Unconditional Security in Cryptography

Stefan Wolf

Computer Science Department
Swiss Federal Institute of Technology (ETH Zürich)
CH-8092 Zürich, Switzerland
E-mail address: wolf@inf.ethz.ch

Abstract. The fact that most presently-used cryptosystems cannot be rigorously proven secure and hence permanently face the risk of being broken motivates the search for schemes with unconditional security. The corresponding proofs however must be based on information theory rather than complexity theory. One reason for this is the lack of known lower bounds on the running time of algorithms solving certain computational problems such as the discrete-logarithm problem or the integer-factoring problem. At the beginning of an information-theoretic analysis of cryptosystems stands Shannon's definition of perfect secrecy, unquestionably the strongest possible security definition, and his well-known inequality giving a lower bound on the key length of every perfectly secret cipher, thus suggesting that such a high level of confidentiality cannot be realized in any practical scheme. This pessimism has later been qualified by several authors who showed that unconditional security can be achieved in many special but realistic scenarios. Some of these approaches are described in this introductory overview article.

1 Computational Versus Information-Theoretic Security

The security of many presently-used cryptosystems, e.g., of all public-key cryptographic schemes, is based on the assumed hardness of computational problems in number theory such as the integer-factoring problem (e.g., RSA [28]) or the problem of computing discrete logarithms in certain finite cyclic groups (e.g., Diffie-Hellman [13]). Such a cryptosystem is called *computationally secure*.

Up to date, no practical cipher has been proven computationally secure. Note first of all that it is an inherent fact that computational security can only hold under certain assumptions on the adversary's computer resources. In other words, a computationally infinitely powerful opponent can break every system of this type by exhaustive search over the key space.

One reason for the lack of proofs of cryptographic security is that in complexity theory, actually proved lower bounds on the running time of algorithms solving specific problems are either rather weak (and useless in cryptography) or valid only in special computational models (e.g., [32]). Unfortunately, such bounds are not directly useful neither since it can never be guaranteed that the adversary is restricted to this particular model. So-called "provable computational security" is always conditional and means that an efficient reduction

from a well-known problem that is believed to be hard, such as the discrete-logarithm problem or the decisional Diffie-Hellman problem, to breaking the proposed system can be given, thus showing that the cryptosystem is secure if some widely-accepted standard complexity assumption is true (e.g., [11]).

Finally, it has been shown that the integer-factoring as well as the discrete-logarithm problem can be solved in polynomial-time by a *quantum computer*, i.e., a computing device that is able to exploit certain effects from quantum mechanics [31]. The security of most public-key cryptographic protocols is based on the hardness of at least one of these problems.

Consequently, practical computational security is always conditional and additionally faces the risk of being broken by progress in the theory of efficient algorithms or in hardware engineering. On the other hand it appears desirable from both a scientific and practical point of view to design cryptosystems whose security is not based on any assumptions and can be proven rigorously. Because of the reasons discussed above, such security proofs must be based on *information theory* (i.e., probability theory) rather than complexity theory. There have been made various attempts at realizing this type of security, some of which we describe in this overview paper.

The outline of the article is as follows. We start with an introduction to some basic definitions and facts from probability and information theory (Section 2). Then, a definition of perfect secrecy, undoubtedly the strongest possible security definition in cryptography, is given (Section 3). Shannon's pessimistic theorem suggests that perfect secrecy is necessarily impractical. However, we describe a number of approaches that could qualify this pessimism. All these constructions have in common that some kind of limitations are needed on the amount of information that an opponent obtains. Realistic scenarios have been described where such an upper bound on the adversary's knowledge can for instance be based on noise, an inherent property of every physical communication channel (Section 4). Motivated by these examples, a model has been presented and analyzed that shows how two parties can generate a secret key from common randomness by communication over an insecure but authentic (or even completely insecure) channel (Sections 5 and 6).

2 Basic Concepts of Information Theory

Information theory goes back to Claude Shannon and his celebrated 1948 paper [30]. Examples of good and detailed introductions into the field are [10] or [5].

2.1 Probability-Theoretic Preliminaries

In this section we introduce some basic probability-theoretic concepts. For a detailed introduction see for example [14].

Let $\mathcal{X}$ be a countable set. The *distribution* P_X of a *discrete random variable* X with *range* $\mathcal{X}$ is a mapping

$$P_X : \mathcal{X} \longrightarrow \mathbf{R}_{\geq 0}$$

with $\sum_{x \in \mathcal{X}} P_X(x) = 1$. If $\mathcal{X} \subset \mathbf{R}$, the expectation of X is defined as

$$\mathrm{E}[X] := \sum_{x \in \mathcal{X}} x \cdot P_X(x) \ .$$

Let f be a convex function. Then we have

$$\mathrm{E}[f(X)] \geq f(\mathrm{E}[X]) \ . \tag{1}$$

Inequality (1) is called *Jensen's inequality*. Most of the basic inequalities in information theory follow directly from this inequality.

The *joint distribution* $P_{X_1 X_2 \cdots X_N}$ of N random variables is a probability distribution over the set $\mathcal{X}_1 \times \mathcal{X}_2 \times \cdots \times \mathcal{X}_N$. The random variables $X_1, X_2, \ldots, X_N$ are called *statistically independent* if

$$P_{X_1 X_2 \cdots X_N}(x_1, x_2, \ldots, x_N) = P_{X_1}(x_1) \cdot P_{X_2}(x_2) \cdots P_{X_N}(x_N)$$

for all $x_1, x_2, \ldots, x_N$, i.e., when the joint distribution equals the product of the *marginal distributions*.

An event $\mathcal{A}$ is a subset of the range of a random experiment. By $\mathrm{Prob}[\mathcal{A}]$ we denote the probability of $\mathcal{A}$, i.e., the sum of the probabilities of all the outcomes belonging to $\mathcal{A}$. The *conditional distribution* of X, given that the event $\mathcal{A}$ (with $\mathrm{Prob}[\mathcal{A}] > 0$) occurs, is defined as

$$P_{X|\mathcal{A}}(x) := \frac{\mathrm{Prob}[\{X = x\} \cap \mathcal{A}]}{\mathrm{Prob}[\mathcal{A}]} \ .$$

As a special case, a random variable can be conditioned on the event

$$\mathcal{A} := \{Y = y\}$$

that another random variable Y takes a particular value y. The resulting distribution

$$P_{X|Y}(x, y) := P_{X|Y=y}(x)$$

is called the *conditional distribution of X given Y*. Note that the function $P_{X|Y}(\cdot, \cdot)$ with two arguments is *not* a probability distribution on $\mathcal{X} \times \mathcal{Y}$, but for every $y \in \mathcal{Y}$, the function $P_{X|Y}(\cdot, y)$ is a distribution on $\mathcal{X}$.

2.2 Bar Kochba, Uncertainty, and Entropy

The following story has been reported about *Bar Kochba* (the "Son of the Star"), leader of the Jews during their independence war in 135 B.C., who defended his fortress heroically against a superior number of Romans [27].

"It is also said that Bar Kochba sent out a scout to the Roman camp who was captured and tortured, having his tongue cut out. He escaped from captivity and reported back to Bar Kochba, but being unable to talk, he could not tell in words what he had seen. Bar Kochba accordingly asked him questions which he could answer by nodding or shaking his head. Thus he acquired from his mute scout

the information he needed to defend the fortress. [...] It occurred to me that, if the story of Bar Kochba were true, then he would have been the forefather of information theory".

In the so-called *Bar-Kochba game*, one player has to find out, by asking yes/no-questions, what the second player has in mind. This game was extremely popular among writers in Budapest at the beginning of this century. Regardless of the (possibly adaptive) strategy of the questioner he cannot, with at most 20 questions, distinguish between more than 2^{20}, i.e., about one million, different objects (because there are only 2^{20} ways of answering the 20 questions differently). On the other hand, given that the object to be found comes from a set of size at most n, then $\lceil \log_2 n \rceil$ questions are always sufficient if the following strategy is used. Let a fixed encoding of all the objects by binary strings of length 20 be defined. Then, the strategy is to ask whether the first, second, ... bit of the encoding is 1.

This example shows the close relationship between the Bar-Kochba game and binary coding. For a random variable X that takes one of $n = 2^k$ values with equal probabilities, the minimal average number of questions in the Bar-Kochba game, as well as the minimal average codeword length of a prefix-free binary code, is k. Note that this bound cannot be beaten even if a strategy is used with variable codeword lengths for the different outcomes. We call this quantity the uncertainty or *entropy* of X, denoted by $H(X)$.

If the size of the range $\mathcal{X}$ of X is not a power of 2, then the average number of questions required obviously lies between $\lfloor \log_2 |\mathcal{X}| \rfloor$ and $\lceil \log_2 |\mathcal{X}| \rceil$. When combining r independent realizations of the random variable X, the optimal average number of questions required to learn all the outcomes together lies between $\lfloor \log_2 |\mathcal{X}|^r \rfloor$ and $\lceil \log_2 |\mathcal{X}|^r \rceil$. Taking such combinations into account, we obtain for the entropy of X that

$$\log_2 |\mathcal{X}| - \frac{1}{r} < \frac{\lfloor \log_2 |\mathcal{X}|^r \rfloor}{r} \leq H(X) \leq \frac{\lceil \log_2 |\mathcal{X}|^r \rceil}{r} < \log_2 |\mathcal{X}| + \frac{1}{r}$$

for all $r \geq 1$, hence

$$H(X) = \log_2 |\mathcal{X}| \ . \tag{2}$$

Equation (2) is called *Hartley's formula* and gives the entropy of a uniformly distributed random variable.

We consider an example of a random variable Y that is *not* uniformly distributed. Let $\mathcal{Y} = \{a, b, c, d\}$, with $P_Y(a) = 1/2$, $P_Y(b) = 1/4$, $P_Y(c) = P_Y(d) = 1/8$. We conclude from the above that two questions are always sufficient, hence $H(Y) \leq 2$. However, there is a better strategy of asking questions or equivalently, a prefix-free code with a shorter average codeword length, namely,

$$a \rightsquigarrow 0 \ , \quad b \rightsquigarrow 10 \ , \quad c \rightsquigarrow 110 \ , \quad d \rightsquigarrow 111 \ .$$

The average number of questions required when asking the bits of the codewords is

$$\frac{1}{2} \cdot 1 + \frac{1}{4} \cdot 2 + \frac{1}{8} \cdot 3 + \frac{1}{8} \cdot 3 = \frac{7}{4} \ (< 2) \ .$$

On the other hand, this code (or strategy of asking questions) is optimal. Note that in this example, the length of the codeword of a letter is $\log_2(1/p)$, where p is the probability of this letter. The quantity $\log_2(1/p)$ is sometimes called the *unexpectedness* of a elementary event with probability p.

A code is optimal if the length of every codeword is equal to the unexpectedness of the corresponding outcome. Hence, for a random variable X for which each probability p_i is of the form $p_i = 2^{-s_i}$ for an integer s_i, we have

$$H(X) = p_1 \log_2(1/p_1) + p_2 \log_2(1/p_2) + \cdots \ . \tag{3}$$

Equation (3) is called *Shannon's formula*, and is a generalization of Hartley's formula (2). By combining independent realizations of the random variable for the encoding, one obtains that this formula gives the entropy of any discrete random variable. The following definition was given by Shannon in 1948.

Definition 1. [**30**] The *entropy* $H(X)$ of a random variable X with distribution P_X is given by

$$H(X) = H(P_X) := \sum_{x \in \mathcal{X}} -P_X(x) \cdot \log_2 P_X(x) = \mathrm{E}\left[-\log_2 P_X\right] .$$

∘

The *joint entropy* of random variables $X_1, X_2, \ldots, X_N$ is the entropy of the joint distribution, i.e.,

$$H(X_1 X_2 \cdots X_N) := H(P_{X_1 X_2 \cdots X_N}) \ .$$

Moreover, Definition 1 also covers the case where the distribution is conditioned on an event $\mathcal{A}$. We write $H(X|\mathcal{A}) := H(P_{X|\mathcal{A}})$ or, if $\mathcal{A} = \{Y = y\}$,

$$H(X|Y = y) := H(P_{X|Y=y}) \ .$$

The entropy of a *binary* random variable with probability distribution $[p, 1-p]$ is given by the *binary entropy function*

$$h(p) := -p \log_2 p - (1-p) \log_2 (1-p)$$

(see Figure 1).

The entropy of a random variable X is always non-negative and upper bounded by the binary logarithm of the cardinality of the range, i.e.,

$$0 \leq H(X) \leq \log_2 |\mathcal{X}| \ . \tag{4}$$

The second inequality, which is intuitively clear when taking into account the discussion above, follows from Jensen's inequality for *concave* functions:

$$H(X) = \mathrm{E}\left[\log_2(1/P_X)\right] \leq \log_2(\mathrm{E}\left[1/P_X\right]) = \log_2 |\mathcal{X}| \ .$$

Equality on the left hand side of (4) holds if and only if there exists an element $x_0 \in \mathcal{X}$ with $P_X(x_0) = 1$, whereas equality on the right hand side is equivalent to the fact that X is uniformly distributed over $\mathcal{X}$, i.e., that $P_X(x) = 1/|\mathcal{X}|$ holds

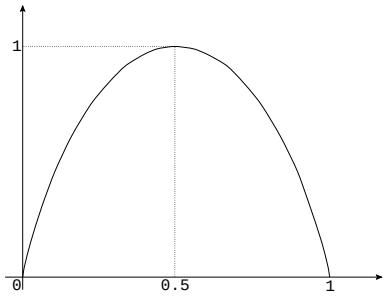

Fig. 1. The Binary Entropy Function

for all $x \in \mathcal{X}$. In the special case where the outcome of the random experiment is a binary string of length n, the second inequality of (4) implies that the entropy of the random variable can be equal to, but not exceed n.

For random variables X and Y, we have

$$H(XY) \leq H(X) + H(Y) \ , \tag{5}$$

with equality if and only if X and Y are statistically independent.

2.3 Conditional Entropy and Mutual Information

When considering inequality (5) it appears natural to interpret the (non-negative) quantity $H(XY) - H(X)$ as the entropy of the random variable Y when X is given.

Definition 2. The *conditional entropy of* Y *when given* X is defined as

$$H(Y|X) := H(XY) - H(X) \ . \tag{6}$$

$\circ$

Note that in contrast to all previously introduced entropies such as $H(X) = H(P_X)$, $H(XY) = H(P_{XY})$, or $H(Y|X = x) = H(P_{Y|X=x})$, the conditional entropy $H(Y|X)$ is *not* the entropy of a specific probability distribution, but rather the expected value of the entropies $H(Y|X = x)$, i.e.,

$$H(Y|X) = \mathrm{E}_X[H(Y|X = x)] \ .$$

Equation (6) can be rewritten as

$$H(XY) = H(X) + H(Y|X) \ .$$

This *chain rule* can be generalized as follows. For random variables $X_1, \ldots, X_N$ and an event $\mathcal{A}$ we have

$$H(X_1 X_2 \cdots X_N|\mathcal{A}) = H(X_1|\mathcal{A}) + H(X_2|X_1, \mathcal{A}) + \cdots + H(X_N|X_1 X_2 \cdots X_{N-1}, \mathcal{A}).$$

It is a fundamental property of the conditional entropy that

$$H(Y|X) \leq H(Y) \ , \tag{7}$$

which is a consequence of inequality (5). (However, note that $H(Y|X = x) > H(Y)$ is possible, as the following example illustrates. Let Y be 100 independent flips of an unfair coin with Prob["heads"] $= 99.9\%$, and let X be the number of "heads" in the sequence. Then, although of course $H(Y|X) < H(Y)$ holds, we have

$$1.141 \approx 100 \cdot h(0.999) = H(Y) < H(Y|X = 50) = \log_2 \left(\binom{100}{50} \right) \approx 96.35 \ .$$

Of course the event $\{X = 50\}$ is extremely unlikely.)

Informally spoken, inequality (7) can be interpreted as the fact that information can never increase uncertainty. More precisely, the quantity

$$I(Y;X) := H(Y) - H(Y|X) = H(X) + H(Y) - H(XY) \geq 0 \tag{8}$$

is the amount of information that X gives about Y. The last expression of (8) shows that $I(Y;X)$ is symmetric in its arguments, i.e., that

$$I(X;Y) = I(Y;X)$$

holds. The quantity $I(X;Y)$ is called *the mutual information between X and Y*. Analogously, one can define $I(X;Y|\mathcal{A}) := H(X|\mathcal{A}) - H(X|Y,\mathcal{A})$ and

$$I(X;Y|Z) := H(X|Z) - H(X|YZ) = \mathrm{E}_Z[I(X;Y|Z=z)] \ .$$

2.4 Graphical Representation of Information-Theoretic Quantities

Let X and Y be random variables. Then the quantities $H(XY)$, $H(X)$, $H(Y)$, $H(X|Y)$, $H(Y|X)$, and $I(X;Y)$ can be graphically represented as shown in Figure 2. The union of all inner regions corresponds to $H(XY)$. The representation

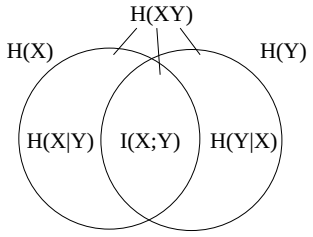

Fig. 2. Two Random Variables

has the property that the quantity corresponding to the disjoint union of some regions equals the sum of the quantities corresponding to these partial regions.

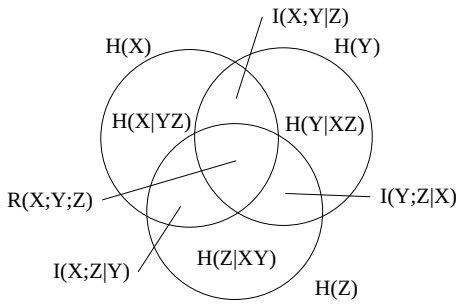

Fig. 3. Three Random Variables

For a detailed discussion of this measure-theoretic representation of information-theoretic quantities see [37].

The case of three random variables is shown in Figure 3. Note that the quantity corresponding to the region in the middle,

$$R(X;Y;Z) := I(X;Y) - I(X;Y|Z) \ ,$$

is symmetric in X, Y, and Z and can be negative. All the other regions represent information-theoretic quantities that are always non-negative.

Figure 4 illustrates independent symmetric bits X and Y and $Z := X \oplus Y$. Figure 5 shows a Markov chain.

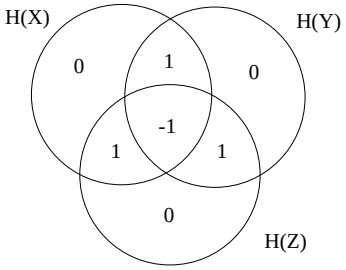

Fig. 4. $Z = X \oplus Y$

3 Perfect Secrecy and Shannon's Pessimistic Theorem

In the following we consider the problem of information-theoretically secure key generation and message transmission over an insecure channel. This section contains Shannon's definition of perfect secrecy of a cipher and his well-known theorem which appears to imply that unconditional security is necessarily completely impractical. In the following sections however it is demonstrated that

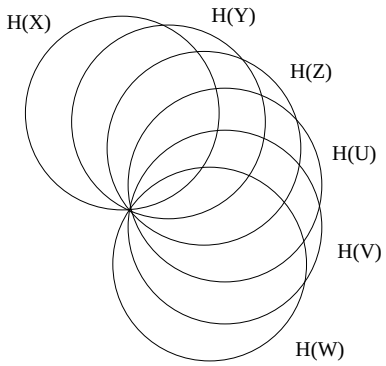

Fig. 5. A Markov Chain $X \to Y \to Z \to U \to V \to W$

information theory cannot be used only to prove such pessimistic results. It is somewhat surprising that when the models and security requirements are only slightly modified, then practical information-theoretic security can be achieved in many realistic scenarios.

Let us start with the classical scenario of a symmetric cryptosystem with message M, key K, and ciphertext C (see Figure 6). The following security

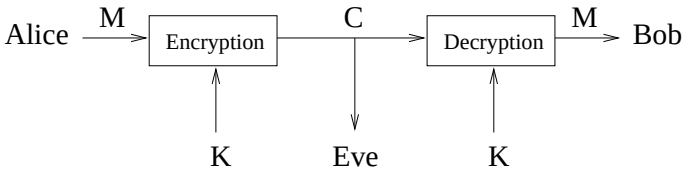

Fig. 6. A Symmetric Cryptosystem

definition appears to be the strongest possible for such a cryptosystem.

Definition 3. [29] A cipher is called *perfectly secret* if the ciphertext reveals no information about the message, i.e., if $I(M;C) = 0$ holds. ◦

Equivalent characterizations of this condition are that M and C are statistically independent, or that the best strategy of an eavesdropper who wants to obtain (information about) the message from the ciphertext is to use only the a priori knowledge about M and to discard C.

Perfect secrecy can even be achieved without any computation, as the example in Figure 7 shows. As everyone can easily see, the ciphertext alone reveals no information about the message at all in this example! (For more on "visual cryptography," see [26].)

This visual cipher is a graphical implementation of the *one-time pad* that was already proposed by Vernam in 1926 [34]. Here, the message is a string

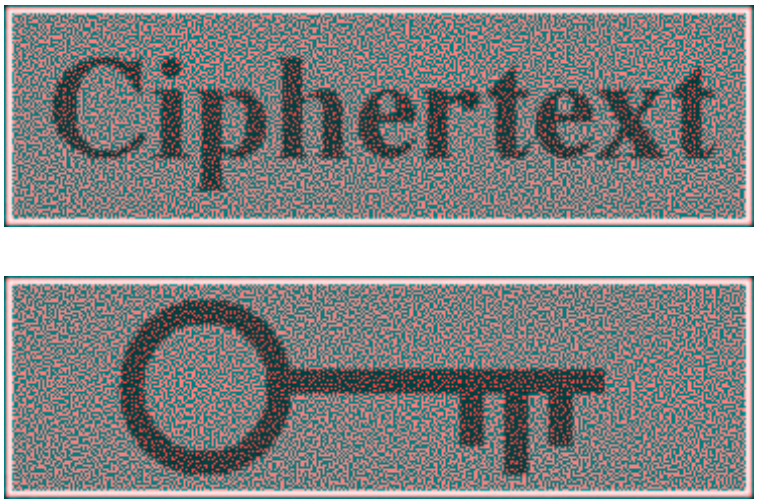

Fig. 7. Visual Decryption

$M = [m_1, m_2, \ldots, m_N]$ of length N, and the key is a uniformly distributed N-bit string $K = [k_1, k_2, \ldots, k_N]$ which is independent of M. The ciphertext C is computed from M and K by

$$C = [c_1, c_2, \ldots, c_N] = [m_1 \oplus k_1, m_2 \oplus k_2, \ldots, m_N \oplus k_N] =: M \oplus K \ .$$

The one-time pad is perfectly secret. To see this, observe first than when given the cleartext and the ciphertext, then the key is uniquely determined, i.e., $H(K|MC) = 0$. Furthermore, $I(K;C|M) = N$ (remember than N is the block length) follows then from $H(K) = N$ and $I(M;K) = 0$. Finally, $I(M;C) = 0$ holds because $H(C) \leq \log_2 |\mathcal{C}| = N$. A graphical representation of the quantities is given in Figure 8.

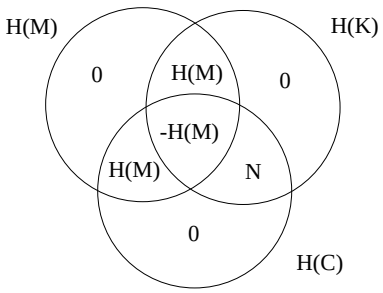

Fig. 8. Perfect Secrecy of the One-Time Pad

Unfortunately, the price one has to pay here for perfect secrecy is that the

communicating parties must share a secret key which is at least as long as the message (and can only be used once). In view of this property, the one-time pad appears to be quite impractical and can only offer an advantage in time: the key can be safely transmitted whenever this is possible, and the message can be secretly sent whenever this is needed.

However, Shannon showed that perfect secrecy cannot be obtained in a cheaper way, i.e., that the one-time pad is optimal with respect to key length.

Theorem 4. [29] *For every perfectly secret cryptosystem (with unique decodability), we have*

$$H(K) \geq H(M) \ .$$

For a proof of Shannon's theorem, note first that unique decodability means $H(M|CK) = 0$. The graphic representation of the involved quantities is given

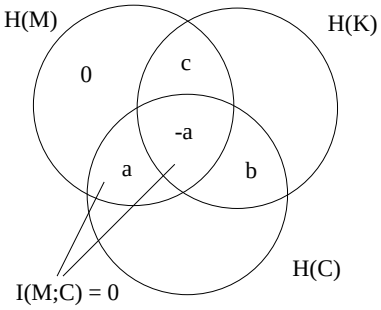

Fig. 9. The Proof of Shannon's Theorem

in Figure 9. We have $b \geq a$ because $I(C;K) \geq 0$, and

$$H(K) \geq b - a + c \geq a - a + c = H(M) \ .$$

This concludes the proof.

4 Optimistic Results by Limiting the Adversary's Information

Unfortunately, Shannon's theorem implies that perfect secrecy is possible only between parties who share a secret key of length at least equal to the entropy of the message to be transmitted. Hence every perfectly secret cipher is necessarily as impractical as the one-time pad. On the other hand, the assumption that the adversary has a *perfect* access to the ciphertext is overly pessimistic and unrealistic in general, since every transmission of a signal over a physical channel is subject to noise.

Motivated by this, many models have been presented and analyzed in which the information the adversary obtains is limited in some way, and which offer the possibility of information-theoretically secure key agreement and, under the assumption that insecure channels are always available, secret message transmission (using the one-time pad with the generated secret key).

The condition that the opponent's knowledge is bounded can for instance be based on noise in communication channels [36],[12],[1],[21], on the fact that the adversary's memory is limited [22],[9], or on the uncertainty principle of quantum mechanics [2]. In this article, we describe a number of models that belong to the first category.

4.1 Wyner's Wire-Tap Channel

Consider the following (simple but generally unrealistic) situation first. Assume that two parties Alice and Bob are connected by an authentic and noiseless binary channel, and that a wiretapper Eve receives the bits sent over the channel with some error probability $\varepsilon > 0$. In other words, her wire-tap channel is a *binary symmetric channel (BSC)* with error probability ε (see Figure 10).

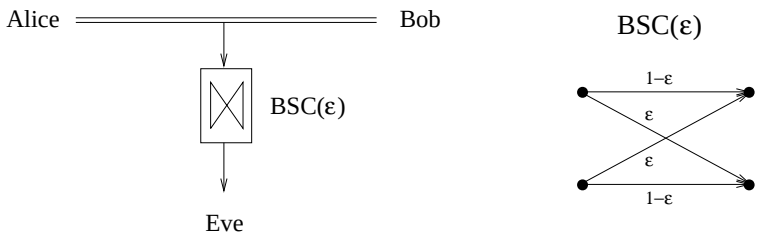

Fig. 10. A Binary-Symmetric Wire-Tap Scenario

In this situation, Alice can send a message bit M to Bob by sending an N-bit block $[X_1, X_2, \ldots, X_N]$, where $X_1, X_2, \ldots, X_{N-1}$ are independent and symmetric bits and X_N is such that

$$X_1 \oplus X_2 \oplus \cdots \oplus X_N = M \ .$$

Eve's error probability when guessing the bit M with the optimal strategy is

$$p = \frac{1 - (1 - 2\varepsilon)^N}{2} \ ,$$

and converges to $1/2$ exponentially fast in N. Moreover, the information that Eve obtains about M from the noisy versions of $X_1, X_2, \ldots, X_N$ does not exceed $1 - h(p)$. By repeating this process, Alice and Bob can agree on a highly secret key of arbitrary length.

The following, more general scenario of the *wire-tap channel* (see Figure 11) was introduced and analyzed by Wyner [36] and simplified by Massey [16]. In

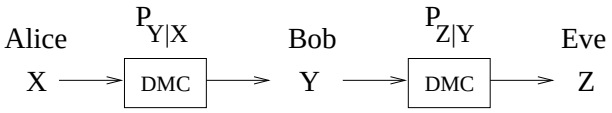

Fig. 11. Wyner's Wire-Tap Channel

this setting, Alice and Bob are connected by a discrete memoryless channel (characterized by its conditional probability distribution $P_{Y|X}$), whereas Eve receives a noisy version Z of Bob's channel output Y. Alice chooses the input to the first channel according to some distribution P_X.

It was shown in [36] that in this scenario, Alice and Bob can agree on a highly secret key at a some rate in many situations (for instance in the case where all the random variables are binary and the channels are binary-symmetric with error probabilities not $1/2$ and not 0 nor 1, respectively). Exact definitions of the security requirements to such a key, as well as of the secret-key generation rate, are given below.

However, the assumption that the adversary only receives a degraded version of the legitimate receiver's information is unrealistic in general. This fact motivated the study of generalizations of Wyner's model.

4.2 Broadcast Channels

Csiszár and Körner [12] considered the situation where the sender Alice is connected to the receiver Bob by a discrete memoryless channel (with conditional distribution $P_{Y|X}$), and where also the adversary Eve receives a noisy version Z of X over a different channel (characterized by $P_{Z|XY}$, i.e., the channels are not necessarily independent). As before, Alice chooses the channels' input X according to some distribution P_X. The broadcast scenario is illustrated in Figure 12.

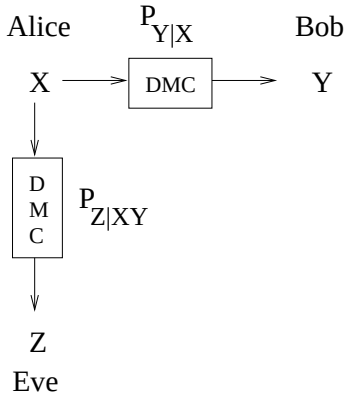

Fig. 12. The Broadcast-Channel Scenario

For this setting, the *secrecy capacity* $C_S(P_{YZ|X})$ has been defined as the maximal rate at which Alice and Bob can generate a virtually secret key. Without going into the details of the definitions and the key-generation protocols, we remark that both the size of the generated secret key as well as the amount of information leaked to the adversary are defined in terms of a *rate*, i.e., measured as average information *per channel use*.

In [12], the following lower bound on the secrecy capacity, depending on the conditional distribution $P_{YZ|X}$, has been proved:

$$C_S(P_{YZ|X}) \geq \max_{P_X} [I(X;Y) - I(X;Z)] \ . \tag{9}$$

In equality (9), the maximum is taken over all possible distributions P_X of X. Intuitively, this condition implies that if the legitimate partners initially have some advantage over Eve in terms of the information about each other's random variables, then this advantage can be fully exploited to generate a secret key.

However, if Alice and Bob have no such advantage to start with, then generally no secret-key agreement is possible in this model. Let us for instance consider the situation where the channels are independent and binary-symmetric with error probabilities ε and δ (see Figure 13). In this special scenario, the secrecy

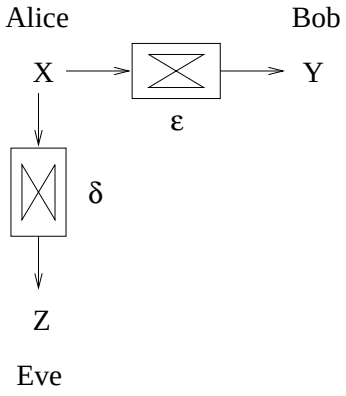

Fig. 13. Independent Binary-Symmetric Broadcast Channels

capacity is given by

$$C_S(\varepsilon, \delta) = \begin{cases} h(\delta) - h(\varepsilon) & \delta > \varepsilon \\ 0 & \text{otherwise} \ . \end{cases}$$

In other words, secret-key agreement is impossible unless Bob's channel is better than Eve's. Unfortunately, it may often be impossible to guarantee that the adversary's channel is noisier than the one of the legitimate partner.

4.3 The Power of Interaction

The following example, given in [21], illustrates how much more powerful interaction can be in contrast to one-way transmission for unconditionally secure key agreement. This is a motivation for the study of a more general model of secret-key agreement from common information by insecure two-way communication. We discuss this model in Section 5.

We start with the situation shown in Figure 13, where $0 < \delta \leq \varepsilon < 1/2$. As mentioned above, no secret-key agreement is possible. However, let us assume an *interactive* variant of this model with an additional noiseless and insecure but authentic channel. (Note that channels with virtually these properties often exist in reality, e.g., telephone lines.) Surprisingly, the situation is now entirely different although the additional channel can be perfectly overheard by Eve.

Observe first that the additional public-discussion channel allows to invert the direction of the noisy channel between Alice and Bob by the following trick. First, Alice chooses a random bit X and sends it over the noisy channel(s). This bit is received by Bob as Y and by Eve as Z. Bob, who wants to send the message bit C to Alice, computes $C \oplus Y$ and sends this over the noiseless public channel. Alice computes $(C \oplus Y) \oplus X$, whereas Eve can compute $(C \oplus Y) \oplus Z$. This perfectly corresponds to the situation where the direction of the main channel is inverted (see Figure 14).

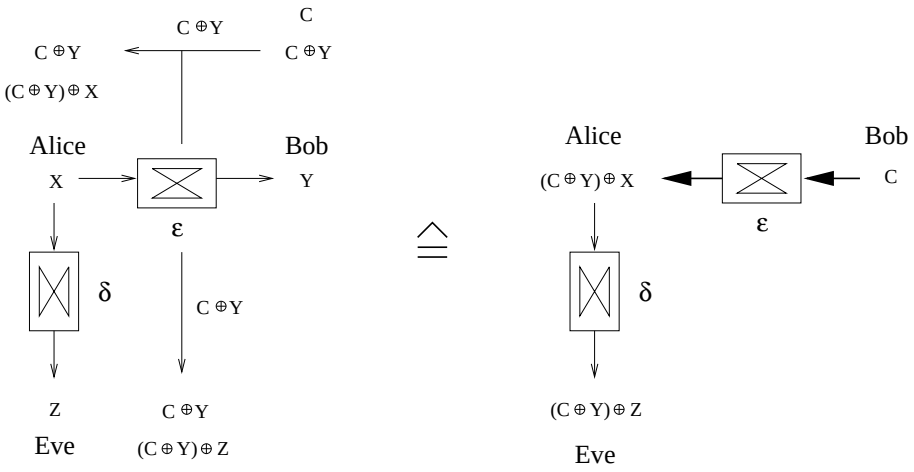

Fig. 14. Inverting the Main Channel

The second crucial observation is that this is exactly the binary-symmetric setting of Wyner's wire-tap channel of Section 4.1, allowing secret-key agreement at some rate. We conclude from this example that the possibility of feedback from Bob to Alice can substantially improve the legitimate partners' situation towards a wire-tapping adversary.

5 Interactive Secret-Key Agreement from Common Randomness

5.1 The Scenario and the Secret-Key Rate

Maurer has proposed the following interactive model of secret-key agreement by public discussion from common information [21]. The parties Alice and Bob who want to establish a mutual secret key have access to realizations of random variables X and Y, respectively, whereas the adversary knows a random variable Z. Let P_{XYZ} be the joint distribution of the random variables. Furthermore, the legitimate partners are connected by an insecure but authentic channel, i.e., a channel that can be passively overheard by Eve but over which no undetected active attacks by the opponent, such as modifying or inserting messages, are possible (see Figure 15).

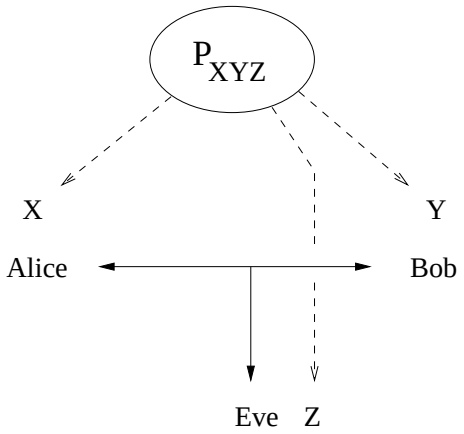

Fig. 15. Secret-Key Agreement by Public Discussion from Common Information

Note that it is natural to consider this model by the following reasons. First, it is an interactive (i.e., two-way) generalization of Wyner's and Csiszár and Körner's models. It is not necessary to assume the existence of noisy communication channels in this interactive setting because equivalents of such channels can be obtained by the same trick as shown in Section 4.3 for inverting the binary-symmetric channel. Secondly, the assumption that the parties have access to correlated randomness appears to be realistic in many contexts. An example of a possible physical implementation is described in Section 5.2.

In analogy to the previous models, where the channels could be used many times independently, we assume here that the parties have access to a number of independent realizations of the corresponding random variables. Consequently, the so-called *secret-key rate* is defined in this model as the maximal rate at which Alice and Bob can generate a highly secret key by communication over the

insecure channel, where the required number of channel uses from the definition of the secrecy capacity is replaced by the amount of randomness (i.e., the number of realizations of X and Y) necessary for the generation of a key of some length.

Definition 5. The *secret-key rate* $S(X;Y||Z)$ of the distribution P_{XYZ} is the maximal number R with the following property. For every $\varepsilon > 0$, there is a number N_0 such that for all $N \geq N_0$, a protocol exists that uses authenticated public discussion and satisfies the following conditions. (We denote the block of the first N realizations of the random variable X, $[X_1, X_2, \ldots, X_N]$, by X^N, and analogous for Y and Z. Furthermore, let U be the entire communication held over the public channel during the execution of the protocol.) There exist k-bit strings S and S' with

$$k > (R - \varepsilon)N \ , \tag{10}$$

$$H(S \,|\, X^N U) = 0 \ , \tag{11}$$

$$H(S' \,|\, Y^N U) = 0 \ , \tag{12}$$

$$\mathrm{Prob}\,[S \neq S'] < \varepsilon \ , \tag{13}$$

$$I(S;\, Z^N U) < \varepsilon \ , \tag{14}$$

$$H(S) > k - \varepsilon \ . \tag{15}$$

In other words, these conditions guarantee that Alice (11) and Bob (12) can generate almost uniformly distributed (15) keys of a certain length (10) that are equal with high probability (13) and about which the adversary has virtually no information (14). ∘

The notion of the secret-key rate is stronger than the one of secrecy capacity in the sense that in the definition of $C_S(P_{YZ|X})$, it was required that *the rate* at which Eve obtains information about the key is small, whereas here, the *total amount of information* about the entire key must be negligible. (However, one can show that the secret-key rates with respect to the weaker and the stronger definitions are equal [19].)

The secret-key rate is a quite fundamental and mathematically interesting property of a distribution P_{XYZ}. One challenging problem in this context is to enlighten the exact relationship between P_{XYZ} and $S(X;Y||Z)$, i.e., to determine the secret-key rate of a given distribution, or at least to decide whether the rate is non-zero and secret-key agreement is possible in principle in a particular situation. We discuss these questions in Section 6.

5.2 The Satellite Scenario and Phases of Secret-Key Agreement Protocols

The following realistic special scenario was proposed in [21] and completely analyzed in [25]. Assume that a satellite sends out random bits at very low signal power and that Alice, Bob, and Eve receive these bits over independent binary-symmetric channels with error probabilities α, β, and ε, respectively (see Figure 16).

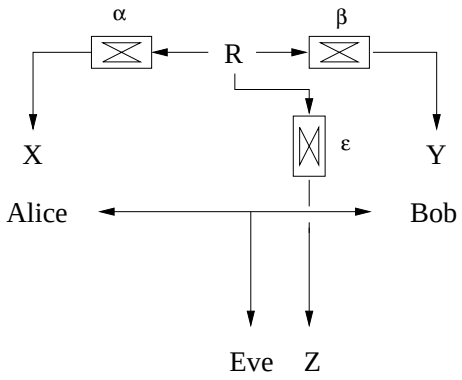

Fig. 16. The Satellite Scenario

In general, we may have to assume that Eve has a better antenna than the legitimate partners, and hence a possibly substantially lower error rate. It is a somewhat surprising fact that secret-key agreement is always possible in this scenario (unless Eve has a noiseless access to the satellite bits or either Alice or Bob obtains no information at all about these bits).

In the following, we describe a protocol for secret-key agreement in the satellite scenario. Such a protocol is often interpreted as consisting of three phases. As mentioned, Alice and Bob possibly start in a situation in which the adversary has an advantage over the legitimate partners with respect to the information about each other's random variables. The objective of the first phase, *advantage distillation*, is to generate an advantage over the opponent by exploiting the authenticity of the public channel. However, Alice and Bob do generally not share a mutual string after this phase. Hence, an interactive error-correction phase, *information reconciliation*, is required. Finally, the resulting mutual but only partially secret string must be transformed into a (shorter) highly secret string. This final phase is called *privacy amplification*. In the illustration of the three phases in Figure 17, the relations between the amounts of information that Bob's and Eve's knowledge provide about Alice's string are shown. The protocol steps are described in detail in the next three sections. An interactive demonstration of the phases is provided on the Internet [7].

5.3 Advantage Distillation

We assume the satellite scenario described in the previous section with error probabilities $0 \leq \alpha, \beta < 1/2$ and $0 < \varepsilon < \min\{\alpha, \beta\}$, i.e., the adversary has an initial advantage over the legitimate partners in terms of the error probabilities. Let us consider N independent realizations of the random variables. Then, we have

$$I(X^N;Y^N) = \sum_{i=1}^{N} I(X_i\,;\,Y_i) \;=\; N \cdot (1 - h(\alpha(1-\beta) + (1-\alpha)\beta)) \;,$$

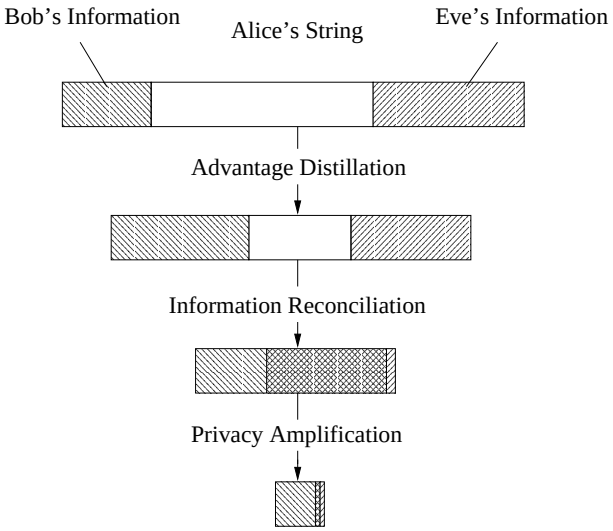

Fig. 17. Phases of a Secret-Key-Agreement Protocol

$$I(X^N;Z^N) = \sum_{i=1}^{N} I(X_i\,;\,Z_i) \;=\; N\cdot(1-h(\alpha(1-\varepsilon)+(1-\alpha)\varepsilon))\;,$$

$$I(Y^N;Z^N) = \sum_{i=1}^{N} I(Y_i\,;\,Z_i) \;=\; N\cdot(1-h(\beta(1-\varepsilon)+(1-\beta)\varepsilon))\;,$$

i.e.,

$$I(X^N;Y^N) < \min\left\{I(X^N;Z^N)\,,\,I(Y^N;Z^N)\right\}\;.$$

The basic idea of the advantage-distillation phase is that Alice and Bob use the noiseless discussion channel for exchanging information about their bits in an insecure but authentic way with the objective of identifying bits that are correct with a higher probability than others. We describe two different protocols that achieve this. The protocols are based on a repeat code and on the exchange of parity-check bits. The repeat-code protocol is simpler to describe, but very inefficient with respect to the required number of realizations of the random variables, whereas the parity-check protocol appears to be quite efficient. For a detailed analysis of the protocols, see for example [21],[20],[24].

Repeat-Code Protocol. The repeat-code protocol works as follows (see also Figure 18). Let N be a fixed parameter. Alice chooses a random bit C and computes

$$C^N \oplus X^N := [C\oplus X_1, C\oplus X_2, \ldots, C\oplus X_N]\;,$$

where C^N stands for the repeat-code block $[C, C, \ldots, C]$ of length N. She sends this "blinded" repeat-code block over the public channel. Bob computes from

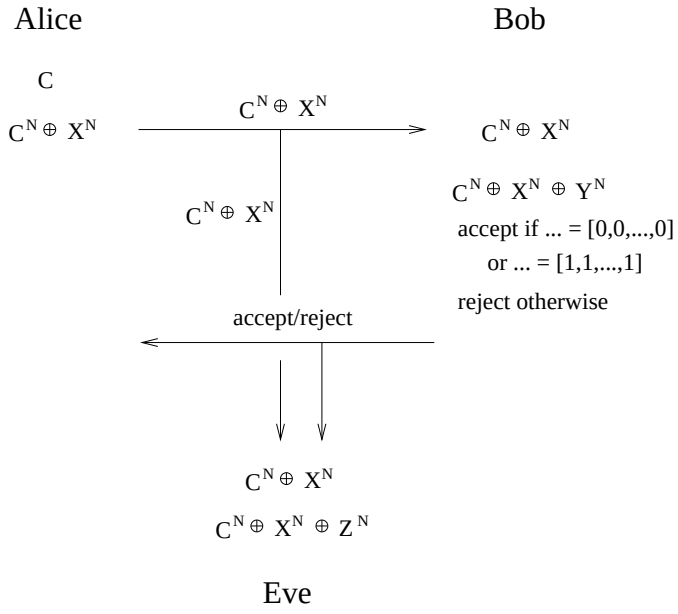

Fig. 18. The Repeat-Code Protocol

this the block $(C^N \oplus X^N) \oplus Y^N$, and sends an "accept" message over the discussion channel if and only if the resulting block is a repeat-code block $(C')^N = [C', C', \ldots, C']$. Note that this is exactly the situation where Alice and Bob have either the same bit in all positions, i.e., $X^N = Y^N$, or opposite values in each position, i.e., $X^N = Y^N \oplus 1^N$. It is intuitively clear that Alice and Bob not only obtain an arbitrarily low probability of the event that $C' \neq C$ for large N this way, but also that they improve their position compared to the opponent by accepting only in situations of apparently highly reliable transmission. However, also the adversary Eve, who can compute $(C^N \oplus X^N) \oplus Z^N$, takes advantage of a greater value of N. It is a somewhat surprising result that for *arbitrary* values of $\alpha, \beta < 1/2$, and $\varepsilon > 0$, Alice and Bob end up in an advantageous situation (both with respect to the error probabilities and to the information about each other's strings) for sufficiently large N.

We show this with respect to the error probabilities of Bob and Eve when guessing the bit C for the special case $\alpha = \beta$ (which we can assume without loss of generality because noise can always be added). We denote by α_{be} the probability that the single bit 0 sent by Alice over the conceptual channel (i.e., $C \oplus X$ is sent over the public channel) is received (i.e., decoded) by Bob as b and by Eve as e. Then we have

$$\begin{aligned}
\alpha_{00} &= (1-\alpha)^2(1-\varepsilon) + \alpha^2\varepsilon \ , \\
\alpha_{01} &= (1-\alpha)^2\varepsilon + \alpha^2(1-\varepsilon) \ , \\
\alpha_{10} &= \alpha_{11} \ = (1-\alpha)\alpha \ .
\end{aligned}$$

We assume that N is an even integer. The probability γ that Bob accepts the N-bit block sent by Alice and that $C' \neq C$ holds is

$$\gamma = (\alpha_{10} + \alpha_{11})^N \ ,$$

whereas the probability δ that Bob accepts and Eve guesses the bit incorrectly is lower bounded by $1/2$ times the probability of the event that the block $(C^N \oplus X^N) \oplus Z^N$ which Eve obtains consists of $N/2$ 0's and the same number of 1's, i.e.,

$$\delta \geq \frac{1}{2} \binom{N}{N/2} \alpha_{01}^{N/2} \cdot \alpha_{01}^{N/2} \approx \frac{1}{2} (2\sqrt{\alpha_{00}\alpha_{01}})^N \ .$$

Clearly, the actual message bit C is statistically independent of the block Eve receives if this event occurs. It is not difficult to see that

$$2\sqrt{\alpha_{00}\alpha_{01}} > \alpha_{10} + \alpha_{11}$$

holds for $\alpha < 1/2$ and $\varepsilon > 0$, meaning that Bob's error probability decreases asymptotically faster than Eve's and is hence smaller for sufficiently large N. One can even show that Eve has less information than Bob about the bit C for sufficiently large N.

Parity-Check Protocol. The second protocol we discuss uses parity-check bits and works as follows. Alice computes the parity bit $X_1 \oplus X_2$ and sends it over the public channel. Bob accepts if and only if $X_1 \oplus X_2 = Y_1 \oplus Y_2$, i.e., if the parities of Alice's and Bob's first two bits are equal. In this case, the values X_1 and Y_1 are chosen by Alice and Bob, respectively, for the next protocol round (whereas otherwise, the bits are discarded). This step is repeated a number of times. After this first round it may be necessary, depending on the initial error probabilities, to carry out some additional rounds (see Figure 19).

It is not difficult to see that r rounds of the parity-check protocol are equivalent to the repeat-code protocol with 2^r-bit blocks with respect to the resulting error probabilities. However, it is obvious that the parity-check protocol is much more efficient.

5.4 Information Reconciliation

During advantage distillation, the partners Alice and Bob compute (possibly distinct) strings S_A and S_B, respectively, about which the adversary also has some information. At the end of the key-agreement protocol however, Alice's and Bob's strings must be equal and highly secure, both with overwhelming probability. The information-reconciliation phase consists of interactive error correction and establishes the first of these two conditions.

After advantage distillation, Bob has more information about Alice's string than Eve has, and after information reconciliation, Bob should exactly know Alice's string. (A more general condition would be that after information reconciliation, Alice and Bob share a string that is equally long as S_A and S_B.) This

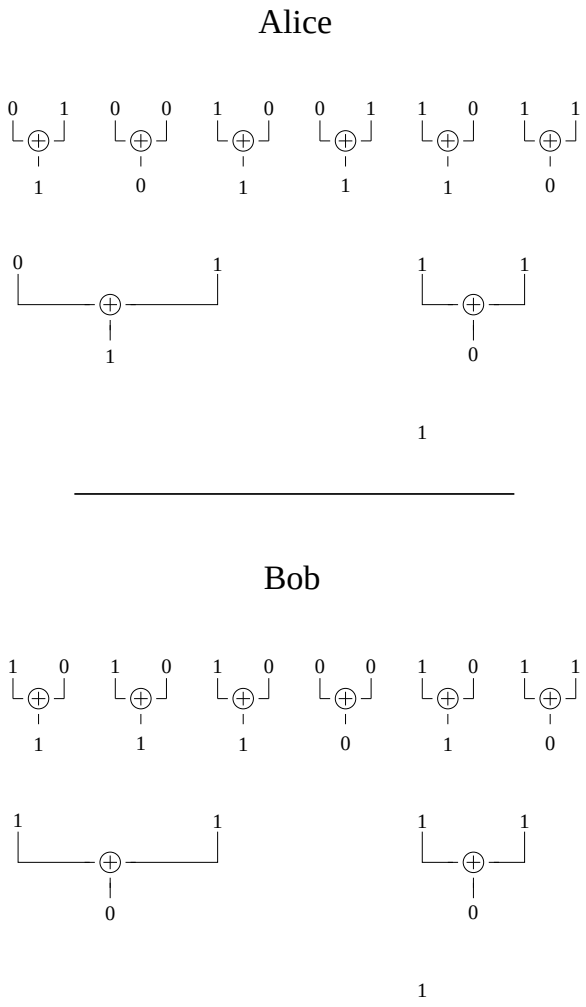

Fig. 19. Three Rounds of the Parity-Check Protocol

leads to a lower bound on the amount of error-correction information E that must be exchanged. Namely, Bob must know S_A completely with overwhelming probability when given S_B and E, i.e.,

$$0 \approx H(S_A \,|\, S_B, E) \geq H(S_A \,|\, S_B) - H(E) \ ,$$

and hence

$$H(E) \gtrsim H(S_A \,|\, S_B) \ .$$

On the other hand, the uncertainty of S_A from Eve's viewpoint can as well be reduced by $H(E)$ in the worst case when Eve learns E (see Figure 20).

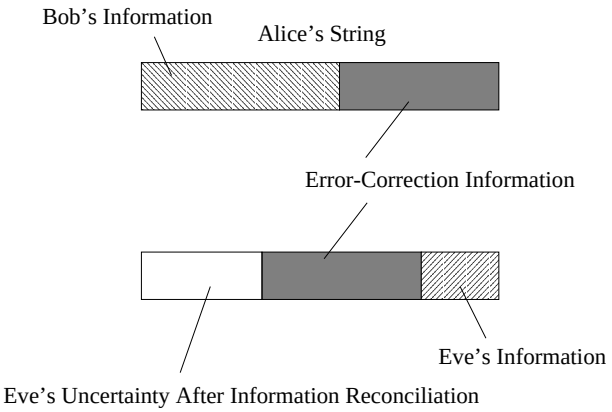

Fig. 20. The Effect of Information Leaked During Information Reconciliation

A good protocol for information reconciliation should both minimize the amount of information leaked to the adversary and be efficient. Examples of protocols satisfying both conditions are given in [6]. We sketch two examples of these protocols. The first is optimal with respect to the information leaked but completely inefficient, whereas the second protocol leaks more information but is more efficient as well. Let us assume that Alice and Bob have finished the advantage-distillation phase in the satellite model. In other words, Bob's string is a (good) estimate about Alice's string, i.e., the same string with a (small) number of errors.

Random-Label Protocol. The first, non-interactive, protocol works as follows. Alice randomly chooses a function f mapping $\{0,1\}^n \to \{0,1\}^m$ (where n is the length of S_A and S_B and m can be roughly equal to $H(S_A|S_B)$) among all such functions and sends (a description of) f together with $f(S_A)$ to Bob, who determines the string S'_A with minimal Hamming distance from S_B that satisfies $f(S'_A) = f(S_A)$. According to the discussion above, and because $m \approx H(S_A|S_B)$, this protocol is optimal with respect to the leaked information. However, it is completely inefficient, hence useless, by the following two reasons. First, the description of the random function f would require $m2^n$ bits. Furthermore, S'_A cannot be efficiently determined from f, $f(S_A)$, and S_B.

Binary-Search Protocol. The idea of the second protocol is to interactively detect the positions where Alice's and Bob's strings differ and to correct these errors. Alice and Bob start by comparing the parity bit, i.e., the XOR-sum, of the bits in randomly but identically chosen substrings S'_A and S'_B of S_A and S_B, respectively. If there are bit errors between the strings S_A and S_B, then the resulting parity bits differ with probability $1/2$ over the choice of the substrings.

If the parities are different, Alice and Bob have detected substrings containing an odd number of errors with respect to each other, and they can locate one of

them by partitioning the substring into two subsets of equal size, one of which clearly contains an odd number of errors as well (and has different parity sums). This splitting procedure is continued until the error is localized and can be corrected by Bob (see Figure 21).

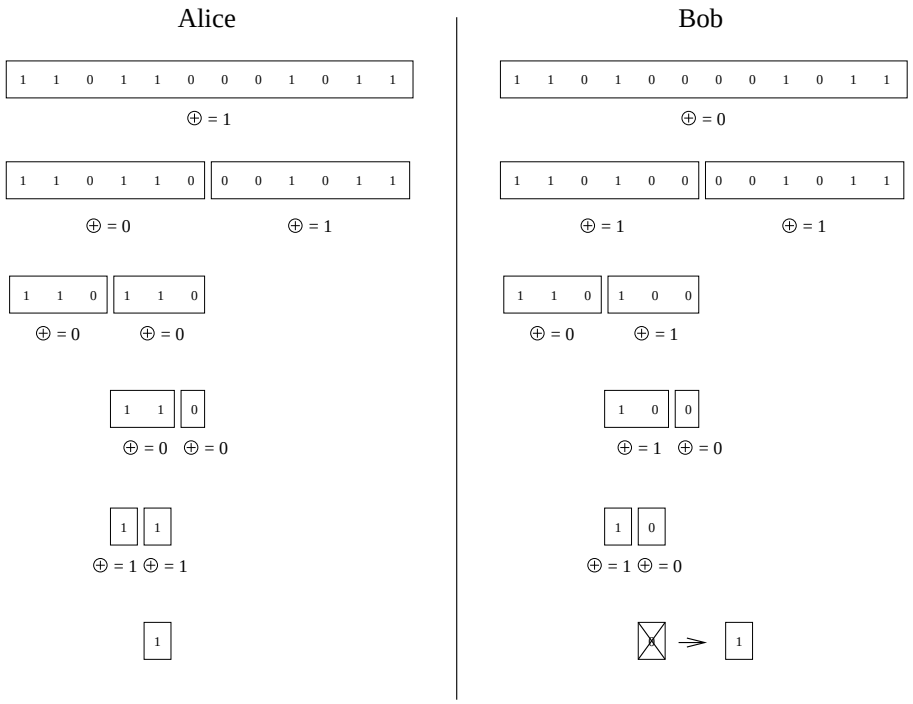

Fig. 21. Finding and Correcting an Error by Comparing Parities

Alice and Bob repeat this procedure until all the errors are found and corrected. If n is the length of the strings S_A and S_B, this protocol requires the exchange of $\lceil \log_2 n \rceil$ bits per error to be corrected. Hence it is efficient if the strings of Alice and Bob differ in only a few bit positions.

After information reconciliation, Alice and Bob have agreed on a mutual string S about which Eve has a possibly considerable amount of information consisting of both a priori knowledge but also information (e.g., physical bits or parities thereof) leaked during information reconciliation.

5.5 Privacy Amplification

Privacy amplification is the art of shrinking a partially secure string S to a highly secret string S' by public discussion. Hereby, the information of the ad-

versary about S can consist of physical bits, of parities thereof, or other types of information (see Figure 22).

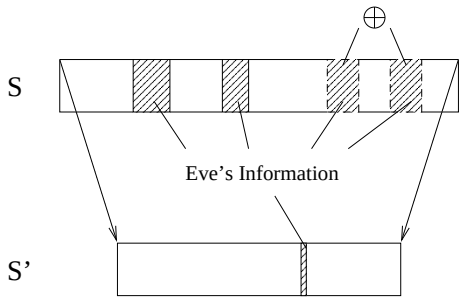

Fig. 22. Eliminating Eve's Knowledge by Privacy Amplification

The following questions related to privacy amplification were studied and answered in [4],[3]. What is a good technique of computing S' from S? What is the possible length of S', depending on this shrinking technique and on the adversary's (type and amount of) information about S?

It is quite clear that the best technique would be to compute S' (of length r) from the n-bit string S by applying a random function $f : \{0,1\}^n \to \{0,1\}^r$. However, Alice and Bob would have to exchange $r2^n$ bits of information to agree on such a function. On the other hand, there exist relatively small classes of functions with "random-like" properties. Examples are so-called *universal classes* of hash functions, which turned out to be useful for privacy amplification.

Definition 6. A class $\mathcal{H}$ of functions h mapping a set $\mathcal{A}$ to a set $\mathcal{B}$ is called *universal* if for all $x, y \in \mathcal{A}$, $x \neq y$, we have

$$\text{Prob}_{h \in_r \mathcal{H}}[h(x) = h(y)] = \frac{1}{|\mathcal{B}|} ,$$

where $h \in_r \mathcal{H}$ stands for the fact that h is chosen randomly in $\mathcal{H}$ according to the uniform distribution. In other words, a function that is chosen randomly from a universal class behaves like a completely random function with respect to *collisions.* ∘

An example of a universal class of functions, mapping $\{0,1\}^n$ to $\{0,1\}^r$, of cardinality $2^{n \cdot r}$ are the *linear* functions. There exist even smaller classes. For more examples and lower bounds on the size of universal classes, see for example [33].

We analyze the following type of privacy amplification protocols. First, Alice chooses a random function h from a fixed universal class $\mathcal{H}$ of hash functions mapping n-bit strings to r-bit strings for some r to be determined, and sends (the description of) h publicly to Bob, i.e., also Eve learns h. Then Alice and Bob both compute $S' := h(S)$.

Let us consider the question how long the virtually secure string S' can be, depending on the type and amount of Eve's knowledge about S. Note first that the fact that Eve has some information about a string S is another way of saying that given Eve's entire knowledge $U = u$ about S, the random variable S is *not* uniformly distributed, i.e.,

$$H(S \,|\, U = u) < n \ .$$

In this case we say that Eve has $n - H(S|U = u)$ bits of (Shannon-) information about S. Because the resulting string S' must satisfy

$$H(S' \,|\, C, U = u) \approx r$$

(where r is the length of S' and C is the communication held over the public channel), privacy amplification can be interpreted as "distribution smoothing."

Intuitively, one might think that if Eve has t bits of information about S, then the length r of the resulting string S' can be roughly $n - t$ (see Figure 23). This fact was shown to be correct if Eve has *deterministic* information about

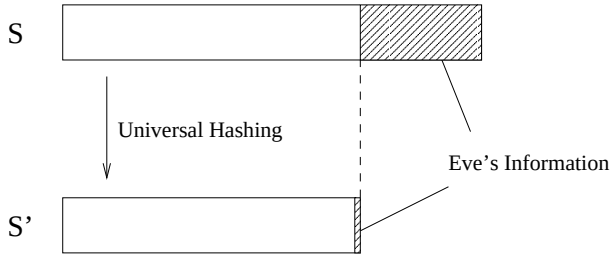

Fig. 23. Can Eve's Knowledge Be Simply Cut Away by Universal Hashing?

S, i.e., if Eve knows the value $g(S)$ for some fixed function g [4]. However, if Eve's information is not deterministic, it is *not* true in general that $n - t$ secure bits can be extracted when Eve has t bits of Shannon information about S, as the following example shows. Let $P_{S|U=u}(s_0) = 1/2$ for some $s_0 \in \{0,1\}^n$, and $P_{S|U=u}(s) = 1/(2 \cdot (2^n - 1))$ for all n-bit strings $s \neq s_0$. Then, we have $H(S|U = u) \approx n/2$, but no secure string S' (of any length, let alone $n/2$) can be extracted because Eve precisely knows S, hence also $S' = h(S)$, with probability $1/2$ (where h is the randomly chosen hash function). This means that S' cannot be highly secure.

The answer to the question what a suitable information (or entropy) measure is with the property that the above intuition (illustrated in Figure 23) is true, was given in [3] as follows.

Definition 7. For a random variable X with distribution P_X, the *collision probability* $P_C(X)$ is defined as

$$P_C(X) := \sum_{x \in \mathcal{X}} P_X(x)^2 \ .$$

The *collision entropy* or *Rényi entropy (of order 2) of* X is

$$H_2(X) := -\log_2(P_C(X)) = -\log_2(\sum_{x \in \mathcal{X}} P_X(x)^2) \ .$$

$\circ$

The collision probability is the probability that two independent realizations of the random variable X show the same value. Equivalently, it is the probability of guessing a realization of X correctly with the optimal strategy on the basis of an independent realization of X, where the distribution of X is unknown. Jensen's inequality implies

$$H_2(X) = -\log_2(\mathrm{E}\,[P_X]) \leq \mathrm{E}\,[-\log_2 P_X] = H(X) \ .$$

It was shown that Rényi entropy is a good information measure in the context of privacy amplification by universal hashing. Theorem 8 (see also Figure 24) implies that the intuitive fact illustrated in Figure 23 is true with respect to Rényi instead of Shannon information.

Theorem 8. [3] *Let S be an n-bit string with conditional distribution $P_{S|U=u}$ (given Eve's knowledge $U = u$ about S) and Rényi entropy $H_2(S|U = u)$, let G be the random variable corresponding to the random choice (with uniform distribution) of a member g of a universal class $\mathcal{H}$ of hash functions mapping n-bit strings to r-bit strings, and let $S' = G(S)$. Then*

$$r \geq H(S'|G, U = u) \geq H_2(S'|G, U = u) \geq r - \frac{2^{r - H_2(S\,|\,U=u)}}{\ln 2} \ .$$

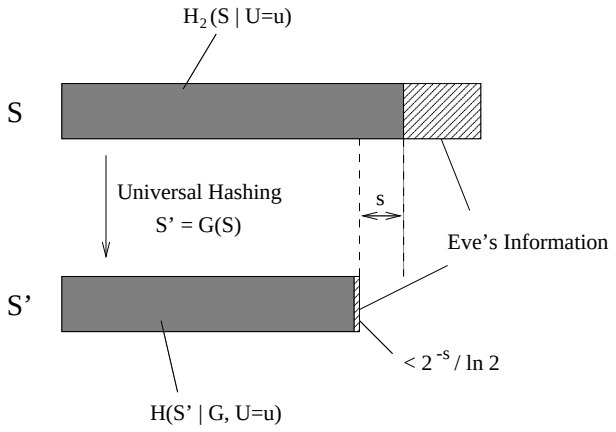

Fig. 24. Rényi Entropy Can Be Extracted by Universal Hashing

Intuitively, Theorem 8 states that if the length r of S' is chosen as

$$r := H_2(S \,|\, U = u) - s \ ,$$

where s is a security parameter, then the resulting string S' is highly secret, where the security increases exponentially in s.

Note that this result in not tight and can be improved in many cases. One reason for this is the counterintuitive fact that Rényi entropy can be increased by giving side information, so called *spoiling knowledge*. By using this property it was shown in [8] that Rényi entropy of order α, for $1 < \alpha < 2$, is a good measure with respect to privacy amplification as well.

One important question finally concerns the influence of the information exchanged during the information-reconciliation phase on the Rényi entropy of S from Eve's point of view, hence on the length of the key that can finally be generated. It was shown in [8] that learning r physical bits cannot reduce Rényi entropy by significantly more than r but with negligible probability.

6 Generalizing the Model

The scenario where the parties receive independent noisy versions of the same random source's signal was completely analyzed in [25],[23]. Possible real-world realizations of the required information source are a satellite sending random bits at low signal power, a pulsar, a deep-space radio source, or randomly polarized photons. However, many more general scenarios can be thought of where the parties receive a different type of correlated information. The assumptions that the parties obtain noisy versions of a common signal or that they have access to a great number of independent realizations of the same random experiment can be modified or dropped. An example is the scenario where Alice, Bob, and Eve obtain a number of playing cards from the same stack [15]. As another generalization, the adversary can be assumed to be more powerful. For instance, it may often be unrealistic to guarantee that the opponent is only a passive wire-tapper.

6.1 Arbitrary Random Variables

Let us have a closer look at the scenario of arbitrary correlated information, i.e., of an arbitrary random experiment P_{XYZ} with many independent realizations (see Figure 15). Note that this is exactly the setting for which the secret-key rate $S(X;Y||Z)$ is defined. In this general case it is a fundamental and natural problem to determine $S(X;Y||Z)$ for a given distribution P_{XYZ}, or at least to decide whether the quantity is non-zero. The following bounds depend on information-theoretic quantities directly derived from P_{XYZ}. The lower bound

$$\max\{I(X;Y) - I(X;Z)\,, I(Y;X) - I(Y;Z)\} \leq S(X;Y||Z)$$

is a consequence of the above-mentioned result by Csiszár and Körner [12] and states that an existing advantage over the adversary can be fully (and even

non-interactively) exploited to generate a secret key. As shown in the previous sections, this bound is *not* tight: Secret-key agreement can also be possible in scenarios where Alice and Bob start in a "bad" situation. On the other hand, the following upper bound was shown in [21]:

$$S(X;Y||Z) \leq \min\{I(X;Y)\,,\, I(X;Y|Z)\}\ . \tag{16}$$

The bound (16) is quite intuitive and states that Alice and Bob cannot extract a larger amount of secret key than the mutual information between their random variables X and Y (with and without giving Eve's random variable Z). However, this bound is not tight neither and can be improved as follows. Trying to reduce the quantity $I(X;Y|Z)$, the adversary Eve can send the random variable Z over a channel, characterized by $P_{\overline{Z}|Z}$, in order to generate the random variable $\overline{Z}$. Clearly,

$$S(X;Y||Z) \leq S(X;Y||\overline{Z}) \leq I(X;Y|\overline{Z}) \tag{17}$$

holds for every such $\overline{Z}$. This motivates the following definition of a new conditional information measure, *the intrinsic conditional mutual information between X and Y when given Z*, which is the infimum of $I(X;Y|\overline{Z})$, taken over all discrete random variables $\overline{Z}$ that can be obtained by sending Z over a channel, characterized by $P_{\overline{Z}|Z}$. The situation is illustrated in Figure 25. (Note that $R(X;Y;Z) \geq 0$ always holds for the particular $\overline{Z}$ which minimizes $I(X;Y|\overline{Z})$.)

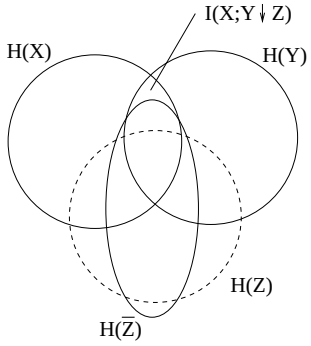

Fig. 25. The Intrinsic Conditional Information

Definition 9. For a distribution P_{XYZ}, the *intrinsic conditional mutual information between X and Y when given Z*, denoted by $I(X;Y\downarrow Z)$, is

$$I(X;Y\downarrow Z) := \inf\left\{ I(X;Y|\overline{Z}) \ :\ P_{XY\overline{Z}} = \sum_{z\in\mathcal{Z}} P_{XYZ}\cdot P_{\overline{Z}|Z} \right\}\ ,$$

where the infimum is taken over all possible conditional distributions $P_{\overline{Z}|Z}$. ∘

Intuitively, the intrinsic conditional information $I(X;Y\downarrow Z)$ measures only the information between X and Y, which is possibly reduced by Z, but not the additional information brought in by giving Z. If for example X and Y are independent symmetric bits and $Z = X \oplus Y$, then we have $I(X;Y|Z) = 1$, but $I(X;Y\downarrow Z) = 0$.

It follows from the above that

$$S(X;Y||Z) \leq I(X;Y\downarrow Z) \ .$$

The fundamental problem of generally determining $S(X;Y||Z)$ for given P_{XYZ} has remained open, but there is some evidence that the intrinsic information is exactly the right quantity linking the secret-key rate with the joint distribution of X, Y, and Z.

Conjecture. $$S(X;Y||Z) = I(X;Y\downarrow Z) \ .$$

However, even the generally easier problem of completely characterizing the distributions P_{XYZ} for which $S(X;Y||Z) > 0$ holds, i.e., for which secret-key agreement is possible in principle, has not been fully answered yet (see Figure 26).

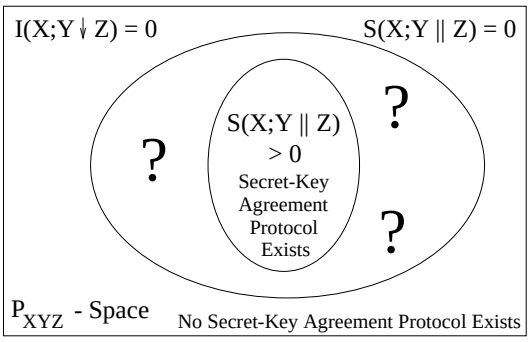

Fig. 26. Characterizing when Secret-Key Agreement Is Possible

6.2 Secret-Key Agreement Secure Against ACTIVE Adversaries

In all the previous models, we have assumed that the adversary is only a passive wire-tapper or equivalently, that the public channel connecting Alice and Bob is authentic. In many cases, secret-key agreement is even possible when dropping this condition, i.e., when the adversary is able to modify or introduce messages without being detected. See [17],[23],[35] for a discussion and analysis of this model.

Note first that a protocol secure against active opponents cannot be guaranteed to work in every situation because Eve, who is assumed to have full control

over the public channel, can block the channel permanently, preventing any communication between the legitimate partners. Hence the best that can be achieved by such a protocol is that Alice and Bob detect an adversary's active attacks and reject the outcome of the protocol unless secret-key agreement is successful (see Figure 27). More precisely, it is required that if Eve chooses to remain passive,

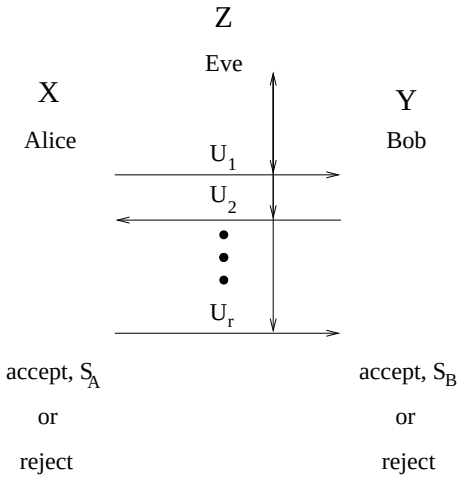

Fig. 27. Unconditional Security Against Active Opponents

then secret-key agreement is successful (as in the passive-adversary model). On the other hand, if Eve is active, then with overwhelming probability either Alice and Bob both reject the outcome of the protocol, or secret-key agreement is successful despite Eve's attacks. (Note that it is not requested that both Alice and Bob accept the outcome in the latter case. Such a perfect synchronization of the acceptance decisions cannot be achieved in the presence of an active adversary, who can always block the final message that makes the second party accept.)

Clearly, secret-key agreement can only be possible in the active-adversary scenario if Alice and Bob have some initial advantage over Eve in terms of the random variables X, Y, and Z. More precisely, this advantage must be such that Eve is not able to perfectly simulate Alice towards Bob and vice versa. In terms of the random variables, this is the condition that she cannot generate, using her random variable Z, a random variable $\overline{X}$ with the property that given only Y, $\overline{X}$ cannot be distinguished from X, and vice versa. Formally, this means that there do not exist conditional distributions $P_{\overline{X}|Z}$ or $P_{\overline{Y}|Z}$ such that either

$$P_{\overline{X}Y} = P_{XY}$$

or $P_{X\overline{Y}} = P_{XY}$ holds, respectively. If one of these distributions existed, secret-key agreement would be impossible because Bob could not tell Alice and Eve apart (or vice versa).

A surprising result however is that if secret-key agreement is possible also in the presence of an active adversary, then asymptotically the same key-generation rate as in the passive-adversary case can be achieved.

Finally, also privacy amplification can be executed in the case where the adversary is active. However, the restrictions on the opponent's knowledge about the partially secret key must be stronger [23],[35]. The idea is to use the string S twice, first as a key for unconditionally authenticating a message containing the description of a randomly chosen hash function, and as the argument for this function.

7 Concluding Remarks

We have described several techniques and results in the context of unconditional security in cryptography. The mentioned possibility and impossibility results can give a rough picture in what settings such provable confidentiality can be achieved. It is an important point in this context that despite Shannon's well-known pessimistic result, unconditional security is not necessarily impractical. A number of fundamental questions in this field are open today. In particular, the ultimate goal is the realization of a system that is practical and provably unconditionally secure simultaneously.

Acknowledgments

It is a pleasure to thank Ueli Maurer for initiating the author's interest in this exciting subject, and for many stimulating discussions. We also thank Ivan Damgård for the invitation to the Cryptology and Data Security Summer School.

References

1. R. Ahlswede and I. Csiszár, "Common randomness in information theory and cryptography - Part I: secret sharing," *IEEE Transactions on Information Theory*, vol. 39, no. 4, pp. 1121–1132, 1993.
2. C. H. Bennett, F. Bessette, G. Brassard, L. Salvail, and J. Smolin, "Experimental quantum cryptography," *Journal of Cryptology*, vol. 5, no. 1, pp. 3–28, Springer-Verlag, 1992.
3. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, "Generalized privacy amplification," *IEEE Transactions on Information Theory*, vol. 41, no. 6, pp. 1915–1923, 1995.
4. C. H. Bennett, G. Brassard, and J.-M. Robert, "Privacy amplification by public discussion," *SIAM Journal on Computing*, vol. 17, pp. 210–229, 1988.
5. R. E. Blahut, *Principles and practice of information theory*, Addison-Wesley Publishing Company, 1988.
6. G. Brassard and L. Salvail, "Secret-key reconciliation by public discussion," *Advances in Cryptology - EUROCRYPT '93*, Lecture Notes in Computer Science, vol. 765, pp. 410–423, Springer-Verlag, 1994.

7. W. Brunner, C. Cachin, U. M. Maurer, and C. Vonäsch, *Demonstration system of secret-key agreement by public discussion*, ETH Zürich, 1996. http://www.inf.ethz.ch/department/TI/um/keydemo/
8. C. Cachin, *Entropy measures and unconditional security in cryptography*, Ph. D. Thesis, ETH Zürich, Hartung-Gorre Verlag, Konstanz, 1997.
9. C. Cachin and U. M. Maurer, "Unconditional security against memory-bounded adversaries," *Advances in Cryptology - CRYPTO '97*, Lecture Notes in Computer Science, vol. 1294, pp. 292–306, Springer-Verlag, 1997.
10. T. M. Cover and J. A. Thomas, *Elements of information theory*, Wiley Series in Telecommunications, 1992.
11. R. Cramer and V. Shoup, "A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack," *Advances in Cryptology - CRYPTO '98*, Lecture Notes in Computer Science, vol. 1462, pp. 13–25, Springer-Verlag, 1998.
12. I. Csiszár and J. Körner, "Broadcast channels with confidential messages," *IEEE Transactions on Information Theory*, vol. IT-24, pp. 339–348, 1978.
13. W. Diffie and M. E. Hellman, "New directions in cryptography," *IEEE Transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, 1976.
14. W. Feller, *An introduction to probability theory and its applications*, 3rd edition, vol. 1, Wiley International, 1968.
15. M. J. Fischer and R. N. Wright, "Bounds on secret key exchange using a random deal of cards," *Journal of Cryptology*, vol. 9, no. 2, pp. 71–99, Springer-Verlag, 1996.
16. J. L. Massey, "A simplified treatment of Wyner's wire-tap channel," *Proceedings of the 21st Annual Allerton Conference of Communication, Control, and Computing*, Monticello, IL, pp. 268–276, 1983.
17. U. M. Maurer, "Information-theoretically secure secret-key agreement by NOT authenticated public discussion," *Advances in Cryptology - EUROCRYPT '97*, Lecture Notes in Computer Science, vol. 1233, pp. 209–225, Springer-Verlag, 1997.
18. U. M. Maurer, "The role of information theory in cryptography," *Codes and Ciphers*: 4th IMA Conference on Cryptography and Coding, Cirencester, UK, Dec. 1993, pp. 49-71, Southend-on-Sea, 1995. The Institute of Mathematics and its Applications.
19. U. M. Maurer, "The strong secret key rate of discrete random triples," *Communication and Cryptography – Two Sides of One Tapestry*, Kluwer Academic Publishers, pp. 271–285, 1994.
20. U. M. Maurer, "Protocols for secret key agreement based on common information," *Advances in Cryptology - CRYPTO '92*, Lecture Notes in Computer Science, vol. 740, pp. 461–470, Springer-Verlag, 1993.
21. U. M. Maurer, "Secret key agreement by public discussion from common information," *IEEE Transactions on Information Theory*, vol. 39, no. 3, pp. 733–742, 1993.
22. U. M. Maurer, "Conditionally-perfect secrecy and a provably-secure randomized cipher," *Journal of Cryptology*, vol. 5, pp. 53–66, Springer-Verlag, 1992.
23. U. M. Maurer and S. Wolf, "Privacy amplification secure against active adversaries," *Advances in Cryptology - CRYPTO '97*, Lecture Notes in Computer Science, vol. 1294, pp. 307–321, Springer-Verlag, 1997.
24. U. M. Maurer and S. Wolf, "The intrinsic conditional mutual information and perfect secrecy," *Proc. of the 1997 IEEE Symp. on Information Theory*, Ulm, Germany, 1997 (abstract). To appear in *IEEE Transactions on Information Theory*.

25. U. M. Maurer and S. Wolf, "Towards characterizing when information-theoretic secret key agreement is possible," *Advances in Cryptology - ASIACRYPT '96*, Lecture Notes in Computer Science, vol. 1163, pp. 196–209, Springer-Verlag, 1996.
26. M. Naor and A. Shamir, "Visual cryptography," *Advances in Cryptology - CRYPTO '94*, Lecture Notes in Computer Science, vol. 950, pp. 1–12, Springer-Verlag, 1995.
27. A. Rényi, *A diary on information theory*, Akadémiai Kiadó, Budapest, 1978.
28. R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
29. C. E. Shannon, "Communication theory of secrecy systems," *Bell System Technical Journal*, vol. 28, pp. 656–715, 1949.
30. C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423, 623–656, 1948.
31. P. W. Shor, "Algorithms for quantum computation: discrete log and factoring," *Proceedings of the 35th IEEE Symposium on the Foundations of Computer Science (FOCS '94)*, pp. 124–134, 1994.
32. V. Shoup, "Lower bounds for discrete logarithms and related problems," *Advances in Cryptology - EUROCRYPT '97*, Lecture Notes in Computer Science, vol. 1233, pp. 256–266, Springer-Verlag, 1997.
33. D. R. Stinson, "Universal hashing and authentication codes," *Advances in Cryptology - CRYPTO '91*, Lecture Notes in Computer Science, vol. 576, pp. 74–85, Springer-Verlag, 1992.
34. G. S. Vernam, "Cipher printing telegraph systems for secret wire and radio telegraphic communications," *Journal of the American Institute for Electrical Engineers*, vol. 55, pp. 109–115, 1926.
35. S. Wolf, "Strong security against active attacks in information-theoretic secret-key agreement," to appear in *Advances in Cryptology - ASIACRYPT '98*, Lecture Notes in Computer Science, Springer-Verlag, 1998.
36. A. D. Wyner, "The wire-tap channel," *Bell System Technical Journal*, vol. 54, no. 8, pp. 1355–1387, 1975.
37. R. W. Yeung, "A new outlook on Shannon's information measures," *IEEE Transactions on Information Theory*, vol. 37, no. 3, pp. 466–474, 1991.