1

Unbreakable Keys from Random Noise

Ueli Maurer, Renato Renner, and Stefan Wolf

[1] Department of Computer Science, ETH Zürich. maurer@inf.ethz.ch.
[2] Centre for Mathematical Sciences, University of Cambridge. r.renner@damtp.cam.ac.uk.
[3] Department of Computer Science, ETH Zürich. wolf@inf.ethz.ch.

Summary. Virtually all presently-used cryptosystems can theoretically be broken by an exhaustive key-search, and they might even be broken in practice due to novel algorithms or progress in computer engineering. In contrast, by exploiting the fact that certain communication channels are inherently noisy, one can achieve encryption provably-secure against adversaries with unbounded computing power, in arguably practical settings. This paper discusses secret key-agreement by public discussion from correlated information in a new definitional framework for information-theoretic reductions.

1.1 Information-Theoretic Cryptographic Security

1.1.1 Motivation

The security of essentially all presently-used cryptosystems is not proven. It is based on at least two assumptions. The first assumption is that the adversary's computational resources, specified within some model of computation, are bounded. This type of assumption can be problematic because it may not even be clear what the right model of computation is, as demonstrated by the recent proposal of a new computational model, a quantum computer, which is believed to be strictly more powerful than classical computers.

The second assumption is that the computational problem of breaking the cryptosystem is computationally infeasible, given the assumed computational resources. Such an assumption could potentially be proven, but the state of the art in complexity theory does not seem to be even close to proving any meaningful lower bound on the hardness of a computational problem. Important computational problems on which the security of cryptographic schemes is based are integer factorisation (e.g., RSA [24]) and computing discrete logarithms in certain finite cyclic groups (e.g., Diffie-Hellman [8]).

A cryptosystem for which the security could be rigorously proven based only on an assumption of the first type would be called *computationally secure*, while a cryptosystem secure under neither assumption, i.e., even against an adversary with unbounded computing power, is called *unconditionally secure* or *information-theoretically secure*. Such a system is even unbreakable by an exhaustive search over the key space.[4] As mentioned, no cryptosystem has been proven to be computationally secure (except of course those that are also information-theoretically secure.)

Many researchers have proposed cryptosystems that are unconditionally secure, with varying degrees of practicality. The most famous (but quite impractical) example is the one-time pad discussed later. There are two types of results on information-theoretic security: impossibility results and constructive results. In this paper we first discuss the impossibility of perfectly-secure message transmission and then focus on the key-agreement problem to show that information-theoretically secure key-agreement is possible in arguably practical settings if noisy information is exploited. We make use of information theory and refer to [6] for an introduction to the subject.

1.1.2 Information-Theoretic Security: Perfect Secrecy and Shannon's Theorem

Let us start with the classical scenario of a symmetric cryptosystem with message M, key K, and ciphertext C (see Fig. 1.1). The following security definition appears to be the strongest possible for such a cryptosystem.

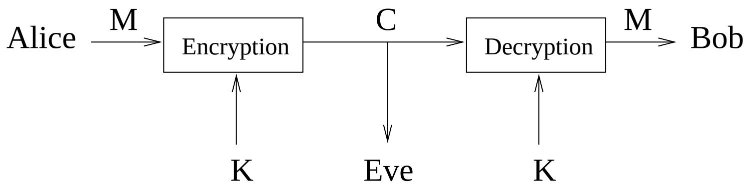

Fig. 1.1. A Symmetric Cryptosystem

Definition 1.1. **[25]** A cipher is called *perfectly secret* if the ciphertext reveals no information about the message, i.e., if $\mathbf{I}(M;C) = 0$ 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.

An example of a perfectly secret cipher is the *one-time pad* that was already proposed by Vernam in 1926 [26]. Here, the message is a string $M = [m_1, m_2, \ldots, m_N]$ of length N, and the key is a uniformly distributed

[4] It should be mentioned that also unconditionally security is based on an assumption, namely that our probabilistic model of Nature is (at least partially) correct.

N-bit string $K = [k_1, k_2, \dots, k_N]$ which is independent of M. The ciphertext C is computed from M and K by

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

A simple proof that the one-time pad is perfectly secret is obtained by using an entropy diagram (see Fig. 1.2) for the three random variables M, K, and C, as proposed in [28]. Any two of these random variables determine the third, hence $\mathsf{H}(C|MK) = 0$, $\mathsf{H}(M|CK) = 0$, and $\mathsf{H}(K|MC) = 0$. The key K is independent of the message M, and hence the entire entropy of K, namely $\mathsf{H}(K) = N$ must be concentrated in the field $\mathbf{I}(K;C|M)$, i.e., $\mathbf{I}(K;C|M) = N$. Since $\mathsf{H}(C) \le N$ and the field $\mathbf{I}(K;C|M)$ already contributes that much, we must have $\mathbf{I}(M;C) = 0$ (since $\mathbf{I}(M;C)$ is non-negative).

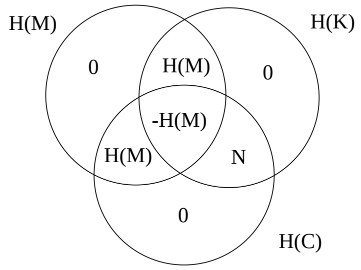

Fig. 1.2. Entropy diagram for one-time pad encryption.

Unfortunately, the price one has to pay 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

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. Maurer [16] proved the stronger statement that the same bound even holds in the more relevant setting where Alice and Bob can interact by (not secret) two-way communication.

Theorem 1.1. [25] *For every perfectly secret cryptosystem (with unique decodability), we have* $\mathsf{H}(K) \ge \mathsf{H}(M)$.

For a proof of Shannon's theorem, note first that unique decodability means $\mathsf{H}(M|CK) = 0$. The entropy diagram of the involved quantities is shown in Fig. 1.3. We have $b \ge a$ because $\mathbf{I}(C;K) \ge 0$, and

$$\mathsf{H}(K) \ge b - a + c \ge a - a + c = \mathsf{H}(M) \ ,$$

which concludes the proof.

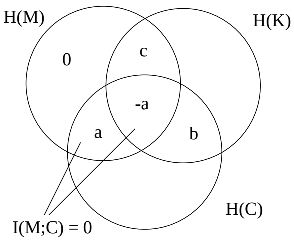

Fig. 1.3. The proof of Shannon's theorem.

1.1.3 Optimistic Results by Limiting the Adversary's Information

Unfortunately, Shannon's and Maurer's above-mentioned results imply 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 *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.

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. If insecure channels are available, this also implies 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 [27],[7],[1],[16], on the fact that the adversary's memory is limited [13] or on the uncertainty principle of quantum mechanics [2, 10].

1.1.4 The Power of Feedback

Wyner [27] showed that in the special case where a (noisy) channel $P_{Y|X}$ from Alice to Bob is available, and where the adversary receives a degraded version Z of Y (through a channel $P_{Z|Y}$, independent of the main channel), secret-key agreement is possible in all non-trivial cases. This setting was generalized by Csiszár and Körner [7] who studied the model of a so-called *noisy broadcast channel* characterized by a probability distribution $P_{YZ|X}$, where Alice's input is X, whereas Bob and Eve receive Y and Z, respectively. They introduced a quantity, called the *secrecy capacity*, measuring Alice and Bob's ability to generate a virtually secret key (asymptotically, per channel use). Their results imply that in case of independent binary symmetric channels, key agreement is possible if and only if Bob's channel has a smaller error probability than Eve's.

However, the following example, given in [16], illustrates the somewhat surprising fact that by using an insecure feedback channel, secret-key agreement is possible in the above setting. We start with the situation on the right-hand side of Fig. 1.4. Here, 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. Surprisingly, the situation is now entirely different although the additional channel is accessible to 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 Fig. 1.4).

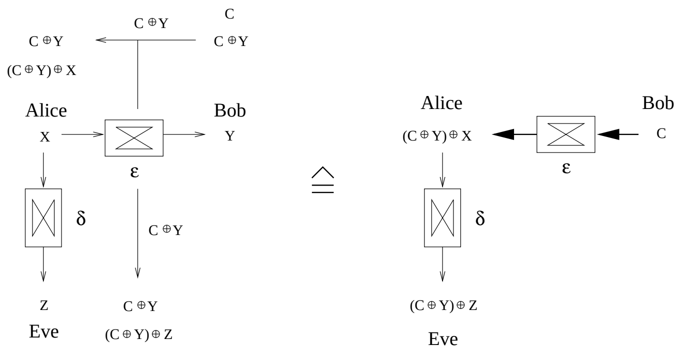

Fig. 1.4. Inverting the main channel.

The second crucial observation is that this is exactly the binary-symmetric setting of Wyner's wire-tap channel [27], allowing secret-key agreement at a positive 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.

Maurer [16] proposed the following interactive model of secret-key agreement by public discussion from common information (see Fig. 1.5). 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.

This model is more general and more natural than noisy channel models since the assumption that the parties have access to correlated information appears to be realistic in many contexts. In the rest of the paper, this model will be the basis of our considerations.

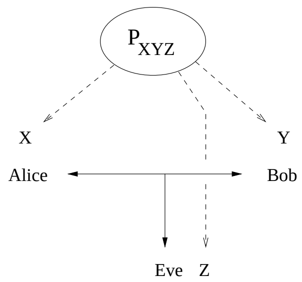

Fig. 1.5. Secret-key agreement by public discussion from common information.

1.2 Smooth Rényi Entropies

Let X and Y be random variables with ranges $\mathcal{X}$ and $\mathcal{Y}$, distributed according to P_{XY}. For any event $\mathcal{E}$ with conditional probabilities $P_{\mathcal{E}|XY}(x,y)$, let[5]

$$\begin{aligned}\mathsf{H}_{\min}(\mathcal{E}X|Y) &:= -\log \max_{(x,y)\in\mathcal{X}\times\mathcal{Y}} P_{\mathcal{E}X|Y}(x,y)\\ \mathsf{H}_{\max}(\mathcal{E}X|Y) &:= \log \max_{y\in\mathcal{Y}} |\{x\in\mathcal{X} : P_{\mathcal{E}XY}(x,y)>0\}| ,\end{aligned}$$

where $P_{\mathcal{E}X|Y}(x,y) := \frac{P_{\mathcal{E}|XY}(x,y)P_{XY}(x,y)}{P_Y(y)}$ (with the convention that $\frac{0}{0}=0$).

The *smooth min-entropy* and the *smooth max-entropy* are then defined by [22, 23] (see also [18] for a generalization of these entropy measures to quantum information theory)

$$\begin{aligned}\mathsf{H}^{\varepsilon}_{\min}(X|Y) &:= \max_{\mathcal{E}:\,\Pr[\mathcal{E}]\geq 1-\varepsilon} \mathsf{H}_{\min}(\mathcal{E}X|Y)\\ \mathsf{H}^{\varepsilon}_{\max}(X|Y) &:= \min_{\mathcal{E}:\,\Pr[\mathcal{E}]\geq 1-\varepsilon} \mathsf{H}_{\max}(\mathcal{E}X|Y) ,\end{aligned}$$

where $\Pr[\mathcal{E}] := \sum_{x,y} P_{\mathcal{E}|XY}(x,y)P_{XY}(x,y)$. Smooth min- and max-entropies can be seen as generalizations of the Shannon entropy in the following sense.

Lemma 1.1. *Let P_{XY} be fixed and let, for any $n\in\mathbb{N}$, $(X_1,Y_1),\ldots(X_n,Y_n)$ be a sequence of random variables distributed according to $(P_{XY})^{\times n}$. Then*

[5] All logarithms are to the base 2.

$$\lim_{\varepsilon\to 0}\lim_{n\to\infty}\frac{1}{n}\mathsf{H}^{\varepsilon}_{\min}(X_1\cdots X_n|Y_1\cdots Y_n)=\mathsf{H}(X|Y)$$
$$\lim_{\varepsilon\to 0}\lim_{n\to\infty}\frac{1}{n}\mathsf{H}^{\varepsilon}_{\max}(X_1\cdots X_n|Y_1\cdots Y_n)=\mathsf{H}(X|Y)$$

where $\mathsf{H}(X|Y)=\mathsf{H}(P_{XY})-\mathsf{H}(P_Y)$ *is the Shannon entropy of* X *conditioned on* Y.

Smooth min- and max-entropies satisfy some basic rules that are very similar to those known from Shannon theory. The following lemma is an analogue of the strong subadditivity, $\mathsf{H}(X|YZ)\leq\mathsf{H}(X|Z)$.

Lemma 1.2. *Let* X, Y, *and* Z *be random variables and let* $\varepsilon\geq 0$. *Then*

$$\mathsf{H}^{\varepsilon}_{\min}(X|YZ)\leq\mathsf{H}^{\varepsilon}_{\min}(X|Z)$$
$$\mathsf{H}^{\varepsilon}_{\max}(X|YZ)\leq\mathsf{H}^{\varepsilon}_{\max}(X|Z)\ .$$

The following is a generalization of the chain rule $\mathsf{H}(XY|Z)=\mathsf{H}(X|YZ)+\mathsf{H}(Y|Z)$.

Lemma 1.3. *Let* X, Y, *and* Z *be random variables and let* $\varepsilon,\varepsilon',\varepsilon''\geq 0$. *Then*

$$\mathsf{H}^{\varepsilon+\varepsilon'}_{\min}(XY|Z)\geq\mathsf{H}^{\varepsilon}_{\min}(X|YZ)+\mathsf{H}^{\varepsilon'}_{\min}(Y|Z)$$
$$\mathsf{H}^{\varepsilon'}_{\min}(XY|Z)<\mathsf{H}^{\varepsilon+\varepsilon'+\varepsilon''}_{\min}(X|YZ)+\mathsf{H}^{\varepsilon''}_{\max}(Y|Z)+\log(1/\varepsilon)$$
$$\mathsf{H}^{\varepsilon+\varepsilon'}_{\max}(XY|Z)\leq\mathsf{H}^{\varepsilon}_{\max}(X|YZ)+\mathsf{H}^{\varepsilon'}_{\max}(Y|Z)$$
$$\mathsf{H}^{\varepsilon'}_{\max}(XY|Z)>\mathsf{H}^{\varepsilon+\varepsilon'+\varepsilon''}_{\max}(X|YZ)+\mathsf{H}^{\varepsilon''}_{\min}(Y|Z)-\log(1/\varepsilon)\ .$$

1.3 Information-Theoretic Reductions

1.3.1 Resources

Consider a set of *players* $\mathcal{P}$ which have access to a certain set of *resources* such as a public communication channel or a source of common randomness. Information theory (and cryptography) is concerned with the question of whether (and how) the players can use a given resource $\mathcal{R}$ (or a set of resources) in order to build the functionality of a new resource $\mathcal{S}$, e.g., a secret communication channel.

On an abstract level, a resource is simply a *(random) system* [15] which can be accessed by each of the players in $\mathcal{P}$. For simplicity, we focus on the restricted class of resources which only take *one single* input U_P from each of the players $P\in\mathcal{P}$ and output one single value V_P to each of them. A resource for an n-player set $\mathcal{P}=\{1,\ldots,n\}$ is then fully specified by a conditional probability distribution $P_{V_1\cdots V_n|U_1\cdots U_n}$. (If the resource does not give an output to player P, V_P is defined to be a constant which we denote by $\perp$.)

For the following, we restrict our considerations to situations with three players and call them *Alice*, *Bob*, and *Eve*. Accordingly, we use the letters A, B, and E to denote the players' inputs and X, Y, and Z for the corresponding outputs. Typically, Alice and Bob will take the role of *honest parties* (which means that they follow a specified protocol) whereas Eve is an *adversary* who might behave arbitrarily. We will use the convention that, if no adversary is present then $E :=\perp$.

Let us have a look at some examples of resources, specified by a conditional probability distribution $P_{XYZ|ABE}$.

Example 1.1. An *authentic public (ℓ-bit) channel taking input from Alice*, denoted $\mathsf{Auth}_\ell^{A\to B}$, is a resource $P_{XYZ|ABE}$ whose outputs X, Y, Z simply take the value of the input A. Formally, for any $a \in \{0,1\}^\ell$,

$$P_{XYZ|ABE}(x,y,z,a,b,e) = \begin{cases} 1 & \text{if } x = y = z = a \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, one can define an *authentic public two-way channel* $\mathsf{Auth}_\ell^{A\leftrightarrow B}$.

Example 1.2. An *authentic secret (ℓ-bit) channel from Alice to Bob*, denoted $\mathsf{Sec}_\ell^{A\to B}$, is a resource $P_{XYZ|ABE}$ whose output Y takes the value of the input A whereas the output Z is constant. Formally, for any $a \in \{0,1\}^\ell$,

$$P_{XYZ|ABE}(x,y,z,a,b,e) = \begin{cases} 1 & \text{if } y = a \text{ and } x = z =\perp \\ 0 & \text{otherwise.} \end{cases}$$

Example 1.3. A *source of correlated randomness (with distribution P_{XYZ})*, denoted $\mathsf{Source}(P_{XYZ})$, is a resource $P_{XYZ|ABE}$ whose outputs X, Y, and Z are jointly distributed according to P_{XYZ} (independently of the inputs). Formally,

$$P_{XYZ|ABE}(x,y,z,a,b,e) := P_{XYZ}(x,y,z) \ .$$

Example 1.4. A *common secret key (of length ℓ)*, denoted SK_ℓ, is a source of correlated randomness $\mathsf{Source}(P_{XYZ})$ where $X = Y$ are uniformly distributed over $\{0,1\}^\ell$ and Z is a constant. Formally,

$$P_{XYZ|ABE}(x,y,z,a,b,e) = \begin{cases} 2^{-\ell} & \text{if } x = y \in \{0,1\}^\ell \text{ and } z =\perp \\ 0 & \text{otherwise.} \end{cases}$$

Example 1.5. A *unreliable common secret key (of length ℓ)*, denoted SK_ℓ^{AB}, is a resource $P_{XYZ|ABE}$ such that the following holds: If $E =\perp$ then the behavior of the resource is the same as for a common secret key SK_ℓ. If $E \neq\perp$ then $X = Y = Z =\perp$. In other words, whenever Alice and Bob get a key then this key is guaranteed to be secure (i.e., unknown to Eve). However, Eve might cause the resource to simply output $\perp$. Formally,

$$P_{XYZ|ABE}(x,y,z,a,b,e) = \begin{cases} 2^{-\ell} & \text{if } x = y \in \{0,1\}^\ell,\ z = \perp, \text{ and } e = \perp \\ 1 & \text{if } x = y = z = \perp \text{ and } e \neq \perp \\ 0 & \text{otherwise.} \end{cases}$$

Example 1.6. An *asymmetric secret key (of length ℓ) with security for Alice*, denoted SK_ℓ^A, is a resource $P_{XYZ|ABE}$ such that the following holds: If $E = \perp$ then the behavior of the resource is the same as for a common secret key SK_ℓ. If $E \neq \perp$ then $Y = E$ and $X = Z = \perp$. In other words, whenever Alice gets a key then it is also known to Bob and secret. However, Eve might cause the resource to give an arbitrary value (chosen by her) to Bob, but this will be detected by Alice. Formally,

$$P_{XYZ|ABE}(x,y,z,a,b,e) = \begin{cases} 2^{-\ell} & \text{if } x = y \in \{0,1\}^\ell,\ z = \perp, \text{ and } e = \perp \\ 1 & \text{if } y = e \text{ and } x = z = \perp \\ 0 & \text{otherwise.} \end{cases}$$

Note that a unreliable common secret key (Example 1.5) models what quantum key distribution (QKD) achieves (using an insecure quantum channel and an authentic classical channel). As long as the adversary is passive, Alice and Bob will get a secret key. On the other hand, Eve might intercept the quantum communication between Alice and Bob, but any severe attack would be detected (with high probability), in which case no key is generated (see [18] for more details on security definitions in quantum cryptography).

Distance Between Resources

In order to compare resources, we will need a notion of distance between two resources. For cryptographic applications, *the* natural distance measure is the *distinguishing advantage*, which is the basis of the following definition.

Definition 1.2. *Two resources $\mathcal{R}$ and $\mathcal{S}$ are said to be ε-close, denoted $\mathcal{R} \overset{\varepsilon}{\approx} \mathcal{S}$, if*

$$\Pr[\mathcal{D}(\mathcal{R}) = 0] - \Pr[\mathcal{D}(\mathcal{S}) = 0] \leq \varepsilon \ .$$

for any system $\mathcal{D}$ which interacts with $\mathcal{R}$ (or $\mathcal{S}$) and gives a binary output $\mathcal{D}(\mathcal{R})$ ($\mathcal{D}(\mathcal{S})$).

This distance measure has an intuitive interpretation. If two resources $\mathcal{R}$ and $\mathcal{S}$ are ε-close then they can be considered equal except with probability (at most) ε.

Using Resources in Parallel

Given two resources $\mathcal{R}$ and $\mathcal{S}$, we denote by $\mathcal{R} \times \mathcal{S}$ the resource which provides both the functionality of $\mathcal{R}$ and $\mathcal{S}$ in parallel. For example, in the three-party case, $\mathsf{Auth}^{A \to B} \times \mathsf{Source}(P_{XYZ})$ describes a situation where Alice and Bob have access to an authentic public channel (from Alice to Bob) *and* a source of correlated randomness (with distribution P_{XYZ}).

1.3.2 Programs and Protocols

Given a resource $\mathcal{R}$, a player P can interact with $\mathcal{R}$ by choosing inputs U_P and processing its outputs V_P. Technically, the way a player P uses $\mathcal{R}$ can be described as a random system π_P which starts with some input X'_P, then interacts with $\mathcal{R}$, and eventually generates an output Y'_P. In the following, we call π_P a *program (for player P)*. Applying a program π_P to a resource $\mathcal{R}$ naturally defines a new resource, which we denote by $\pi_P(\mathcal{R})$. More generally, if players $P_1, \ldots, P_k$ apply programs $\pi_{P_1}, \ldots, \pi_{P_k}$ to $\mathcal{R}$, we denote the resulting resource by $\pi_{P_1} \circ \cdots \circ \pi_{P_k}(\mathcal{R})$. Note that, because all programs act on different inputs/outputs of $\mathcal{R}$, the order in which the programs are written is irrelevant.

1.3.3 Realizing Resources

For the following definition, we again restrict to the three-party case with two honest players (Alice and Bob) and a malicious player (Eve). A pair of programs (π_A, π_B) for Alice and Bob is called a *protocol*. Moreover, we denote by $\perp_E$ the program for Eve which inputs $\perp$ to the resource and outputs $\perp$.

Definition 1.3. *Let $\mathcal{R}$ and $\mathcal{S}$ be resources and let $\pi = (\pi_A, \pi_B)$ be a protocol. We say that π ε-*realizes* $\mathcal{S}$ *from* $\mathcal{R}$, denoted*

$$\mathcal{R} \xrightarrow{\pi}_{\varepsilon} \mathcal{S}$$

if the following holds:

- $\pi_A \circ \pi_B \circ \perp_E (\mathcal{R}) \stackrel{\varepsilon}{\approx} \perp_E (\mathcal{S})$;
- *there exists a program τ_E for Eve such that $\pi_A \circ \pi_B(\mathcal{R}) \stackrel{\varepsilon}{\approx} \tau_E(\mathcal{S})$.*

Note that the definition imposes two conditions. The first corresponds to a situation where the adversary is passive. In this case, Alice and Bob apply their programs π_A and π_B to the resource $\mathcal{R}$, whereas Eve does nothing. The resulting resource should then be a good approximation of $\mathcal{S}$ where, again, Eve does nothing.

The second condition of the definition corresponds to a situation where the adversary is active. In this case, it should be guaranteed that Eve could run some simulator[6] τ_E on $\mathcal{S}$ which would give her the same information as she would get when accessing $\mathcal{R}$.

The relation $\longrightarrow$ is transitive in the following sense.

Lemma 1.4. *Given two protocols Π and Π' such that $\mathcal{R} \xrightarrow{\Pi}_{\varepsilon} \mathcal{S}$ and $\mathcal{S} \xrightarrow{\Pi'}_{\varepsilon'} \mathcal{T}$ then the sequential concatenation $\Pi' \circ \Pi$ satisfies $\mathcal{R} \xrightarrow{\Pi' \circ \Pi}_{\varepsilon + \varepsilon'} \mathcal{T}$.*

In cryptography, this transitivity is also called *composability* (most notably in the frameworks for computational cryptography proposed in [17, 5]).

[6] The concept of simulators in cryptographic security definitions has been introduced in [11].

1.3.4 Examples Protocols

Example 1.7. A *one-time pad*, denoted OTP, is a protocol which (perfectly) realizes a secret channel using an authentic channel and a secret key (see Section 1.1.2), i.e.,

$$\mathsf{Auth}_\ell^{A\to B} \times \mathsf{SK}_\ell \xrightarrow{\mathtt{OTP}}_0 \mathsf{Sec}_\ell^{A\to B} \ . \tag{1.1}$$

The protocol $\mathtt{OTP} = (\pi_A, \pi_B)$ is defined as follows: π_A takes as input A', computes the XOR between A' and the secret key, and then gives the result as input to the public communication channel. π_B computes the XOR between the output of the public channel and the key and outputs the result.

To prove (1.1), we need to verify that the two criteria of Definition 1.3 are satisfied. For the first, it suffices to check that the resources on both sides of the identity

$$\pi_A \circ \pi_B \circ \perp_E (\mathsf{Auth}_\ell^{A\to B} \times \mathsf{SK}_\ell) = \perp_E (\mathsf{Sec}_\ell^{A\to B})$$

are equal. In fact, both of them take an ℓ-bit input from Alice and output the value of this input to Bob, whereas Eve gets a constant.

The second criterion (for $\varepsilon = 0$) reads

$$\pi_A \circ \pi_B(\mathsf{Auth}_\ell^{A\to B} \times \mathsf{SK}_\ell) = \tau_E(\mathsf{Sec}_\ell^{A\to B})$$

where τ_E is some appropriately chosen program. Note that Eve's output of the resource defined by the left hand side of this equality is simply the XOR between the message and the key. This value is uniformly distributed and independent of the message. It thus suffices to define τ_E as a program which simply outputs some uniformly distributed ℓ-bit string.

Example 1.8. The following protocol Π uses a public authentic channel to transform an asymmetric secret key with security for Alice into a unreliable common secret key, i.e.,

$$\mathsf{SK}_\ell^{A} \times \mathsf{Auth}_1^{A\to B} \xrightarrow{\Pi}_0 \mathsf{SK}_\ell^{AB} \ .$$

Let X and Y be the outputs of SK_ℓ^{A} on Alice and Bob's side, respectively. The protocol $\Pi = (\pi_A, \pi_B)$ is then defined as follows: Alice's program π_A uses the public channel to announce whether $X = \perp$ and outputs X. Bob's program π_B outputs $\perp$ if $X = \perp$ and Y otherwise.

Example 1.9. A *hashing* or *privacy amplification protocol* HA is a protocol which uses a source of correlated randomness and an authentic public channel to realize a secret key (see Section 1.4.5 below). Formally,[7]

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Auth}_\infty^{A\to B} \xrightarrow{\mathtt{HA}}_\varepsilon \mathsf{SK}_\ell \ ,$$

[7] $\mathsf{Auth}_\infty^{A\to B}$ denotes an authentic public channel for arbitrarily long messages.

for any distribution P_{XYZ} with $X = Y$ and $\mathsf{H}_{\min}^{\varepsilon'}(X|Z) \geq \ell + 2\log(1/(\varepsilon - \varepsilon'))$. The protocol $\mathtt{HA} = (\pi_A, \pi_B)$ is defined as follows: π_A chooses at random a function f from a two-universal (see Definition 1.9 below) set of functions from X to a string of size ℓ, sends a description of f over the channel, and outputs $f(X)$. π_B outputs $f(Y)$.

The proof of the above statement follows immediately from the security of privacy amplification [4, 12, 3] (see Corollary 1.3). This example will be further elaborated on in Section 1.4.5.

Example 1.10. An *information reconciliation protocol* IR uses a public channel to transform a source of correlated randomness P_{XYZ} into another source of correlated randomness $P_{X'Y'Z'}$ such that $X' = Y'$ in such a way that the decrease of Eve's uncertainty on Alice's data is minimal (see Section 1.4.4 below).[8] Formally,

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Auth}_{\infty}^{A \to B} \xrightarrow{\mathtt{IR}}_{\varepsilon} P_{X'Y'Z'} ,$$

where $X' = Y'$ and $\mathsf{H}_{\min}^{\varepsilon'+\varepsilon''}(X'|Z') > \mathsf{H}_{\min}^{\varepsilon'}(X|Z) - \mathsf{H}_{\max}(X|Y) - \log(1/\varepsilon\varepsilon'')$ (see Corollary 1.2). In order to achieve this, Alice sends some error correcting information C to Bob which, together with his knowledge Y, allows him to guess Alice's value X. A conceptually simple way to generate C is by two-universal hashing of X (in practice, one typically uses error correcting codes that have more structure in order to allow computationally efficient decoding on Bob's side).

More precisely, the protocol IR works a follows: Alice's program π_A chooses at random a function f from a two-universal set of hash functions which map X to a string of length roughly $\mathsf{H}_{\max}(X|Y) + \log(1/\varepsilon)$. A description of f as well as $C = f(X)$ is then sent to Bob using the public channel. Moreover, Alice outputs $X' := X$. Bob's program π_B, upon receiving C, outputs some string Y' which satisfies $P_{X|Y}(Y'|Y) > 0$ (i.e., has non-zero probability from his point of view) and $f(Y') = C$ (i.e., is compatible with C). For information reconciliation, see also Section 1.4.4.

1.3.5 Independent and Identically Distributed Resources

In information-theoretic cryptography, one often assumes that Alice and Bob can use many independent realizations of a given resource (e.g., many independently distributed pairs of correlated random values (X, Y)) in order to produce many independent realizations of another resource (e.g., a secret key bit). Under this so-called *i.i.d. assumption*, one can study asymptotic quantities such as *key rates*.

[8] A *secure sketch* as defined in [9] can be seen as a special case of an information reconciliation protocol where the *sketching* and the *recovery procedure* correspond to Alice and Bob's programs, respectively.

Definition 1.4. *An* asymptotic protocol Π *is a sequence of pairs* (Π_k, τ_k) *where, for any* $k \in \mathbb{N}$*,* Π_k *is a protocol and* $\tau_k \in \mathbb{N}$*. The* rate of Π *is defined by*

$$\mathrm{rate}(\Pi) := \lim_{k \to \infty} \frac{k}{\tau_k} .$$

Definition 1.5. *We say that an asymptotic protocol* $\{(\Pi_k, \tau_k)\}_{k \in \mathbb{N}}$ realizes $\mathcal{S}$ from $\mathcal{R}$*, denoted*

$$\mathcal{R} \overset{\Pi_k, \tau_k}{\Longrightarrow} \mathcal{S} ,$$

if there exists a zero-sequence $\{\varepsilon_k\}_{k \in \mathbb{N}}$ *(i.e.,* $\lim_{k \to \infty} \varepsilon_k = 0$*) such that, for any* $k \in \mathbb{N}$*,*

$$\mathcal{R}^{\times \tau_k} \xrightarrow{\Pi_k}_{\varepsilon_k} \mathcal{S}^{\times k} .$$

See below for examples of asymptotic protocols.

Definition 1.6. *Let* $\Pi = \{(\Pi_k, \tau_k)\}_{k \in \mathbb{N}}$ *and* $\Pi' = \{(\Pi'_k, \tau'_k)\}_{k \in \mathbb{N}}$ *be asymptotic protocols. The concatenation* $\bar{\Pi} := \Pi' \circ \Pi$ *is then defined by the protocol* $\{\bar{\Pi}_k, \bar{\tau}_k\}_{k \in \mathbb{N}}$ *where* $\bar{\Pi}_k := \Pi'_{\tau_k} \circ \Pi_k$ *and* $\bar{\tau}_k := \tau'_{\tau_k}$*.*

Definition 1.7. *Let* $\mathcal{R}$ *and* $\mathcal{S}$ *be resources. The* rate of $\mathcal{R} \Longrightarrow \mathcal{S}$ *is defined by*

$$\mathrm{rate}(\mathcal{R} \Longrightarrow \mathcal{S}) := \max_{\Pi} \mathrm{rate}(\Pi)$$

where the maximum ranges over all protocols Π *such that* $\mathcal{R} \overset{\Pi}{\Longrightarrow} \mathcal{S}$*.*

It is straight-forward to prove that composability also holds for this asymptotic definition:

Lemma 1.5. *Given two asymptotic protocols* Π *and* Π' *such that* $\mathcal{R} \overset{\Pi}{\Longrightarrow} \mathcal{S}$ *and* $\mathcal{S} \overset{\Pi'}{\Longrightarrow} \mathcal{T}$ *then the sequential concatenation* $\pi' \circ \pi$ *satisfies*

$$\mathcal{R} \overset{\Pi' \circ \Pi}{\Longrightarrow} \mathcal{T} .$$

Moreover $\mathrm{rate}(\mathcal{R} \Longrightarrow \mathcal{T}) \geq \mathrm{rate}(\mathcal{R} \Longrightarrow \mathcal{S}) \cdot \mathrm{rate}(\mathcal{S} \Longrightarrow \mathcal{T})$*.*

1.4 Turning Correlated Randomness into Keys

1.4.1 Generic One-Way Key Agreement

A generic way to generate a key from weakly correlated and only partially secure randomness is to employ an error correction protocol and then apply privacy amplification. This process only requires communication in one direction (e.g., from Alice to Bob). This simple protocol for one-way key agreement is also used in more complex key agreement protocols (including QKD protocols) as a last step.[9]

Theorem 1.2. *Let $a, b \in \mathbb{N}$ and $\varepsilon, \varepsilon', \varepsilon'' > 0$ be fixed and let $\ell := a - b - 3\log(2/\varepsilon)$. Then there exists a protocol* $\mathtt{KA} = \mathtt{KA}_{a,b}$ *(called* one-way key agreement protocol*) such that*

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Auth}_{\infty}^{A\to B} \xrightarrow{\mathtt{KA}}_{\bar{\varepsilon}} \mathsf{SK}_{\ell} \ , \tag{1.2}$$

for any P_{XYZ} such that $\mathsf{H}_{\min}^{\varepsilon'}(X|Z) \geq a$ and $\mathsf{H}_{\max}^{\varepsilon''}(X|Y) \leq b$, and $\bar{\varepsilon} \geq 2\varepsilon + \varepsilon' + \varepsilon''$.

In particular, for any distribution P_{XYZ} with

$$\mathsf{H}_{\min}^{\varepsilon'}(X|Z) - \mathsf{H}_{\max}^{\varepsilon''}(X|Y) \geq \ell + 3\log(2/\varepsilon) \ ,$$

there exists a protocol $\mathtt{KA}$ *satisfying* (1.2)*.*

Note that a very similar result also holds in a quantum world, where Eve's information is encoded into the state of a quantum system [18].

Proof. The protocol $\mathtt{KA}$ can be defined as the concatenation of the information reconciliation protocol $\mathtt{IR}$ and the hashing protocol $\mathtt{HA}$ as in Examples 1.10 and 1.9, respectively. The assertion then follows from the composition lemma (Lemma 1.4).

1.4.2 Independent Repetitions

Definition 1.8. *The* secret-key rate *of a tripartite probability distribution P_{XYZ} is defined by*

$$S(P_{XYZ}) := \mathrm{rate}(\mathsf{Source}(P_{XYZ}) \times \mathsf{Auth}_{\infty}^{A\leftrightarrow B} \stackrel{\mathtt{KA}}{\Longrightarrow} \mathsf{SK}_1) \ .$$

Similarly, the one-way secret-key rate *is*

$$S_{\rightarrow}(P_{XYZ}) := \mathrm{rate}(\mathsf{Source}(P_{XYZ}) \times \mathsf{Auth}_{\infty}^{A\to B} \stackrel{\mathtt{KA}}{\Longrightarrow} \mathsf{SK}_1) \ .$$

The following lemma gives some basic properties of the secret-key rate.

Lemma 1.6. *Let X, Y, and Z be random variables with joint distribution P_{XYZ}. Then*

1. *Local operations by Alice can only decrease the rate: $S(P_{XYZ}) \geq S(P_{X'YZ})$ for any X' such that $(Y, Z) \to X \to X'$ is a Markov chain.*

[9] A *fuzzy extractors* as defined in [9] can generally be seen as a one-way key agreement protocol where the *generation* and the *reproduction* procedure correspond to Alice and Bob's programs, respectively. The *helper string* generated by the generation procedure of a secure sketch corresponds to the message sent from Alice to Bob in a one-way key agreement protocol.

2. *Local operations by Bob can only decrease the rate:* $S(P_{XYZ}) \geq S(P_{XY'Z})$ *for any* Y' *such that* $(X,Z) \to Y \to Y'$ *is a Markov chain.*
3. *Local operations by Eve can only increase the rate:* $S(P_{XYZ}) \leq S(P_{XYZ'})$ *for any* Z' *such that* $(X,Y) \to Z \to Z'$ *is a Markov chain.*
4. *Giving information* U *to Eve can only decrease the rate by at most the entropy of* U*:* $S(P_{XYZ}) \leq S(P_{XYZU}) + \mathsf{H}(U|Z)$ *for any random variable* U*.*

The following is an immediate consequence of Theorem 1.2 and Lemma 1.1.

Corollary 1.1. *The one-way secret-key rate is lower bounded by*

$$S_{\to}(P_{XYZ}) \geq \mathsf{H}(X|Z) - \mathsf{H}(X|Y) \ . \tag{1.3}$$

The one-way key agreement protocol presented above is not optimal. In fact, there are distributions P_{XYZ} where $\mathsf{H}(X|Z) - \mathsf{H}(X|Y)$ equals zero but $S(P_{XYZ})$ is still positive, as the following example illustrates.

Example 1.11. Let X be a uniformly distributed random bit, let Y be the output of a binary symmetric channel with noise δ on input X, and let Z be the output of an erasure channel with erasure probability δ' on input X.

Let $\delta := h^{-1}(\frac{1}{2}) \approx 0.11$ and $\delta' := 1/2$. Then

$$\mathsf{H}(X|Z) - \mathsf{H}(X|Y) = 0$$

Consider now the random variable X' obtained by sending X through a binary symmetric channel with noise μ. Then, for any $\mu > 0$,

$$\mathsf{H}(X'|Z) - \mathsf{H}(X'|Y) > 0$$

and hence $S(P_{X'YZ}) > 0$. It thus follows from statement 1 of Lemma 1.6 that $S(P_{XYZ}) > 0$.

1.4.3 Advantage Distillation

There exist situations where the expression in (1.3) is negative, but key agreement is—somewhat surprisingly—nevertheless possible. An example is the "special satellite scenario" [16], where all involved parties have conditionally-independent noisy versions of a binary signal. In this case, key agreement has been shown possible in *all non-trivial scenarios*, i.e., even when the information the adversary obtains about this signal is arbitrarily larger than the legitimate parties'.

In this case, however, *two-way communication* is necessary. A concrete example of such an *advantage-distillation protocol* is the following: Alice and Bob repeatedly compare parities of bits publicly and continue the process only in case of equal parities. Intuitively speaking, this allows them to use the authentic channel for locating positions where an error is less likely, until they finally end up in a situation where they know more about each other's pieces of information than the adversary (although the latter also learns the public communication of the protocol).

1.4.4 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 leads to a lower bound on the amount of error-correction information C that must be exchanged. Namely, Bob must know S_A completely with overwhelming probability when given S_B and C. Hence, the amount of error-correction information is at least the uncertainty of S_A given S_B. On the other hand, the uncertainty of S_A from Eve's viewpoint can as well be reduced by $\mathsf{H}(C)$ in the worst case when Eve learns C (see Fig. 1.6).

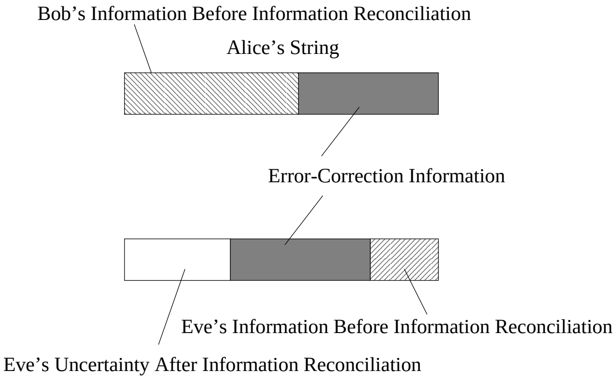

Fig. 1.6. The effect of information leaked during information reconciliation

Lemma 1.7 and Corollary 1.2 link information reconciliation to the conditional smooth max-entropy of Alice's information, given Bob's.

Definition 1.9. *A class $\mathcal{H}$ of functions h mapping a set $\mathcal{A}$ to a set $\mathcal{B}$ is called* two-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 two-universal class behaves like a completely random function with respect to* collisions. ◦

An example of a two-universal class of functions, mapping $\{0,1\}^n$ to $\{0,1\}^r$, of cardinality $2^{n \cdot r}$ are the *linear* functions.

Lemma 1.7. *Let X be a random variable on $\mathcal{X}$, let $\mathcal{E}$ be an event, and let F be chosen at random (independently of X) from a two-universal family of hash functions from $\mathcal{X}$ to $\mathcal{U}$. Then there exists a function d_F depending on F such that*

$$\Pr\big[\mathcal{E} \wedge (d_F(F(X)) \neq X)\big] \leq \frac{|\mathcal{X}_\mathcal{E}|}{|\mathcal{U}|} ,$$

where $\mathcal{X}_\mathcal{E} := \{x \in \mathcal{X} : P_{\mathcal{E}X}(x) > 0\}$.

Proof. For any hash function F, define $\mathcal{D}_F(u) := \mathcal{X}_\mathcal{E} \cap F^{-1}(u)$. Moreover, let d_F be any function from $\mathcal{U}$ to $\mathcal{X}$ such that $d_F(u) \in \mathcal{D}_F(u)$ if $\mathcal{D}_F(u) \neq \emptyset$.

Observe that, whenever the event $\mathcal{E}$ occurs then $X \in \mathcal{X}_\mathcal{E}$ and thus $X \in \mathcal{D}_F(F(X))$. It thus suffices to show that $\mathcal{D}_F(F(X))$ contains no other element, except with probability $\frac{|\mathcal{X}_\mathcal{E}|}{|\mathcal{U}|}$, i.e.,

$$\Pr\big[|\mathcal{D}_F(F(X))| > 1\big] \leq \frac{|\mathcal{X}_\mathcal{E}|}{|\mathcal{U}|} . \tag{1.4}$$

By the definition of two-universality, $\Pr\big[F(x) = F(x')\big] \leq \frac{1}{|\mathcal{U}|}$, for any $x \neq x'$. Consequently, by the union bound, for any fixed $x \in \mathcal{X}$,

$$\Pr\big[\exists x' \in \mathcal{X}_\mathcal{E} : (x \neq x') \wedge (F(x) = F(x'))\big] \leq \frac{|\mathcal{X}_\mathcal{E}|}{|\mathcal{U}|} .$$

This implies (1.4) and thus concludes the proof.

Corollary 1.2. *Let X and Y be random variables and let F be chosen at random from a two-universal family of functions from $\mathcal{X}$ to $\mathcal{U}$, where $|\mathcal{U}| = 2^\ell$. Then there exists a function d_F depending on F such that, for any $\varepsilon \geq 0$,*

$$\Pr\big[d_F(F(X), Y) \neq X\big] \leq 2^{-(\ell - \mathsf{H}^\varepsilon_{\max}(X|Y))} + \varepsilon .$$

Remark 1.1. In order to ensure that errors are corrected except with probability $\bar{\varepsilon}$, it suffices to use a hash function with output length ℓ such that

$$\ell \geq \mathsf{H}^\varepsilon_{\max}(X|Y) + \log\left(\frac{1}{\bar{\varepsilon} - \varepsilon}\right) ,$$

for some $\varepsilon < \bar{\varepsilon}$.

Proof (of Corollary 1.2). Let $\mathcal{E}$ be an event with $\Pr[\mathcal{E}] = 1 - \varepsilon$ such that

$$\mathsf{H}_{\max}(\mathcal{E}X|Y) = \mathsf{H}^\varepsilon_{\max}(X|Y) . \tag{1.5}$$

Let $\mathcal{Y}$ be the range of the random variable Y. By Lemma 1.7, there exists a function d_F from $\mathcal{U} \times \mathcal{Y}$ to $\mathcal{X}$ such that, for any $y \in \mathcal{Y}$,

$$\begin{aligned} \Pr\big[\mathcal{E} \wedge (d_F(F(X), Y) \neq X)\big|Y = y\big] &\leq \frac{|\{x \in \mathcal{X} : P_{\mathcal{E}XY}(x, y) > 0\}|}{|\mathcal{U}|} \\ &= 2^{\mathsf{H}_{\max}(\mathcal{E}X|Y) - \ell} . \end{aligned}$$

Moreover, we have

$$\begin{aligned} \Pr\big[d_F(F(X), Y) \neq X\big] &\leq \Pr\big[\mathcal{E} \wedge (d_F(F(X), Y) \neq X)\big] + (1 - \Pr[\mathcal{E}]) \\ &\leq \max_y \Pr\big[\mathcal{E} \wedge (d_F(F(X), Y) \neq X)\big|Y = y\big] + \varepsilon \end{aligned}$$

Combining this with the above and (1.5) concludes the proof.

1.4.5 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 adversary about S can consist of physical bits, of parities thereof, or other types of information (see Fig. 1.7).

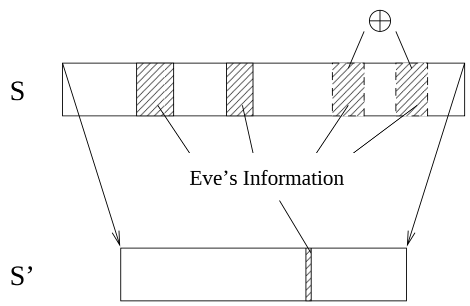

Fig. 1.7. Eliminating Eve's knowledge by privacy amplification.

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.

We analyze the following type of privacy amplification protocols. First, Alice chooses a random function h from a fixed two-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 $S' := h(S)$.

The following is a slightly generalized version of the so-called *Leftover Hash Lemma* [12, 3].

Lemma 1.8. *Let X be a random variable on $\mathcal{X}$, let $\mathcal{E}$ be an event, and let F be chosen at random (and independently of X) from a two-universal family of functions from $\mathcal{X}$ to $\mathcal{U}$. Then*

$$\|P_{\mathcal{E}F(X)F} - P_{\mathcal{E}U} \times P_F\|_1 \leq \sqrt{|\mathcal{U}| \max_x P_{\mathcal{E}X}(x)} ,$$

where $P_{\mathcal{E}U}$ is defined by $P_{\mathcal{E}U}(v) := \frac{\Pr[\mathcal{E}]}{|\mathcal{U}|}$, for any $u \in \mathcal{U}$.

Proof. Using the Cauchy-Schwarz inequality and Jensen's inequality, we find

$$\begin{aligned} \|P_{\mathcal{E}F(X)F} - P_{\mathcal{E}U} \times P_F\|_1 &= \mathbb{E}_F\left[\|P_{\mathcal{E}F(X)} - P_{\mathcal{E}U}\|_1\right] \\ &\leq \mathbb{E}_F\left[\sqrt{|\mathcal{U}| \cdot \|P_{\mathcal{E}F(X)} - P_{\mathcal{E}U}\|_2^2}\right] \\ &\leq \sqrt{|\mathcal{U}| \cdot \mathbb{E}_F\left[\|P_{\mathcal{E}F(X)} - P_{\mathcal{E}U}\|_2^2\right]} . \end{aligned} \tag{1.6}$$

The L_2-norm under the square root can be rewritten as

$$\|P_{\mathcal{E}F(X)} - P_{\mathcal{E}U}\|_2^2 = \|P_{F(X)\mathcal{E}}\|_2^2 - \frac{\Pr[\mathcal{E}]^2}{|\mathcal{U}|} \tag{1.7}$$

where

$$\begin{aligned} \|P_{\mathcal{E}F(X)}\|_2^2 &= \sum_u P_{\mathcal{E}F(X)}(u)^2 \\ &= \sum_u \sum_{\substack{x \in F^{-1}(u) \\ x' \in F^{-1}(u)}} P_{\mathcal{E}X}(x) P_{\mathcal{E}X}(x') \\ &= \sum_{x,x'} P_{\mathcal{E}X}(x) P_{\mathcal{E}X}(x') \delta_{F(x),F(x')} \\ &\leq \sum_x P_{\mathcal{E}X}(x)^2 + \sum_{x \neq x'} P_{\mathcal{E}X}(x) P_{\mathcal{E}X}(x') \delta_{F(x),F(x')} . \end{aligned}$$

Because F is chosen from a two-universal family of functions with range $\mathcal{U}$, we have $\mathbb{E}_F[\delta_{F(x),F(x')}] \leq \frac{1}{|\mathcal{U}|}$, for any $x \neq x'$. Hence,

$$\mathbb{E}_F\left[\|P_{F(X)}\|_2^2\right] = \sum_x P_{\mathcal{E}X}(x)^2 + \frac{\Pr[\mathcal{E}]^2}{|\mathcal{U}|} \leq \max_x P_{\mathcal{E}X}(x) + \frac{\Pr[\mathcal{E}]^2}{|\mathcal{U}|} .$$

Combining this with (1.6) and (1.7) concludes the proof.

Corollary 1.3. *Let X and Z be random variables and let F be chosen at random from a two-universal family of functions from $\mathcal{X}$ to $\mathcal{U}$, where $|\mathcal{U}| = 2^\ell$. Then, for any $\varepsilon \geq 0$,*

$$\|P_{F(X)ZF} - P_U \times P_Z \times P_F\|_1 \leq 2^{-\frac{\mathsf{H}^\varepsilon_{\min}(X|Z)-\ell}{2}} + 2\varepsilon ,$$

where P_U is the uniform distribution on $\mathcal{U}$.

Remark 1.2. Note that the distinguishing probability between a perfect key U and the function output $F(X)$ is given by half the L_1-distance on the left-hand side of the corollary. Hence, in order to get an ℓ-bit key which is $\bar{\varepsilon}$-indistinguishable from a perfect key, it suffices to ensure that

$$\mathsf{H}^\varepsilon_{\min}(X|Z) \geq \ell + 2\log(1/(\bar{\varepsilon} - \varepsilon)) ,$$

for some $\varepsilon < \bar{\varepsilon}$.

Proof (of Corollary 1.3). Let $\mathcal{E}$ be an event with $\Pr[\mathcal{E}] = 1 - \varepsilon$ such that

$$\mathsf{H}_{\min}(\mathcal{E}X|Z) = \mathsf{H}^\varepsilon_{\min}(X|Z) . \tag{1.8}$$

Then, by Lemma 1.8,

$$\|P_{\mathcal{E}F(X)F|Z=z} - \Pr(\mathcal{E}|Z=z)P_U \times P_F\|_1 \leq \sqrt{|\mathcal{U}| \max_x P_{\mathcal{E}X|Z=z}(x)} ,$$

for any value z of the random variable Z. Moreover, by the triangle inequality for the L_1-norm,

$$\begin{aligned}&\|P_{F(X)F|Z=z} - P_U \times P_F\|_1 \\ &\qquad \leq \|P_{\mathcal{E}F(X)F|Z=z} - \Pr(\mathcal{E}|Z=z)P_U \times P_F\|_1 + 2(1 - \Pr[\mathcal{E}|Z=z]) .\end{aligned}$$

Hence,

$$\begin{aligned}\|P_{F(X)ZF} - P_U \times P_Z \times P_F\|_1 &= \mathbb{E}_Z\left[\|P_{F(X)F|Z=z} - P_U \times P_F\|_1\right] \\ &\leq \mathbb{E}_Z\left[\sqrt{|\mathcal{U}| \max_x P_{\mathcal{E}X|Z=z}(x)} + 2(1 - \Pr[\mathcal{E}|Z=z])\right] \\ &\leq \sqrt{|\mathcal{U}| \max_{x,z} P_{\mathcal{E}X|Z=z}(x)} + 2\varepsilon = \sqrt{2^{\ell - \mathsf{H}_{\min}(\mathcal{E}X|Z)}} + 2\varepsilon .\end{aligned}$$

The assertion then follows from (1.8).

1.4.6 Protocol Monotones and Upper Bounds

The described protocol techniques lead to *lower* bounds on the quantity of interest, the secret-key rate S. One is, on the other hand, interested in *upper* bounds on S and, ultimately, determining S precisely; the latter, however, has been successfully done in trivial cases only in the two-way-communication setting.

Characterization of the One-Way Key Rate

In contrast to this, the *one-way* communication scenario has been completely solved [7].

Lemma 1.9. *Let X, Y, and Z be random variables with joint distribution P_{XYZ}. Then*

$$S_{\to}(P_{XYZ}) = \sup_{(U,V)\leftarrow X\leftarrow (Y,Z)} \mathsf{H}(U|ZV) - \mathsf{H}(U|YV) \ .$$

General Properties of Upper Bounds

Before we discuss concrete upper bounds on S, we observe that any quantity which is a so-called *monotone*, i.e., cannot be increased by any protocol and has some additional properties described in Lemma 1.10.

Lemma 1.10. *Let $M(X,Y|Z) = M(P_{XYZ})$ be a real-valued quantity such that the following holds:*

1. *M can only decrease under local operations, i.e., $M(X,Y|Z) \geq M(X',Y|Z)$ if $(Y,Z) \to X \to X'$ is a Markov chain (and likewise for Y).*
2. *M can only decrease if public communication is used, i.e., $M(XC,Y|Z) \geq M(XC,YC|ZC)$, for any random variable C.*
3. *M is asymptotically continuous (as a function of P_{XYZ}).*
4. *M equals one for one key bit, i.e., $M(P_{SS\perp}) = 1$ if $P_{SS\perp}$ denotes the distribution of two identical and uniformly distributed bits.*

Then M is an upper bound on the key rate, i.e., $S(P_{XYZ}) \leq M(P_{XYZ})$.

1.4.7 Intrinsic Information, Information of Formation, and a Gap

In this section, we propose two protocol monotones. The *information of formation* measures the amount of secret-key bits necessary to generate a certain partially secret correlation between Alice and Bob. The *intrinsic information*, on the other hand, measures, intuitively speaking, the correlation two parties share and that is unaccessible to and indestructible by an adversary. Finally, we mention a result showing that an arbitrarily large gap can separate the secrecy required for constructing the distribution from the amount of extractable secrecy.

Information of Formation

Instead of transforming weakly correlated and partially secure data into a secure key, one could also do the opposite [19].

Definition 1.10. *The* information of formation *(also called* key cost*) of a tripartite probability distribution P_{XYZ} is defined by*

$$\mathbf{I}_{\text{form}}(P_{XYZ}) := \text{rate}(\mathsf{SK}_1 \times \mathsf{Auth}_\infty^{A\rightarrow B} \xRightarrow{\texttt{Form}} \mathsf{Source}(P_{XYZ}))^{-1} .$$

It is easy to verify that the information of formation satisfies the assumptions of Lemma 1.10, i.e., $S(P_{XYZ}) \leq \mathsf{I}_{\text{form}}(P_{XYZ})$. Alternatively, the same conclusion can be obtained using Lemma 1.5.

Intrinsic Information

It is straightforward to verify that the mutual information $\mathbf{I}(X;Y)$ as well as the conditional mutual information $\mathbf{I}(X;Y|Z)$ satisfy the assumptions of Lemma 1.10, i.e., they are both upper bounds on the secret-key rate. The following definition is motivated by this observation.

Definition 1.11. *Let X, Y, and Z be random variables with joint distribution P_{XYZ}. The* intrinsic information *is defined by*

$$\mathbf{I}(X;Y \downarrow Z) := \inf_{Z' \leftarrow Z \leftarrow (X,Y)} \mathbf{I}(X;Y|Z') .$$

Again, it is straightforward to verify that $\mathbf{I}(X;Y \downarrow Z)$ satisfies the assumptions of Lemma 1.10, i.e., it is an upper bound on the secret-key rate, $S(P_{XYZ}) \leq \mathsf{I}(X;Y \downarrow Z)$. On the other hand, we have $\mathsf{I}(X;Y \downarrow Z) \leq \mathbf{I}_{\text{form}}(P_{XYZ})$.

The Gap

Interestingly, it has been shown [19] that the gap between S and $\mathbf{I}_{\text{form}}$ can be arbitrarily large, whereas it is still unknown whether there exists a classical analog to *bound quantum entanglement*, i.e., undistillable entanglement: a distribution satisfying $S = 0$ and $\mathbf{I}_{\text{form}} \neq 0$.

Lemma 1.11. *For any $\delta > 0$ there exists a probability distribution such that $S(P_{XYZ}) < \delta$ whereas $\mathbf{I}_{\text{form}}(P_{XYZ}) \geq 1$.*

1.5 Secrecy from Completely Insecure Communication

So far in this chapter, we have considered scenarios where the channel connecting Alice and Bob is authentic. In this section, we show results demonstrating that unconditionally secure key agreement can even be possible from *completely* insecure communication, i.e., a channel over which the adversary has complete control. Clearly, she can always choose to *prevent* key agreement in this case, but it should not happen that Alice or Bob believe that key agreement was successful although it was not. We consider three special scenarios in this setting: Independent repetitions, privacy amplification, and general one-way key agreement.

Independent Repetitions: An All-Or-Nothing Result

In the scenario where a random experiment P_{XYZ} is independently repeated a great number of times, an all-or-nothing result has been shown: Either key agreement is possible at the same rate as in the authentic channel scenario, or completely impossible.

The *robust secret-key rate* $S^*(X;Y||Z)$ is the rate at which a key can be generated in this scenario:

$$S^*(P_{XYZ}) := \mathrm{rate}(\mathsf{Source}(P_{XYZ}) \times \mathsf{Chan}_\infty^{A\leftrightarrow B} \stackrel{\mathtt{KA}}{\Longrightarrow} \mathsf{SK}_1) \ .$$

Here, $\mathsf{Chan}_\infty^{A\leftrightarrow B}$ stands for a completely insecure bidirectional channel.

Theorem 1.3. [14] *If there exists a channel $P_{\overline{X}|Z}$ such that $P_{\overline{X}Y} = P_{XY}$ holds or a channel $P_{\overline{Y}|Z}$ such that $P_{X\overline{Y}} = P_{XY}$, then $S^*(P_{XYZ}) = 0$. Otherwise, $S^*(P_{XYZ}) = S(P_{XYZ})$.*

1.5.1 Privacy Amplification: Authentication is for Free

In the special case of privacy amplification, it has been shown [20] that the loss of the communication channel's authenticity does not (substantially) decrease the length of the extractable key, but the obtained key is only asymmetrically secure, i.e., only one party (Alice) knows whether it is secret. (For a formal definition of such an asymmetry, see Example 1.6)

Theorem 1.4. *Let P_{XYZ} be a distribution where X and Y are identical n-bit strings such that $\mathsf{H}_{\min}(X|Z) \geq tn$, for some fixed $t > 0$. Then there exists a protocol Π such that*

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Chan}_\infty^{A\leftrightarrow B} \xrightarrow{\Pi}_\varepsilon \mathsf{SK}_\ell^A$$

where $\ell = (1 - o(1))tn$ and ε is exponentially small in n.

1.5.2 Robust General One-Way Key Agreement

The result of Theorem 1.4 has been generalized in [21] to the case where Alice's and Bob's strings are not identical initially, i.e., where information reconciliation and privacy amplification have to be combined.

Theorem 1.5. *Let P_{XYZ} be a distribution where X and Y are n-bit strings such that $\mathsf{H}_{\min}(X|Z) - \mathsf{H}_{\max}(X|Y) \geq tn$, for some fixed $t > 0$. Then there exists a protocol Π such that*

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Chan}_\infty^{A\leftrightarrow B} \xrightarrow{\Pi}_\varepsilon \mathsf{SK}_\ell^A$$

where $\ell = (1 - o(1))tn$ and ε is exponentially small in n.

Roughly speaking, Theorem 1.5 states that the length of the extractable key is the difference between Eve's and Bob's uncertainties about Alice's string. Similarly, it is also possible to generate an (unreliable) common secret key.

Theorem 1.6. *Let P_{XYZ} be a distribution where X and Y are n-bit strings such that $\mathsf{H}_{\min}(X|Z) - \mathsf{H}_{\max}(X|Y) - \mathsf{H}_{\max}(Y|X) \geq tn$, for some fixed $t > 0$. Then there exists a protocol Π such that*

$$\mathsf{Source}(P_{XYZ}) \times \mathsf{Chan}_{\infty}^{A \leftrightarrow B} \xrightarrow{\Pi}_{\varepsilon} \mathsf{SK}_{\ell}^{AB}$$

where $\ell = (1 - o(1))tn$ and ε is exponentially small in n.

References

1. R. Ahlswede and I. Csiszár. Common randomness in information theory and cryptography - part i: secret sharing. *IEEE Transactions on Information Theory*, 39:1121–1132, 1993.
2. C. H. Bennett and G. Brassard. Quantum cryptography: Public-key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.
3. C. H. Bennett, G. Brassard, C. Crépeau, and U. Maurer. Generalized privacy amplification. *IEEE Transaction on Information Theory*, 41(6):1915–1923, 1995.
4. C. H. Bennett, G. Brassard, and J.-M. Robert. Privacy amplification by public discussion. *SIAM Journal on Computing*, 17(2):210–229, 1988.
5. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 136–145, 2001.
6. T. M. Cover and J. A. Thomas. *Elements of information theory.* Wiley Series in Telecommunications, 1992.
7. I. Csiszár and J. Körner. Broadcast channels with confidential messages. *IEEE Transactions on Information Theory*, 24:339–348, 1978.
8. W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22:644–654, 1976.
9. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *Advances in Cryptology — EUROCRYPT '04*, Lecture Notes in Computer Science,. Springer, 2004.
10. A. K. Ekert. Quantum cryptography based on Bell's theorem. *Phys. Rev. Lett.*, 67:661, 1991.
11. O. Goldreich, S. Micali, and C. Rackoff. The knowledge complexity of interactive proofs. *SIAM Journal of Computing*, 18(1):186–208, 1989.
12. R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions (extended abstract). In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pages 12–24, 1989.
13. U. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5:53–66, 1992.
14. U. Maurer. Information-theoretically secure secret-key agreement by not authenticated public discussion. In *Advances in Cryptology - EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 209–225. Springer, 1997.

15. U. Maurer. Indistinguishability of random systems. In Lars Knudsen, editor, *Advances in Cryptology — EUROCRYPT '02*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. Springer, 2002.
16. U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39(3):733–742, 1993.
17. B. Pfitzmann and M. Waidner. Composition and integrity preservation of secure reactive systems. In *7th ACM Conference on Computer and Communications Security*, pages 245–254. ACM Press, 2000.
18. R. Renner. *Security of Quantum Key Distribution*. PhD thesis, Swiss Federal Institute of Technology (ETH) Zurich, 2005. Available at `http://arxiv.org/abs/quant-ph/0512258`.
19. R. Renner and S. Wolf. New bounds in secret-key agreement: The gap between formation and secrecy extraction. In *Advances in Cryptology — EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 562–577. Springer-Verlag, May 2003.
20. R. Renner and S. Wolf. Unconditional authenticity and privacy from an arbitrarily weak secret. In *Advances in Cryptology — CRYPTO '03*, Lecture Notes in Computer Science, pages 78–95. Springer, 2003.
21. R. Renner and S. Wolf. The exact price for unconditionally secure asymmetric cryptography. In *Advances in Cryptology — EUROCRYPT '04*, Lecture Notes in Computer Science, pages 109–125. Springer, 2004.
22. R. Renner and S. Wolf. Smooth Rényi entropy and applications. In *Proceedings of 2004 IEEE International Symposium on Information Theory*, page 233. IEEE, 2004.
23. R. Renner and S. Wolf. Simple and tight bounds for information reconciliation and privacy amplification. In *Advances in Cryptology — ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 199–216. Springer-Verlag, 2005.
24. R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1978.
25. Claude E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28:656–715, 1949.
26. G. S. Vernam. Cipher printing telegraph systems for secret wire and radio telegraphic communications. *J. Am. Inst. Elec. Eng.*, 55:109–115, 1926.
27. A. D. Wyner. The wire-tap channel. *Bell System Technical Journal*, 54:1355–1387, 1975.
28. R. W. Yeung. A new outlook on Shannon's information measures. *IEEE Transactions on Information Theory*, 37:466–474, 1991.