblivious transfer (*right*)

7.3 (Randomized) Oblivious Transfer

The specific 2PC problem we study here is known as *oblivious transfer* (OT). In OT, Alice has as input two messages M^0 and M^1 in $\{0,1\}^\ell$, and Bob has a "choice bit" C; as a result of the cooperation, Bob is supposed to learn the message M^C of his choice but nothing more, whereas Alice should learn nothing at all; see Fig. 8 (left). OT is interesting since it is conceptually simple, yet it is very strong: one can show that OT is *complete* for secure 2PC, meaning that in principle one can implement *any* secure 2PC when given an OT [37, 38].

Actually, we consider and construct a *randomized* version (see Fig. 8, right), where the two messages M^0 and M^1 are not *input* by Alice, but they are produced as part of the OT and then *output* to Alice (and as such denoted as W^0 and W^1), where we still require that Bob only learns one of them, W^C, whereas the other is completely random to him. It is easy to see that given such a randomized OT, the ordinary OT can be obtained simply by running the randomized OT and having Alice send the one-time-pad encryptions $E^0 := E^0 \oplus W^0$ and $W^1 := M^1 \oplus W^1$ to Bob, who then can compute M^C as $E^C \oplus W^C$ (where $\oplus$ is bit-wise addition mod 2), whereas he learns no information on M^{1-C} due to the randomness of W^{1-C}. Thus, it suffices to construct a scheme for such a randomized OT.

7.4 Towards Quantum OT

In order to try to construct a randomized OT with the help quantum mechanics, let us again start with an EPR pair, which turned out to be a useful stepping stone towards QKD. As discussed in Sect. 4.2, if the two qubits of an EPR pair $|\Phi\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$ are both measured in the computational basis or both in the Hadamard basis, then a common random bit is observed. However, it is easy to see that if the two qubits are measured in "opposite" bases, i.e., one in the computational and the other in the Hadamard basis, or the other way round, then two random and *independent* bits are observed. Indeed, if the first qubit is measured, say, in the computational basis, then the second qubit collapses to $|0\rangle$ or $|1\rangle$, depending on the bit observed, and as such produces a random and independent outcome when being measured in the Hadamard basis.

Consider that, as in Sect. 4.2, Alice and Bob share a sequence of EPR pairs, i.e., of each pair, Alice controls one and Bob the other qubit. For instance Alice, or Bob, could produce the EPR pairs and of each pair keep one qubit and send the other to the other party (as in `EPR-QKD`). Even though we may give the impression here that Alice and Bob *store* their qubits, we actually mean that they will measure each of their qubits immediately upon generation/arrival, in order to avoid the need for quantum storage; the bases to be used are described below. Note that here, we do not have

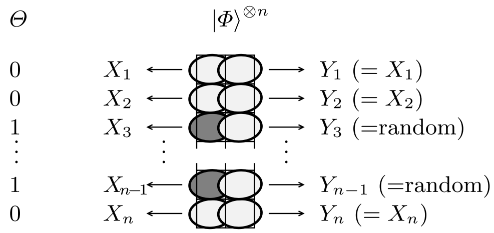

Fig. 9 Measuring EPR pairs in random (for Alice) and fixed (for Bob) bases

to worry that the qubits may get disturbed during the communication (in the setting we consider here there is no adversarial Eve), but rather that the sender of the EPR pairs is not honest and might send qubits in a different state. Now, assume that Alice measures some of her qubits in the computational basis and some in the Hadamard basis, and that the bits obtained from measuring in the computational basis, let us call them X^0, determine her outcome W^0 of the randomized OT, and the remaining bits, X^1, obtained by measuring in the Hadamard basis, determine W^1. Obviously, Bob should not know for which positions Alice used the computational and for which the Hadamard basis, since otherwise he could do the exact same measurements on his side and learn the exact same information as Alice, and thus would learn both W^0 *and* W^1. Note, however, that Bob can learn the information needed for one of the two: for instance if he is interested in learning W^0, then he can simply measure *all* qubits in the computational basis. This guarantees that he gets all the bits from X^0 right, whereas he gets randomness for the remaining bits (see Fig. 9). However, he does not know at which positions he finds the “good” bits. Thus, at this point, Alice needs to tell Bob for which positions she had used which basis, so that Bob can “filter out” the correct bits to obtain X^0 and compute W^0.

It is easy to see that if Alice and Bob act as described above, then Alice indeed learns W^0 and W_1 and nothing more, and Bob learns W^C but not W^{1-C}. But what if one of Alice and Bob is *dishonest* and deviates from the prescribed actions in order to try to obtain more information than allowed? It is obvious that Alice cannot learn any information on C, even if she is dishonest and actively attacks the scheme, since there is no information sent to her at all, except maybe for her parts of the EPR pairs, which do not depend on C. Arguing security against dishonest Bob is more delicate. For instance, dishonest Bob may not measure all the qubits in the same basis as required; this allows him to learn partial information on X^0 and on X^1. In order to take care of this, we let W^0 and W^1 be extracted from X^0 and on X^1, respectively, by means of privacy amplification, so that partial information on, say, X^0 is not sufficient to learn significant information on W^0. A more serious problem is that Bob may not measure all the qubits when he is supposed to. Indeed, if he could store all his qubits and delay his measurements until he learns the bases used by Alice, then he could measure every qubit in the same basis as Alice and would learn *both* X^0 and X^1, and thus W^0 and W^1, correctly. The crucial point now is that, by assumption, he *cannot* store all his qubits but is forced to do some (possibly collective) measurement on (some of) them. Intuitively, since at this point he does not know yet the bases that Alice will use for her measurements, one expects that this forces Bob to lose some information on X^0 or X^1. This intuition is indeed true, and thus guarantees security for this approach; however, formally proving this intuition in a rigorous and quantitative way is non-trivial.

7.5 OT in the Bounded-Quantum-Storage Model

The scheme for randomized OT follows the approach sketched above, except that we specify that the quantum communication is done from Bob to Alice. Furthermore, similar to the transition from `EPR-QKD` to `BB84-QKD`, we avoid the need for EPR pairs: instead of preparing an EPR pair and measuring one part in basis $H^C\{|0\rangle, |1\rangle\}$, with the result that Bob observes a random bit $Y_i \in \{0, 1\}$ and the other part collapses to $H^C|Y_i\rangle$, Bob chooses Y_i at random and prepares and sends to Alice qubit $H^C|Y_i\rangle$; and Alice measures all qubits immediately upon arrival. The resulting scheme, which is due to Damgård, Fehr, Renner, Salvail and Schaffner [23], is given in Fig. 10.[26]

It follows by trivial inspection that if Alice and Bob honestly follow the scheme, then Bob's output indeed coincides with his "string of his choice" W^C. It is also easy to see that a dishonest Alice, no matter in what way she may deviate from the scheme, learns no information on Bob's input bit C. Indeed, the only information Alice obtains are the qubits in step 1; however, this n-qubit state is described by the density matrix $\frac{1}{2^n}\mathbb{I} \in (\mathbb{C}^2)^{\otimes n}$, *independent of* C, and thus contains no information on C.

`QOT`(C):

1. Bob picks a random string $Y = (Y_1, \ldots, Y_n) \in \{0, 1\}^n$ and sends the qubits $H^C|Y_1\rangle \cdots H^C|Y_n\rangle$ to Alice.
2. Alice picks a random string $\Theta = (\Theta_1, \ldots, \Theta_n) \in \{0, 1\}^n$, and she obtains $X = (X_1, \ldots, X_n) \in \{0, 1\}^n$ by measuring the i-th qubit that she receives in basis $H^{\Theta_i}\{|0\rangle, |1\rangle\}$ (where $i = 1, \ldots, n$).
3. Alice sends Θ to Bob. Alice divides X into sub-strings X^0 and X^1 by collecting those X_i's with $\Theta_i = 0$ and with $\Theta_i = 1$, and Bob computes the sub-string Y^C of Y by only keeping those Y_i with $\Theta_i = C$. (Note: we expect $Y^C = X^C$.)
4. Alice chooses random, independent seeds S^0, S^1 for suitable universal hash functions f^0 and f^1, sends S^0 and S^1 to Bob, and she takes $W^0 := f^0(S^0, X^0)$ and $W^1 := f^1(S^1, X^1)$ as her output strings.
5. Bob takes $W^C := f^C(S^C, Y^C)$ as his output.

Fig. 10 The OT scheme in the bounded-quantum-storage model

[26] Actually, via the above EPR-pair based version, it is easy to see that the quantum communication could just as well be from Alice to Bob, with Alice sending $H^\Theta|X\rangle$ for random $X, \Theta \in \{0, 1\}^n$ and Bob measuring all qubits in basis $HC\{|0\rangle, |1\rangle\}$. It is interesting to note that in case of an *imperfect* quantum source, which produces two (or more) identical qubits in one time slot with some probability (like when using a weak coherent-pulse implementation), then the scheme from Fig. 10 actually becomes insecure, whereas the version with the quantum communication from Alice to Bob remains secure (with an appropriate adaptation of the parameters). However, we chose to present the scheme as given in Fig. 10 since its security proof (assuming an ideal quantum source) works more directly.

It remains to argue that a dishonest Bob with a bounded quantum memory can only learn one of the two strings, and has essentially no information on the other (and neither does he have any joint information on the two strings). This is nontrivial because we have no control over the n-qubit state that dishonest Bob sends in step 1; it may be an arbitrary n-qubit quantum state, possibly entangled with an auxiliary quantum state that Bob holds. Yet, still we need to argue quantitatively that Bob cannot have full information on both X^0 *and* X^1. We do this in Sect. 7.7; the main ingredient is a specific quantum uncertainty relation, where the uncertainty is expressed in terms of min-entropy.

7.6 Tool: Entropic Quantum Uncertainty Relation

Heisenberg's uncertainty principle [34] and its generalizations due to Robertson [53] and Schrödinger [55] state that for any two non-commuting observables, there exists no quantum state for which the measurement outcome for both observables is certain. In these original uncertainty relations, uncertainty is measured by the standard deviation of the outcome. Deutsch [27] proposed to express uncertainty relations using Shannon entropy H instead. Maassen and Uffink then proved a tight entropic uncertainty relation for any pair of observables [44]. Maassen and Uffink's uncertainty relation in particular implies the following bound on the (conditional) Shannon entropy of the outcome when measuring a qubit in the computational or the Hadamard basis.

Theorem 2 *Let* $\rho_\circ \in \mathcal{D}(\mathbb{C}^2)$ *be an arbitrary qubit, and let* $X_\circ \in \{0,1\}$ *be obtained by measuring* $\rho_\circ$ *in basis* $H^{\Theta_\circ}\{|0\rangle,|1\rangle\}$ *for a random* $\Theta_\circ \in \{0,1\}$. *Then* $\mathrm{H}(X_\circ|\Theta_\circ) \geq \frac{1}{2}$.

This bound is tight; e.g. the qubit $\rho_\circ = |0\rangle\langle 0|$ reaches it: $\mathrm{H}(X_\circ|\Theta_\circ{=}0) = 0$ and $\mathrm{H}(X_\circ|\Theta_\circ{=}1) = 1$ so that indeed $\mathrm{H}(X_\circ|\Theta_\circ) = \frac{1}{2}$.

By the chain rule for Shannon entropy, this scales up to $\mathrm{H}(X|\Theta) \geq \frac{1}{2}n$ when an arbitrary n-qubit state is measured qubit-wise in random basis $H^{\Theta_i}\{|0\rangle,|1\rangle\}$ to obtain $X = (X_1,\dots,X_n)$. However, as pointed out in Sect. 5.2, Shannon entropy is typically not a strong enough uncertainty measure for cryptographic purposes. In [23], Damgård et al. showed that, for large n, Theorem 2 (and similarly any uncertainty relation expressed in terms of Shannon entropy) scales up to essentially the same bound $\frac{1}{2}n$, also for the min-entropy: $\mathrm{H}_\infty(X|\Theta) \gtrsim \frac{1}{2}n$, up to some small loss and except with small error probability. The formal statement is as follows.

Theorem 3 *Let* ρ *be an arbitrary* n*-qubit state. Let* Θ *be uniformly distributed over* $\{0,1\}^n$ (*independent of the state*), *and let* $X \in \{0,1\}^n$ *be obtained by measuring* ρ *qubit-wise in bases* $H^{\Theta_i}\{|0\rangle,|1\rangle\}$, $i = 1,\dots,n$. *Then, for any* $\varepsilon > 0$ *there exists an event* $\mathcal{E}$ *such that*

$$\mathrm{H}_\infty(X|\Theta\mathcal{E}) \geq \left(\frac{1}{2} - \varepsilon\right)n$$

and the probability $P[\mathcal{E}]$ *of* $\mathcal{E}$ *is exponentially* (*in* n) *close to* 1.

The proof is highly non-trivial; it is given in (the full version of) [23].

7.7 Security Against Dishonest Bob

The security of the scheme QOT against a quantum-memory-bounded dishonest Bob can now be proven quite easily. The goal is to show that there exists $C \in \{0, 1\}$ so that X^{1-C} has high min-entropy from dishonest Bob's point of view, even when given C and W^C, so that privacy amplification (Theorem 1) implies that W^{1-C} is close to random-and-independent for Bob, even when given C and W^C. We point out that giving C and W^C to Bob "for free" in the security definition ensures that Bob can also not learn any *joint* information on W^0 and W^1, like their bit-wise XOR.

First, note that the entropic uncertainty relation of Theorem 3 is tailored to the way Alice obtains X in QOT. It follows that no matter what n-qubit state Bob sends in step 1, the string X Alice obtains satisfies

$$\mathrm{H}_\infty(X|\Theta) \geq \frac{1}{2}n,$$

up to an arbitrary small linear loss and except with exponentially small error probability in the sense of Theorem 3 (which we ignore for simplicity). Since, for given Θ, X is uniquely determined by X^0 and X^1, it follows that

$$\mathrm{H}_\infty(X^0X^1|\Theta) \geq \frac{1}{2}n.$$

From this one can show that there exists a random variable C such that

$$\mathrm{H}_\infty(X^{1-C}|\Theta C) \geq \frac{1}{4}n - 1.$$

The intuition here is that if the pair (X^0, X^1) carries t bits of entropy then one expects at least one of X^0 and X^1 to carry $t/2$ bits of entropy. This intuition is indeed true if the choice of which of the two has $t/2$ bits of entropy is allowed to be randomized and one sacrifices 1 bit of entropy. A slightly differently formalized version of this claim, called entropy-splitting lemma, can be found in [23].

Note that so far we have not taken into account that the n-qubit state Bob has sent in step 1 may be entangled with Bob's quantum system E, and as such Bob may have *quantum* information on X^{1-C}. However, the assumed bound on the dimension of E limits the amount of quantum information Bob may have. Indeed, from the chain rule, Lemma 1, it follows that

$$\mathrm{H}_\infty(X^{1-C}|\Theta C E) \geq \frac{1}{4}n - \log\bigl(\dim(\mathcal{H}_E)\bigr) - 1.$$

By the independent choice of the seeds, and by again using the chain rule, we further conclude that

$$\mathrm{H}_\infty(X^{1-C}|\Theta C S^C W^C E) \geq \frac{1}{4}n - \log\bigl(\dim(\mathcal{H}_E)\bigr) - \ell - 1,$$

where ℓ denotes the bit-size of W^0 and of W^1. It follows that if the qubit-size q of E is upper bounded by γn for some $0 \leq \gamma < \frac{1}{4}$, and if ℓ is chosen to be upper bounded by λn for some $0 < \lambda < \frac{1}{2}(\frac{1}{4} - \gamma)$, then

$$\mathrm{H}_\infty(X^{1-C}|\Theta C S^C W^C E) \geq \ell + \varepsilon n - 1$$

for some $\varepsilon > 0$. It follows by the privacy amplification theorem that $W^{1-C} = f^{1-C}(S^{1-C}, X^{1-C})$ is indeed exponentially close to random-and-independent from dishonest Bob's information (S^0, S^1, Θ and E) and C and W^C:

$$\delta\big(\rho_{W^{1-C} W_C C S^0 S^1 \Theta E}, \mu_{W^{1-C}} \otimes \rho_{W_C C S^0 S^1 \Theta E}\big) \leq 2^{-\varepsilon n/2}.$$

This proves security of QOT against a dishonest Bob whose quantum memory is bounded by γn with $\gamma < \frac{1}{4}$. In other words, if q denotes the qubit-size of Bob's quantum memory, then it suffices to communicate $n > 4q$ qubits in order to have provable security.

Note that, intuitively, one would expect QOT to be secure as long as $q < n$ with a linear gap: if Bob cannot store all the qubits then he inevitably misses some information. However, proving this rigorously is still an open problem. As a matter of fact, for any q between $\frac{1}{4}n$ and n, it is not know whether QOT is secure or not: no attack but also no security proof is known.

7.8 Variants of QOT, and Schemes for Other 2-Party Cooperation Problems

In the above description of QOT and its analysis it is assumed that the quantum channel is noise-free, so that whenever $\Theta_i = C$ then $X_i = Y_i$ with probability 1. As mentioned in Sect. 4.5, current technology does not allow for noise-free quantum communication. By doing error correction on the strings X^0 and X^1 by means of the same technique as in EPR-QKD and BB84-QKD, it is rather straightforward to make QOT robust against noisy quantum communication. The error-correction information then has to be taken into account in the security analysis, leading to a smaller bound on dishonest Bob's quantum memory.

Another imperfection that occurs with real-life implementations is multi-qubit emissions, meaning that when the qubit source is triggered to produce a qubit in a certain state, it may actually produces two (or more) qubits in the prescribed state. As mentioned in footnote 26, this renders protocol QOT insecure, whereas the version of QOT in which the quantum communication goes from Alice to Bob remains secure (with an appropriate adaption of the parameters).

In additional work by Damgård et al. (see [23–26] for the complete line of research), secure quantum-cryptographic schemes in the bounded-quantum-storage model for a sequence of additional 2PC tasks have been proposed and analyzed: for Rabin's OT, one-out-of-many OT, bit and string commitment, identification, and QKD with implicit identification.

In another line of work [54, 59], Schaffner, Terhal and Wehner analyzed and proved-secure QOT in a slightly different model. Instead of assuming a strict bound on the size of dishonest Bob's quantum memory, they allow Bob to store all the qubits but assume the storage to be *noisy*.

8 Conclusions

Quantum cryptography is an active research area with the goal to construct cryptographic schemes with provably strong security guarantees. It combines "cryptographic thinking" with elements of quantum mechanics. Designing and analyzing quantum cryptographic schemes often leads to interesting new questions in quantum information theory, whose answers may find other applications as well.

Current research activities include the design of new QKD schemes that require less trust in the devices used to prepare and measure the involved quantum states, or whose security relies on fewer properties from quantum mechanics, for instance only on the non-signaling property. Another current research activity is the design of quantum 2PC schemes relying on other reasonable technological assumptions on the adversary than bounding his quantum memory, like the noisy-quantum-storage model mentioned. On the implementational side, great effort is put into the problem that over long distances, quantum communication currently is too noisy for QKD to work.

With further advances in theory and practice, quantum cryptography could very well be the first real application of quantum mechanics at the single-quantum level.

Acknowledgements The author would like to thank the anonymous referees. They had put a real great effort into reviewing this article, and they had provided a lot of useful feedback that helped to improve the presentation of this article.

References

1. Ambainis, A.: A new protocol and lower bounds for quantum coin flipping. J. Comput. Syst. Sci. **68**(2), 398–416 (2004)
2. Ambainis, A., Nayak, A., Ta-Shma, A., Vazirani, U.: Dense quantum coding and quantum finite automata. J. ACM **49**(4), 496–511 (2002)
3. Bell, J.S.: On the Einstein-Podolsky-Rosen paradox. Physics **1**(3), 195–290 (1964)
4. Bennett, C.H.: Quantum cryptography using any two nonorthogonal states. Phys. Rev. Lett. **68**(21), 3121–3124 (1992)
5. Bennett, C.H., Brassard, G.: Quantum cryptography and its application to provably secure key expansion, public-key distribution, and coin-tossing. In: IEEE International Symposium on Information Theory (ISIT), p. 91 (1983)
6. Bennett, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. In: IEEE International Conference on Computers, Systems, and Signal Processing, pp. 175–179 (1984)
7. Bennett, C.H., Brassard, G., Breidbart, S.: Quantum cryptography II: How to re-use a one-time pad safely even if P = NP. Unpublished Manuscript (1982)
8. Bennett, C.H., Brassard, G., Breidbart, S., Wiesner, S.: Quantum cryptography, or unforgeable subway tokens. In: CRYPTO 1982, pp. 267–275. Plenum, New York (1982)
9. Bennett, C.H., Brassard, G., Crépeau, C., Maurer, U.M.: Generalized privacy amplification. IEEE Trans. Inf. Theory **41**, 1915–1923 (1995)
10. Bennett, C.H., Brassard, G., Crépeau, C., Skubiszewska, M.-H.: Practical quantum oblivious transfer. In: CRYPTO 1991. Lecture Notes in Computer Science, vol. 576, pp. 351–366. Springer, Berlin (1991)
11. Bennett, C.H., Brassard, G., Mermin, N.D.: Quantum cryptography without Bell's theorem. Phys. Rev. Lett. **68**(5), 557–559 (1992)
12. Bennett, C.H., Brassard, G., Robert, J.-M.: How to reduce your enemy's information. In: CRYPTO 1985. Lecture Notes in Computer Science, vol. 218, pp. 468–476. Springer, Berlin (1985)
13. Bennett, C.H., Brassard, G., Robert, J.-M.: Privacy amplification by public discussion. SIAM J. Comput. **17**(2), 210–229 (1988)

14. Biham, E., Boyer, M., Boykin, P.O., Mor, T., Roychowdhury, V.: A proof of the security of quantum key distribution. In: 32rd Annual ACM Symposium on Theory of Computing (STOC), pp. 715–724 (2000)
15. Bouman, N., Fehr, S.: Sampling in a quantum population, and applications. http://arxiv.org/abs/0907.4246 (2009)
16. Brassard, C., Crépeau, C.: Quantum bit commitment and coin tossing protocols. In: CRYPTO 1990. Lecture Notes in Computer Science, vol. 537, pp. 49–61. Springer, Berlin (1990)
17. Brassard, G.: Brief history of quantum cryptography: A personal perspective. In: IEEE Information Theory Workshop on Theory and Practice in Information-Theoretic Security, pp. 19–23. IEEE, New York (2005)
18. Brassard, G., Crépeau, C., Jozsa, R., Langlois, D.: A quantum bit commitment scheme provably unbreakable by both parties. In: 34th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 362–371 (1993)
19. Bruß, D.: Optimal eavesdropping in quantum cryptography with six states. Phys. Rev. Lett. **81**(14), 3018–3021 (1998)
20. Carter, J.L., Wegman, M.N.: Universal classes of hash functions. In: 9th Annual ACM Symposium on Theory of Computing (STOC), pp. 106–112 (1977)
21. Christandl, M., König, R., Renner, R.: Post-selection technique for quantum channels with applications to quantum cryptography. Phys. Rev. Lett. **101**(2), 020504 (2009)
22. Damgård, I.B., Fehr, S., Lunemann, C., Salvail, L., Schaffner, C.: Improving the Security of Quantum Protocols via Commit-and-Open. In: CRYPTO 2009. Lecture Notes in Computer Science, vol. 5677, pp. 408–427. Springer, Berlin (2009)
23. Damgård, I.B., Fehr, S., Renner, R., Salvail, L., Schaffner, C.: A tight high-order entropic quantum uncertainty relation with applications. In: CRYPTO 2007. Lecture Notes in Computer Science, vol. 4622, pp. 360–378. Springer, Berlin (2007)
24. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Cryptography in the bounded quantum-storage model. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 449–458 (2005). Full version available at: http://arxiv.org/abs/quant-ph/0508222v2
25. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Secure identification and QKD in the bounded-quantum-storage model. In: CRYPTO 2007. Lecture Notes in Computer Science, vol. 4622, pp. 342–359. Springer, Berlin (2007)
26. Damgård, I.B., Fehr, S., Salvail, L., Schaffner, C.: Cryptography in the bounded-quantum-storage model. SIAM J. Comput. **37**(6), 1865–1890 (2008)
27. Deutsch, D.: Uncertainty in quantum measurements. Phys. Rev. Lett. **50**(9), 631–633 (1983)
28. Einstein, A., Podolsky, B., Rosen, N.: Can quantum-mechanical description of physical reality be considered complete? Phys. Rev. **47**(10), 777–780 (1935)
29. Ekert, A.K.: Quantum cryptography based on Bell's theorem. Phys. Rev. Lett. **67**(6), 661–663 (1991)
30. Even, S., Goldreich, O., Lempel, A.: A randomized protocol for signing contracts. In: CRYPTO 1982. Plenum, New York (1982)
31. Gisin, N., Ribordy, G., Tittel, W., Zbinden, H.: Quantum cryptography. Rev. Mod. Phys. **74**, 145–195 (2002)
32. Gottesman, D., Lo, H.-K.: Proof of security of quantum key distribution with two-way classical communications. IEEE Trans. Inf. Theory **49**(2), 457–475 (2003). quant-ph/0105121
33. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM J. Comput. **28**(4), 1364–1396 (1999)
34. Heisenberg, W.: Schwankungserscheinungen und quantenmechanik. Z. Phys. **40**, 501–506 (1927)
35. Hoeffding, W.: Probability inequalities for sums of bounded random variables. J. Am. Stat. Assoc. **58**(301), 13–30 (1963)
36. Impagliazzo, R., Levin, L.A., Luby, M.: Pseudo-random generation from one-way functions. In: 21st Annual ACM Symposium on Theory of Computing (STOC), pp. 12–24 (1989)
37. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer—efficiently. In: CRYPTO 2008. Lecture Notes in Computer Science, vol. 5157, pp. 572–591. Springer, Berlin (2008)
38. Kilian, J.: Founding cryptography on oblivious transfer. In: 20th Annual ACM Symposium on Theory of Computing (STOC), pp. 20–31 (1988)
39. König, R., Renner, R., Bariska, A., Maurer, U.: Small accessible quantum information does not imply security. Phys. Rev. Lett. **98**, 140502 (2007)
40. König, R., Renner, R., Schaffner, C.: The operational meaning of min- and max-entropy. IEEE Trans. Inf. Theory **55**(9), 4337–4347 (2009)

41. Lo, H.-K.: Insecurity of quantum secure computations. Phys. Rev. A **56**(2), 1154–1162 (1997)
42. Lo, H.-K., Chau, H.F.: Is quantum bit commitment really possible? Phys. Rev. Lett. **78**(17), 3410–3413 (1997)
43. Lo, H.-K., Chau, H.F., Ardehali, M.: Efficient quantum key distribution scheme and a proof of its unconditional security. J. Cryptology **18**(2), 133–165 (2005)
44. Maassen, H., Uffink, J.B.M.: Generalized entropic uncertainty relations. Phys. Rev. Lett. **60**(12), 1103–1106 (1988)
45. Mayers, D.: Quantum key distribution and string oblivious transfer in noisy channels. In: CRYPTO 1996. Lecture Notes in Computer Science, vol. 1109, pp. 343–357. Springer, Berlin (1996)
46. Mayers, D.: Unconditionally secure quantum bit commitment is impossible. Phys. Rev. Lett. **78**(17), 3414–3417 (1997)
47. Mayers, D.: Unconditional security in quantum cryptography. J. ACM **48**(3), 351–406 (2001)
48. Rabin, M.: How to exchange secrets by oblivious transfer. Technical Report, Harvard Aiken Computation Lab (1981)
49. Renner, R.: Security of Quantum Key Distribution. Ph.D. Thesis, ETH Zürich (Switzerland), September 2005. http://arxiv.org/abs/quant-ph/0512258
50. Renner, R.: Symmetry of large physical systems implies independence of subsystems. Nat. Phys. **3**, 645–649 (2007)
51. Renner, R., Gisin, N., Kraus, B.: An information-theoretic security proof for QKD protocols. Phys. Rev. Lett. A **72**, 012332 (2005)
52. Renner, R., König, R.: Universally composable privacy amplification against quantum adversaries. In: TCC 2005. Lecture Notes in Computer Science, vol. 3378, pp. 407–425. Springer, Berlin (2005)
53. Robertson, H.P.: The uncertainty principle. Phys. Rev. **34**(1), 163–164 (1929)
54. Schaffner, C., Terhal, B.M., Wehner, S.: Robust cryptography in the noisy-quantum-storage model. Quantum Inf. Comput. **9**(11&12), 963–996 (2009)
55. Schrödinger, E.: Zum Heisenbergschen Unschärfeprinzip. In: Sitzungsberichte der Preussischen Akademie der Wissenschaften, physikalisch-mathematische Klasse, pp. 296–303 (1930)
56. Shannon, C.E.: A mathematical theory of communication. Bell Syst. Tech. J. **27**, 379–423 (1948) Also 623–656
57. Shor, P.W., Preskill, J.: Simple proof of security of the BB84 quantum key distribution protocol. Phys. Rev. Lett. **85**(2), 441–444 (2000)
58. Simmons, G.J.: Authentication theory/coding theory. In: CRYPTO 1984. Lecture Notes in Computer Science, vol. 196, pp. 411–431. Springer, Berlin (1984)
59. Wehner, S., Schaffner, C., Terhal, B.M.: Cryptography from noisy storage. Phys. Rev. Lett. **100**(22), 220502 (2008)
60. Wiesner, S.: Conjugate coding. SIGACT News **15**(1), 78–88 (1983). Original manuscript written circa 1970
61. Yao, A.: Protocols for secure computations. In: 23rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 160–164 (1982)