Secret-Key Agreement over Unauthenticated Public Channels – Part III: Privacy Amplification

Ueli Maurer, *Fellow, IEEE*   Stefan Wolf

***Abstract*— This is the third part of a three-part paper on secret-key agreement secure against active adversaries. Here, we consider the special case where the legitimate partners already share a mutual string which might, however, be partially known to the adversary. The problem of generating a secret key in this case has been well studied in the passive-adversary model—for instance in the context of quantum key agreement—under the name of *privacy amplification*. We consider the same problem with respect to an active adversary and propose two protocols, one based on universal hashing and one based on extractors, allowing for privacy amplification secure against an adversary whose knowledge about the initial partially secret string is limited to one third of the length of this string. Our results are based on novel techniques for authentication secure even against adversaries knowing a substantial amount of the "secret" key.**

Keywords. Cryptography, unconditional security,

secret-key agreement, authentication, privacy amplification, quantum key agreement.

I. Motivation, Definition, and Preliminaries

A. Introduction

A special case of the general key agreement scenario defined in [15] is the situation where the parties Alice and Bob already share a string $X = Y = S$, about which, however, the adversary has possibly substantial information. The problem of transforming this partially secret string into a virtually secret key S' is called *privacy amplification*; it is the final phase of many key-agreement protocols.

Privacy amplification was first described in the context of quantum key agreement by Bennett *et al.* [2], where universal hashing was shown to be a good technique in the case where the adversary possesses *deterministic* information about S. More precisely, it was shown that the key S' must be shorter than S, and that $\mathrm{len}(S) - \mathrm{len}(S')$ must be equal to the amount of information the adversary has about S, plus a security parameter. This result was generalized by Bennett *et al.* [1] to *probabilistic* information about S. The length of S' can in this case be roughly equal to the *Rényi entropy* of S from Eve's viewpoint.

In this paper we investigate the same problem under the assumption that the communication between Alice and Bob is *not* authenticated. Note that, in contrast to the model where many independent repetitions of the involved random variables are given [15], the same piece of information S must be used here both for authentication and as the input for privacy amplification. Two problems that arise in this context are authentication with an only partially secret key, and the fact that this authentication leaks information about S, hence potentially also about S', to the adversary. We show that for our purpose, a new, interactive, authentication method is better than one-way authentication by strongly universal hashing, and that so-called *extractors*, requiring fewer random bits (i.e., shorter messages to be communicated) are a better technique for privacy amplification than universal hashing.

The outline of this paper is as follows. In Section I-B we define the notion of a protocol for privacy amplification by completely insecure communication. This is a modified version of the protocol definition for the scenario of independent realizations as given in [15]. In Section I-C we show some impossibility results. Section I-D analyzes, as a preparation, the effect of side information on certain important entropy measures, and connects the entropy of strings and parts thereof. Then, we present two different protocols for privacy amplification secure against active adversaries.

Protocol UH (Section II) is based on universal hashing, whereas Protocol EX (Section III) uses extractors for transforming S to S'. It is shown in Section V that each of these protocols can be better than the other in certain situations. The used techniques for authentication and identification are introduced in Sections II-A, II-B, and III-A.

B. Protocol Definition

The protocol definition for privacy amplification secure against active adversaries can be strengthened in two respects as compared to the definition in the general case [15]. First, we require that Alice and Bob both accept and end up with the same string with probability 1 if Eve is passive. Moreover, the protocols can work for an entire class of distributions P_{XYZ} instead of only one distribution. More precisely, Eve's knowledge about the mutual n-bit string S is limited by assuming that $P_{S|Z=z}$ is, for all $z \in \mathcal{Z}$, contained in some subset $\mathcal{D}$ of all possible distributions over the set $\{0,1\}^n$. Typically, $\mathcal{D}$ consists of the distributions satisfying a certain condition in terms of[1] Rényi entropy or min-entropy. We denote, for every t and $\alpha = 2$ or $\alpha = \infty$, by $\mathcal{D}_{n,\alpha,t}$ the subset $\{P_X \mid H_\alpha(X) \geq t\}$ of distributions over n-bit strings.

[1] For a random variable X with range $\mathcal{X}$ and distribution P_X, the *Rényi entropy* $H_2(X)$ is defined as $H_2(X) := -\log(\sum_{x\in\mathcal{X}} P_X(x)^2)$. The *min-entropy* $H_\infty(X)$ is $H_\infty(X) := -\log \max_{x\in\mathcal{X}} P_X(x)$. All logarithms here and in the rest of the paper are binary, unless stated otherwise.

Definition 1: Assume that Alice and Bob both know an n-bit random variable S, and that the random variable Z summarizes Eve's entire knowledge about S. Let $\mathcal{D}$ be a subset of all probability distributions on the set of n-bit strings, let r be an integer, and let $\varepsilon, \delta > 0$. An *$(n, \mathcal{D}, r, \varepsilon, \delta)$-protocol for privacy amplification by communication over an insecure and unauthenticated channel* (a *robust $(n, \mathcal{D}, r, \varepsilon, \delta)$-PA-protocol* for short) is a key-agreement protocol, as defined in [15], with the following properties.

1. Correctness and Privacy. Let Eve be a passive wire-tapper receiving a particular value $Z = z$ satisfying $P_{S|Z=z} \in \mathcal{D}$. Then both Alice and Bob must accept at the end of the protocol, and there must exist an r-bit string S' such that $S' = S'_A = S'_B$ and $H(S'|C, Z = z) \geq r - \varepsilon$ hold. In this case, we say that privacy amplification has been *successful*.

2. Robustness. Let $P_{S|Z=z} \in \mathcal{D}$. For every possible strategy of Eve, the probability that either *both* Alice and Bob reject the outcome of the protocol, or privacy amplification has been successful, must be at least $1 - \delta$.

C. Impossibility Results

Clearly, the impossibility results of [15] immediately carry over to privacy amplification secure against active opponents (where the non-simulatability condition is fulfilled in all non-trivial cases). There exists neither a protocol with perfect synchronization of the accepting states (i.e., both accept or both reject in every case), nor a one-way-transmission protocol satisfying the required properties.

Theorem 1: Let $\alpha = 2$ or $\alpha = \infty$. Assume that a robust $(n, \mathcal{D}_{n,\alpha,t}, r, \varepsilon, \delta)$-PA-protocol either with perfect synchronization or using only one-way transmission exists. Then either $\varepsilon \geq r-1$, or $\varepsilon > n - t + \log(1 - 2^{\varepsilon - r+1})$, or $\delta = 1$ holds.

Proof. Assume $\varepsilon < r - 1$ and $\varepsilon \leq n - t + \log(1 - 2^{\varepsilon - r+1})$. We show that there exists, for every fixed function $f : \{0,1\}^n \to \{0,1\}^r$ (on which Alice and Bob could agree without any communication) a distribution $P_{S|Z=z} \in \mathcal{D}$ such that

$$H(f(S)|Z = z) < r - \varepsilon \ .$$

Let $\mathcal{A} \subseteq \{0,1\}^r$ be the particular set of size

$$2^r - \lfloor 2^{r-\varepsilon} \rfloor + 1$$

which minimizes the cardinality of the set

$$f^{-1}(\mathcal{A}) := \{s \in \{0,1\}^n \ : \ f(s) \in \mathcal{A}\} \ .$$

Then

$$\begin{aligned} |f^{-1}(\mathcal{A})| &\leq 2^{n-r}(2^r(1 - 2^{-\varepsilon}) + 2) \\ &= 2^n(1 - 2^{-\varepsilon} + 2^{-r+1}) \ . \end{aligned}$$

For $\mathcal{B} := \{0,1\}^n \setminus f^{-1}(\mathcal{A})$ we have hence

$$|\mathcal{B}| \geq 2^{n-\varepsilon}(1 - 2^{\varepsilon - r+1}) \ .$$

Let $P_{S|Z=z}$ be the uniform distribution on $\mathcal{B} \subseteq \{0,1\}^n$. Then we have

$$\begin{aligned} H_\alpha(S|Z=z) &= \log|\mathcal{B}| \\ &\geq n - \varepsilon + \log(1 - 2^{\varepsilon - r+1}) \\ &\geq t \end{aligned}$$

and

$$H(f(S)|Z = z) \leq \log(2^r - |\mathcal{A}|) < r - \varepsilon$$

by construction and by the assumption. This contradicts the protocol definition, hence at least one message must be sent in the protocol. The rest of the argument is as in the proofs of Theorems 8 and 9 in [15]. □

D. The Effect of Side Information and Knowledge About Partial Strings

In this section we provide some facts necessary for the analysis of Protocols UH and EX for privacy amplification described below. We derive bounds on the amount of knowledge (e.g., of an adversary) in terms of Rényi entropy and min-entropy about a partial string, depending on the amount of knowledge about the entire string. This is done both for the cases where the adversary does (Corollary 2) or does not (Lemma 1) obtain information about the remaining part of the string. In both cases, the result is roughly the intuitive fact that (with high probability) one cannot know (substantially) more about a part than about the whole (see Figure 1). In the case where the adversary

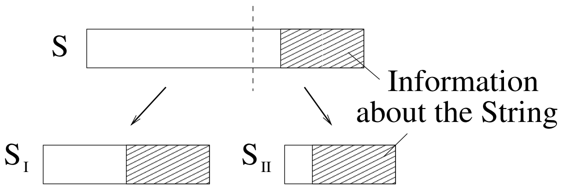

Fig. 1. Information About Partial Strings

obtains information about the remaining part of the string, the result follows from a general upper bound on the reduction of Rényi entropy and min-entropy of a random variable when side information is given (Lemma 2).

Lemma 1: Let $S = (S_1, S_2, \ldots, S_n)$ be a random variable consisting of n binary random variables. For any k-tuple $\underline{i} = (i_1, i_2 \ldots, i_k)$ with $1 \le i_1 < i_2 < \cdots < i_k \le n$, let $S_{\underline{i}}$ be the string $(S_{i_1}, S_{i_2}, \ldots, S_{i_k})$. Then

$$H_\alpha(S_{\underline{i}}) \ge H_\alpha(S) - (n-k)$$

holds for $\alpha = 2$ and $\alpha = \infty$.

Proof. Let first $\alpha = 2$. Consider a fixed string $(s_{i_1}, \ldots, s_{i_k})$. This particular value of the random variable $S_{\underline{i}}$ corresponds to exactly 2^{n-k} values $(s_1, \ldots, s_n)$ of the random variable S. Let $p_1, \ldots, p_{2^{n-k}}$ be the probabilities of these strings, and let $p_0 := \sum_{i=1}^{2^{n-k}} p_i$. Now we have

$$\begin{aligned}
\log\left(\sum_{i=1}^{2^{n-k}}\left(\frac{p_i}{p_0}\right)^2\right) &= \log\left(\mathrm{E}\left[\frac{p_i}{p_0}\right]\right) \\
&\ge \mathrm{E}\left[\log\left(\frac{p_i}{p_0}\right)\right] \\
&= -\mathrm{E}\left[\log\left(\frac{p_0}{p_i}\right)\right] \\
&\ge -\log\left(\mathrm{E}\left[\frac{p_0}{p_i}\right]\right) \\
&= -\log\left(\sum_{i=1}^{2^{n-k}}\frac{p_i}{p_0}\cdot\frac{p_0}{p_i}\right) \\
&= -\log(2^{n-k}) \\
&= k-n\ .
\end{aligned}$$

Here, the expectation is with respect to the probability distribution p_i/p_0 over the 2^{n-k} strings. We have made double use of the fact that the logarithm is a concave function, and of Jensen's inequality. We conclude that

$$\sum_{i=1}^{2^{n-k}} p_i^2 \ge \frac{p_0^2}{2^{n-k}}\ . \qquad (1)$$

Because inequality (1) holds for every particular string $(s_{i_1}, \ldots, s_{i_k})$, we have for the collision probabilities[2] P_C of the random variables

[2]For a random variable X with range $\mathcal{X}$, the *collision probability* $P_C(X)$ is the probability of getting the same outcome twice in two independent realizations, i.e., $P_C(X) = \sum_{x\in\mathcal{X}} P_X(x)^2$. The Rényi entropy of X is then $H_2(X) = -\log(P_C(X))$.

S and $S_{\underline{i}}$:

$$\begin{aligned} P_C(S_{\underline{i}}) &= \sum_{\{0,1\}^k} P_{S_{\underline{i}}}((s_{i_1},\ldots,s_{i_k}))^2 \\ &\leq 2^{n-k}\cdot \sum_{\{0,1\}^n} P_S((s_1,\ldots,s_n))^2 \\ &= 2^{n-k}\cdot P_C(S)\ . \end{aligned}$$

Hence

$$H_2(S_{\underline{i}}) \geq H_2(S) - (n-k)\ .$$

For the case $\alpha = \infty$, the inequality follows directly from the fact that the maximal probability of a k-bit string is at most 2^{n-k} times the maximal probability of a string in S. □

Remark. Note that when the string S is split into two parts S_l and S_r, then the bounds of Lemma 1 applied to S_l and S_r are tight *simultaneously*. For example let $\alpha = \infty$ and $s = (s_l, s_r)$ be a particular n-bit string (where n is even), and let s_l and s_r be the first and second halves of s. Define (for some $v \leq n/2-1$) $P_S((s_l,\overline{s})) = P_S((\overline{s},s_r)) := 2^{v-n}$ for all $n/2$-bit strings $\overline{s}$ (and a uniform distribution for the remaining n-bit strings), i.e., $H_\infty(S) = n - v$. Then

$$H_\infty(S_l) = H_\infty(S_r) = n/2 - v = H_\infty(S) - n/2\ .$$

Intuitively speaking, Eve's information about S in terms of min-entropy appears entirely in both substrings S_l and S_r, a fact that might contradict one's intuition.

Lemma 2 gives an upper bound on the reduction of the Rényi entropy and min-entropy $H_2(P)$ and $H_\infty(P)$ of a random variable P when side information $[Q,R]$ (consisting of a pair of random variables) is given, where $I(P;R) = 0$. It states that this reduction exceeds $\log|\mathcal{Q}|$ (where $\mathcal{Q}$ is the range of Q) substantially only with small probability in both cases. (Note that it is not a trivial fact that no additional reduction is induced by R if $I(P;R) = 0$. For instance, $I(P;Q) = 0$ and $I(P;R) = 0$ together do *not* imply that $H_2(P|Q=q,R=r) = H_2(P)$, as the example $P = Q \oplus R$ shows.)

Lemma 2: Let P, Q, and R be random variables with $I(P;R) = 0$. Then

$$\mathrm{P}_{QR}[H_2(P\,|\,Q=q,R=r) \geq H_2(P) - \log|\mathcal{Q}| - s] \geq 1 - 2^{-(s/2-1)}$$

for all $s > 2$, and

$$\mathrm{P}_{QR}\,[H_\infty(P\,|\,Q=q,R=r) \geq H_\infty(P) - \log|\mathcal{Q}| - s] \geq 1 - 2^{-s}$$

for $s > 0$.

Proof. We first prove the statement concerning Rényi entropy. The argument is a generalization of the proof of Theorem 4.17 in [3]. Let $r \in \mathcal{R}$ be fixed. It is straight-forward that

$$2^{-H_2(PQ|R=r)} = \mathrm{E}_Q\left[2^{\log P_{Q|R=r} - H_2(P\,|\,Q=q,R=r)}\right]\ .$$

Hence the probability that $\log P_{Q|R=r} - H_2(P|Q=q, R=r)$ exceeds $-H_2(PQ|R=r)$ by more than $s/2$ is at most $2^{-s/2}$, i.e., $\mathrm{P}_Q[H_2(P\,|\,Q=q, R=r) \leq H_2(PQ\,|\,R=r) + \log P_{Q|R=r} - s/2] \leq 2^{-s/2}$. Furthermore,

$$\mathrm{P}_Q\left[\log P_{Q|R=r} \leq -\log|\mathcal{Q}| - s/2\right] \leq 2^{-s/2} \ .$$

These inequalities together imply $\mathrm{P}_Q[H_2(P\,|\,Q=q, R=r) \leq H_2(PQ\,|\,R=r) - \log|\mathcal{Q}| - s] \leq 2^{-(s/2-1)}$. Finally, $H_2(PQ|R=r) \geq H_2(P|R=r) = H_2(P)$ holds because of

$$\begin{aligned}
&\sum_{p\in\mathcal{P},\, q\in\mathcal{Q}} P_{PQ|R=r}(p,q)^2 \\
&= \sum_{p\in\mathcal{P}} \left(P_{P|R=r}(p)^2 \cdot \sum_{q\in\mathcal{Q}} P_{Q|P,R=r}(q,p)^2 \right) \\
&\leq \sum_{p\in\mathcal{P}} P_{P|R=r}(p)^2 \\
&= \sum_{p\in\mathcal{P}} P_P(p)^2 \ .
\end{aligned}$$

We have used for the last equality that P and R are statistically independent, i.e., that $P_{P|R=r} = P_P$. We conclude that

$$\begin{aligned}
&\mathrm{P}_Q[H_2(P\,|\,Q=q, R=r) \leq H_2(P) - \log|\mathcal{Q}| - s] \\
&\leq 2^{-(s/2-1)}
\end{aligned}$$

holds for all $r \in \mathcal{R}$, and the first statement of the lemma follows.

Let us address the second statement. Let $p_0 := 2^{-s}/|\mathcal{Q}|$. Then we have for all $r \in \mathcal{R}$

$$P_{Q|R=r}[\{q \ : \ P_{Q|R=r} < p_0\}] < 2^{-s} \ ,$$

and hence

$$P_{QR}\left[\{(q,r) \in \mathcal{Q}\times\mathcal{R} \ : \ P_{Q|R}(q,r) < p_0\}\right] < 2^{-s} \ .$$

This inequality implies that

$$\begin{aligned}
P_{P|QR}(p,q,r) &= \frac{P_{PQR}(p,q,r)}{P_{QR}(q,r)} \\
&= \frac{P_P(p)\cdot P_R(r)\cdot P_{Q|PR}(q,p,r)}{P_R(r)\cdot P_{Q|R}(q,r)} \\
&\leq \frac{P_P(p)}{P_{Q|R}(q,r)} \\
&\leq \frac{P_P(p)}{p_0} \\
&= P_P(p)\cdot|\mathcal{Q}|\cdot 2^s
\end{aligned}$$

holds with probability greater than $1 - 2^{-s}$ (taken over Q and R). The statement follows by maximizing over all $p \in \mathcal{P}$, and by taking negative logarithms. □

The following corollary is a direct consequence of Lemma 2. It states that a formally slightly weaker result than that of Lemma 1, concerning the knowledge (in terms of H_2 and of H_∞) of a partial string, even holds when the rest of the string is made public.

Corollary 2: Let S be an n-bit string, and let a partition of S into two strings S' and S'' of lengths l and $n-l$, respectively, be given. Let $s > 2$ be a security parameter. Then the probability, taken over s'', that

$$H_2(S'\,|\,S''=s'') \geq H_2(S) - (n-l) - s$$

holds is at least $1 - 2^{-(s/2-1)}$. Furthermore, for $s > 0$, the probability, taken over s'', that

$$H_\infty(S'\,|\,S''=s'') \geq H_\infty(S) - (n-l) - s$$

holds is at least $1 - 2^{-s}$.

II. PROTOCOL UH BASED ON UNIVERSAL HASHING

A. Message Authentication with a Partially Secret Key I: Strongly Universal Hashing

All previous results on unconditionally secure authentication require a key that is completely secret, i.e., its probability distribution is uniform from the opponent's viewpoint. In this section we prove a result on authentication where the opponent is allowed to have some partial information about the key. These techniques are used in the protocols described in the following sections.

There exists a variety of possibility and impossibility results on information-theoretically secure authentication (see for example [20], [11], or [21]). The following two types of attacks are possible. In an *impersonation attack*, the opponent tries to generate a (correctly authenticated) message, and in a *substitution attack*, the adversary observes a correctly authenticated message and tries to replace it by a different correctly authenticated message. The success probabilities are denoted by p_{imp} and p_{sub}, respectively. (General lower bounds on these probabilities are given in [11].)

One possibility for realizing information-theoretically secure authentication is by using strongly universal classes of hash functions (see for example [21]).

Definition 2: A class $\mathcal{H}$ of (hash) functions $\mathcal{A} \to \mathcal{B}$ is called *strongly universal* (or SU_2 for short) if for all distinct $a_1, a_2 \in \mathcal{A}$ and for all $b_1, b_2 \in \mathcal{B}$, the number of functions $h \in \mathcal{H}$ for which both $h(a_1) = b_1$ and $h(a_2) = b_2$ hold is $|\mathcal{H}|/|\mathcal{B}|^2$.

Remark. Note that a strongly universal class has in particular the following property. For every $a \in \mathcal{A}$ and $b \in \mathcal{B}$, the number of functions $h \in \mathcal{H}$ such that $h(a) = b$ holds is $|\mathcal{H}|/|\mathcal{B}|$. This is true because for all $a, a' \in \mathcal{A}$, $a' \neq a$, and $b \in \mathcal{B}$, we have

$$
\begin{aligned}
&|\{h \in \mathcal{H} \,:\, h(a) = b\}| \\
&= \left| \bigcup_{b' \in \mathcal{B}} \{h \in \mathcal{H} \,:\, h(a) = b\,,\ h(a') = b'\} \right| \\
&= \sum_{b' \in \mathcal{B}} |\{h \in \mathcal{H} \,:\, h(a) = b\,,\ h(a') = b'\}| \\
&= |\mathcal{B}| \cdot \frac{|\mathcal{H}|}{|\mathcal{B}|^2} \ = \ \frac{|\mathcal{H}|}{|\mathcal{B}|}\ .
\end{aligned}
$$

By roughly the same argument one can also show that a strongly universal class is in particular universal (see Definition 3); a fact that is suggested by the names of the properties.

A strongly universal class of hash functions can immediately be used for authentication: the secret key determines a hash function of the class, and the message is authenticated by its hash value. The authentication code corresponding to an SU_2 class of hash function satisfies

$$p_{imp} = 1/|\mathcal{B}|$$

(because of the property mentioned in the above remark) and

$$p_{sub} = 1/|\mathcal{B}|$$

(which follows directly from the definition). An SU_2 class of functions mapping N-bit strings to N-bit strings can be constructed similarly to the universal class described in [15]. Namely, the class

$$\mathcal{H} = \{h_{ab} \,:\, (a,b) \in (GF(2^N))^2\} \ , \quad (2)$$

where

$$h_{ab}(x) := a \cdot x + b \ ,$$

is an SU_2 class of hash functions $\{0,1\}^N \to \{0,1\}^N$ with 2^{2N} elements, i.e., with a key $S = a||b$ of length $2N$.

Let us now investigate the scenario in which the key is not entirely secret, i.e., where the opponent Eve has a certain amount of information about this key. The following result states that information-theoretically secure authentication is possible under the condition that the Rényi entropy of the key from the adversary's viewpoint is greater than half the length of the key.

Theorem 3: Let S be a binary string of (even) length n. Assume that S is used by two parties as the key in the authentication scheme based on strongly universal hashing with respect to the class (2), that an adversary knows a random variable Z, jointly distributed with S according to some probability distribution, and that the opponent has no further information about S. Let

$$H_2(S|Z = z) \geq (1/2 + R) \cdot N$$

for a particular z in the range $\mathcal{Z}$ of Z. Then the probabilities of successful impersonation and substitution attacks, given $Z = z$, are upper bounded by

$$p_{imp} \leq 2^{-RN/2}$$

and

$$p_{sub} \leq 3 \cdot 2^{-RN/4} \ , \quad (3)$$

respectively.

Remark. Note that in Theorem 3 it need not be assumed that the message observed by Eve be independent of S (but independent of S *given* $Z = z$). For example, inequality (3) holds even when the message is selected by Eve herself.

Proof. First we prove the upper bound on the success probability p_{imp} of the impersonation attack. For every possible message $m \in GF(2^{N/2})$ and for every authenticator $y \in GF(2^{N/2})$ there exist exactly $2^{N/2}$ possible keys such that y is the correct authenticator for m. The probability of such a set of keys, given that $Z = z$, can be upper bounded as follows. In the worst case (i.e., the best case for the impersonating attacker) the $2^{N/2}$ keys all have the same probability, say p. Then p must satisfy

$$2^{N/2} \cdot p^2 \leq P_C(S) \leq 2^{-(1/2+R)N} ,$$

i.e.,

$$p \leq 2^{-(1/2+R/2)N} .$$

Hence

$$p_{imp} \leq 2^{N/2} \cdot 2^{-(1/2+R/2)N} = 2^{-RN/2} .$$

Let us now consider the substitution attack. The crucial argument is that the key s is uniquely determined by $(m, h_s(m))$ and $(m', h_s(m'))$ if $m \neq m'$. Hence the probability of a successful substitution attack is not greater than the probability of guessing S correctly when given $(M, h_S(M))$. From Lemma 2, and because $I(S; M|Z = z) = 0$, we can conclude that

$$H_2(S \,|\, M = m, h_S(M) = h_s(m), Z = z) \geq RN/2 \tag{4}$$

holds with probability at least $1 - 2^{-(RN/4-1)}$. On the other hand, if inequality (4) holds, then the maximal probability of a particular key s is at most

$$\sqrt{2^{-H_2(S \,|\, M=m, h_S(M)=h_s(m), Z=z)}} \leq 2^{-RN/4} .$$

Thus we have, by the union bound,

$$p_{sub} \leq 2^{-(RN/4-1)} + 2^{-RN/4} = 3 \cdot 2^{-RN/4} .$$

$\square$

Remark. It has been proposed to use smaller but "weaker" classes of functions, so-called ε-almost strongly universal (ε-ASU) hash functions, instead of strongly universal hashing for authentication [21]. Such classes allow for authentication with a substantially smaller secret key at the price of a somewhat greater success probability of a substitution attack. However, for the purpose of authentication with a *partially* secret key, these classes of functions do not lead to better results. Whenever the Rényi entropy of the partially secret key is smaller than half the length of the key, then no unconditionally secure authentication is possible with this key by using (ε-A)SU hashing because one correct message-authenticator pair can reveal the remaining information necessary to uniquely determine the key.

B. Challenge-Response Identification with a Highly Insecure Key

In the previous section we have shown that message authentication is possible with a partially secret key, or more precisely, with a key the Rényi entropy of which (from the adversary's viewpoint) is more than half its length. In this section on the other hand we prove that a certain type of security against active attacks can even be achieved when the key shared by the legitimate partners is highly insecure (e.g., in terms of Rényi entropy). A challenge-response scheme is described that can successfully be attacked only by an adversary having almost complete knowledge about the secret key. This method is used as the final step in both the Protocols UH (Section II-D) and EX (Section III-C). The purpose of this step is to prevent the party sending the final message that is needed for successful secret-key agreement from accepting although key agreement has failed. In Section III-A, a related result is proved that shows how the same scheme can be used for authenticating short messages.

Lemma 3: Let N and ℓ be integers such that ℓ divides N and $2^\ell \geq N/\ell$ holds, and let K be a random variable with range $\mathcal{K} \subseteq GF(2^N)$. Let further for any $d \in GF(2^\ell)$ the function $f_d : \{0,1\}^N \to \{0,1\}^\ell$ be defined as

$$f_d(x) := \sum_{i=0}^{N/\ell-1} d^i x_i \ ,$$

where $(x_0, \ldots, x_{N/\ell-1}) \in (GF(2^\ell))^{N/\ell}$ is a representation of $x \in GF(2^N)$ with respect to a fixed basis of $GF(2^N)$ over $GF(2^\ell)$, where the computations are carried out in the field $GF(2^\ell)$, and where the elements of $GF(2^\ell)$ are represented as ℓ-bit strings with respect to a fixed basis of $GF(2^\ell)$ over $GF(2)$. Assume that for $d \in_R GF(2^\ell)$, the value $f_d(K)$ can be guessed correctly (with some strategy) with probability $\alpha \geq (N/\ell)/2^\ell$, taken over the distribution of K, the choice of d, and the coin tosses of the guessing strategy. Then

$$H_2(K) \leq -\frac{2N}{\ell} \cdot \log\left(\alpha - \frac{N/\ell}{2^\ell}\right) \tag{5}$$

or, equivalently,

$$\alpha \leq 2^{-(\ell/2N)\cdot H_2(K)} + \frac{N/\ell}{2^\ell} \ .$$

Proof. First, we can assume without loss of generality that the strategy of guessing $f_d(k)$ is deterministic, since for every possible strategy there exists a deterministic strategy that is at least as good (since every randomized strategy can be seen as a combination of deterministic strategies, of which the optimal one can be chosen).

We give a lower bound on the probability β that for randomly and independently chosen distinct arguments $d_1, \ldots, d_{N/\ell}$ of $GF(2^\ell)$, all the values $f_{d_i}(k)$ are guessed correctly. Let $g(x)$ be the function

$$g(x) := x \cdot \left(x - \frac{1}{2^\ell}\right) \cdot \left(x - \frac{2}{2^\ell}\right) \cdots \left(x - \frac{N/\ell - 1}{2^\ell}\right)$$

if $x \geq (N/\ell - 1)/2^\ell$ and $g(x) := 0$ otherwise. Let, for every $k \in \mathcal{K}$, n_k denote the number of distinct $d \in GF(2^\ell)$ for which $f_d(k)$ is guessed correctly by the (deterministic) guessing strategy. Then we have

$$\begin{aligned} \beta &= \mathrm{E}_K\left[g\left(\frac{n_k}{2^\ell}\right)\right] \\ &\geq g\left(\mathrm{E}_K\left[\frac{n_k}{2^\ell}\right]\right) \\ &\geq \left(\alpha - \frac{N/\ell}{2^\ell}\right)^{N/\ell} . \end{aligned}$$

(Here, we have made use of Jensen's inequality. Note that g is a convex function.)

Hence there exist N/ℓ distinct $d_1,\ldots,d_{N/\ell}$ such that the values $f_{d_i}(k)$ are simultaneously guessed correctly with probability at least $(\alpha-(N/\ell)/2^\ell)^{N/\ell}$, taken over k.

On the other hand, k is uniquely determined by the correct values $f_{d_i}(k)$, $k=1,\ldots,N/\ell$. In order to see this, note first that

$$\begin{pmatrix} f_{d_1}(k) \\ f_{d_2}(k) \\ \vdots \\ f_{d_{N/\ell}}(k) \end{pmatrix} =$$

$$\begin{pmatrix} d_1^0 & d_1^1 & \cdots & d_1^{N/\ell-1} \\ d_2^0 & d_2^1 & \cdots & d_2^{N/\ell-1} \\ \vdots & \vdots & & \vdots \\ d_{N/\ell}^0 & d_{N/\ell}^1 & \cdots & d_{N/\ell}^{N/\ell-1} \end{pmatrix} \cdot \begin{pmatrix} k_0 \\ k_1 \\ \vdots \\ k_{N/\ell-1} \end{pmatrix} ,$$

and second, that the determinant of the matrix, called *Vandermonde determinant*, is equal to

$$\prod_{1\le i<j\le N/\ell} (d_j - d_i) \neq 0 ,$$

hence the matrix is invertible, and $k = (k_0,k_1,\ldots,k_{N/\ell-1})$ is uniquely determined by the $f_{d_i}(k)$'s. An alternative way to see this fact is by interpreting $f_d(x)$ as a polynomial $P_x(d)$ of degree at most $N/\ell-1$ over $GF(2^\ell)$, which is uniquely determined, thus also x is, by its evaluation at N/ℓ distinct points $d_1,\ldots,d_{N/\ell}$. Hence there must be an element $k_0\in\mathcal{K}$ with

$$P_K(k_0) \ge \left(\alpha - \frac{N/\ell}{2^\ell}\right)^{N/\ell} .$$

Because of $P_C(K) \ge P_K(k_0)^2$, we can conclude that (5) holds. □

C. Privacy Amplification by Universal Hashing

Assume that Alice and Bob share an N-bit string w about which an eavesdropper Eve has incomplete information characterized by a probability distribution P_W over the N-bit strings, and that Alice and Bob have some knowledge of this distribution P_W, but that they do not know exactly in which way the secrecy of their string is compromised. Using the public-discussion channel they wish to agree on a function $g:\{0,1\}^N\to\{0,1\}^M$ (for some suitable M) such that Eve, despite her partial knowledge about w and complete knowledge of g, almost certainly knows nearly nothing about $g(w)$. This process transforms a partially secret N-bit string w into a highly secret but shorter M-bit string $g(w)$.

The two natural questions in this context are what a good technique is for computing the compressed from the initial string, and how long the virtually secret string can be, depending on this technique and on P_W. Bennett, Brassard, and Robert [2] considered the case where Eve receives *deterministic* information, i.e., where the key is, from Eve's viewpoint, uniformly distributed over a subset of the set of all possible keys. They used universal hashing as the technique for compressing the string.

Definition 3: **[5]** A class G of functions g :

$\mathcal{A} \longrightarrow \mathcal{B}$ is *universal*$_2$ ("universal" for short) if, for any distinct x_1 and x_2 in $\mathcal{A}$, the probability that $g(x_1) = g(x_2)$ holds is at most $1/|\mathcal{B}|$ when g is chosen at random from G according to the uniform distribution.

The following is an example of a universal class of functions from $\{0,1\}^N$ to $\{0,1\}^M$, for $M \leq N$, with 2^N elements [1].

Example 1: Let a be an element of $GF(2^N)$, and interpret $x \in \{0,1\}^N$ as an element of $GF(2^N)$ with respect to a fixed basis of the extension field over the prime field $GF(2)$. Consider the function $h_a : \{0,1\}^N \to \{0,1\}^M$ assigning to an argument x the first M bits (with respect to this basis representation) of the element ax of $GF(2^N)$, i.e.,

$$h_a(x) := \mathrm{LSB}_M(a \cdot x) \ .$$

The class

$$\{h_a \,|\, a \in GF(2^N)\}$$

is a universal class of functions for $1 \leq M \leq N$.

The results of [2] were generalized by Bennett, Brassard, Crépeau, and Maurer [1] to scenarios in which Eve's information about w is not deterministic, but where the probability distribution P_W satisfies a constraint in terms of Rényi entropy. The main result of [1] is the following theorem (see also Figure 2).

Theorem 4: **[1]** Let P_W be a probability distribution over $\mathcal{W}$ with Rényi entropy $H_2(W)$,

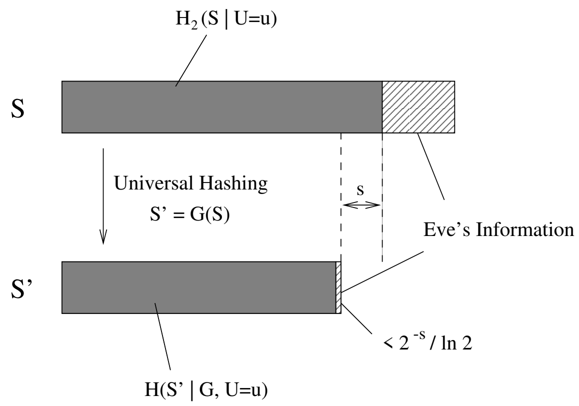

Fig. 2. Universal Hashing Allows to Extract Rényi Entropy

and let G be the random variable corresponding to the random choice, with respect to the uniform distribution, of an element of a universal class of functions mapping $\mathcal{W}$ to $\{0,1\}^M$. Then

$$\begin{aligned} H(G(W)|G) &\geq H_2(G(W)|G) \\ &\geq M - \frac{2^{M-H_2(W)}}{\ln 2} \ . \end{aligned}$$

Theorem 4 states that if Alice and Bob share a particular string S and Eve's information about S corresponds to the distribution $P_{S|Z=z}$ (where z denotes the particular value of her information Z) about which Alice and Bob know nothing except a lower bound R on the Rényi entropy, i.e., $H_2(S|Z=z) \geq R$, then Alice and Bob can generate a secret key S' of roughly R bits. More precisely, if Alice and Bob compress S slightly more to an $(R-s)$-bit key for some security parameter $s > 0$, then Eve's total information about this key is exponentially small in s.

A natural problem that arises when combining information reconciliation and privacy amplification with universal hashing is to determine the effect of the error-correction information (leaked also to the adversary) on the Rényi entropy of the partially secret string, given Eve's information. The following result, which was shown by Cachin [3] as an improvement of an earlier result by Cachin and Maurer [4], states that leaking t physical bits of arbitrary side information about a random variable cannot reduce its Rényi entropy by substantially more than t except with exponentially small probability. Note that the statement of Lemma 4 is a special case (namely, if $|\mathcal{R}| = 1$) of the first statement of Lemma 2.

Lemma 4: **[3]** Let X and Q be random variables, and let $s > 0$. Then with probability at least $1 - 2^{-(s/2-1)}$ (taken over $q \in \mathcal{Q}$), we have

$$H_2(X) - H_2(X|Q = q) \leq \log |\mathcal{Q}| + s \ .$$

D. Protocol UH

We are now ready to give a first protocol for privacy amplification secure against active adversaries. The ingredients of this protocol are universal hashing (for privacy amplification), strongly universal hashing (for the authentication of the message, i.e., the random bits determining the hash function), and the challenge-response scheme of Section II-B.

For parameters n and ℓ, where 3 divides n and ℓ divides $2n/3$, Protocol UH is defined as follows. (Here, as well as in Protocol EX, the `reject` states are the default states, and are valid initially and until "`accept`" appears in the protocol specification.)

Protocol UH (Universal Hashing)

Alice		**Bob**
$S = S_I \| S_{II} \| S_{III}$		$S = S_I \| S_{II} \| S_{III}$
$S_I \| S_{II} =$		$S_I \| S_{II} =$
$S_0 \| S_1 \| \cdots \| S_{2n/3\ell-1}$		$S_0 \| S_1 \| \cdots \| S_{2n/3\ell-1}$
$h \in_R GF(2^{n/3})$		
$S'_A := \mathrm{LSB}_r(h \cdot S_{III})$		
$a := h \cdot S_I + S_{II}$		
$u \in_R GF(2^\ell)$		
	$\xrightarrow{(h,a,u)}$	
		`accept` if
		$a = h \cdot S_I + S_{II}$
		`stop` otherwise
		$S'_B := \mathrm{LSB}_r(h \cdot S_{III})$
		$v := f_u(S_I \| S_{II})$
	$\xleftarrow{v}$	
`accept` if		
$v = f_u(S_I \| S_{II})$		

Here, S_I, S_{II}, and S_{III} are $(n/3)$-bit strings, whereas $S_0, \ldots, S_{N/\ell-1}$ are ℓ-bit strings. Recall that $h \in_R GF(2^{n/3})$ means that h is chosen randomly from $GF(2^{n/3})$ according to the uniform distribution. All the computations are carried out in the fields $GF(2^{n/3})$ and

$GF(2^\ell)$, respectively.

Theorem 5 states that Protocol UH allows for privacy amplification secure against active adversaries whenever the Rényi entropy, from Eve's point of view, of S is greater than two thirds of the length of S. Moreover, the length of the resulting secret key S' can be roughly equal to the excess, i.e., to

$$H_2(S) - (2/3) \cdot \mathrm{len}(S)$$

(see Figure 3).

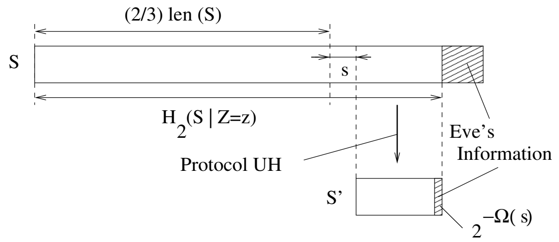

Fig. 3. Analysis of Protocol UH

Theorem 5: Let n, t, ℓ, and s be positive integers such that 3 divides n, ℓ divides $2n/3$, and $n > tn > 2n/3 + s$ holds. Then Protocol UH is a robust

$$(n\,,\, \mathcal{D}_{n,2,tn}\,,\, (t-2/3)n - s\,,\, \varepsilon\,,\, \delta)\text{-PA-protocol}$$

for

$$\begin{aligned} \varepsilon &= r \cdot 2^{-(s/3-1)} + 2^{-s/3}/\ln 2\,, \\ \delta &= 2^{-(t-2/3)n/2} + 3 \cdot 2^{-(t-2/3)n/4} \\ &\qquad + 3 \cdot 2^{-(3\ell/4n)(1-3\ell/2n)(t-2/3)n} + \frac{2n}{3\ell 2^\ell}\,. \end{aligned}$$

Proof. Let $z \in \mathcal{Z}$ be the particular value known to Eve. We first assume that Eve is a *passive* wire-tapper. Let $(h, a) = (h, h \cdot S_I + S_{II})$ be the message sent from Alice to Bob, and let $\mathcal{E}$ be the event that

$$\begin{aligned} &H_2(S_{III} \,|\, S_I = s_I, S_{II} = s_{II}, Z = z) \\ &\qquad \geq (t - 2/3)n - 2s/3 \end{aligned} \tag{6}$$

holds. According to Lemma 2, the event $\mathcal{E}$ has probability at least $1 - 2^{-(s/3-1)}$. Let $r := (t - 2/3)n - s$, and let $S' := \mathrm{LSB}_r(h \cdot S_{III})$. Because of (6), Theorem 4 implies that

$$\begin{aligned} &H(S' \,|\, HA, \mathcal{E}, Z = z) \\ &\quad \geq H(S' \,|\, HAS_IS_{II}, \mathcal{E}, Z = z) \\ &\quad = H(S' \,|\, HS_IS_{II}, \mathcal{E}, Z = z) \\ &\quad \geq r - \frac{2^{-s/3}}{\ln 2}\,. \end{aligned}$$

We have used that $I(S_{III}; HA|S_IS_{II}, Z = z) = 0$ holds. We conclude that

$$\begin{aligned} &H(S'|HA, Z = z) \\ &\quad \geq \mathrm{P}[\mathcal{E}] \cdot \left(r - \frac{2^{-s/3}}{\ln 2} \right) \\ &\quad \geq r - r \cdot 2^{-(s/3-1)} - \frac{2^{-s/3}}{\ln 2} \\ &\quad =: r - \varepsilon\,. \end{aligned}$$

Let us now consider the case where Eve is an *active* attacker. First, Lemma 1 implies that

$$H_2(S_IS_{II} \,|\, Z = z) \geq (t - 1/3)n = n/3 + (t - 2/3)n\,.$$

Therefore, by Lemma 3, the probability of a successful active attack of the message authentication with strongly universal hashing is upper bounded by

$$2^{-(t-2/3)n/2} + 3 \cdot 2^{-(t-2/3)n/4}\,.$$

On the other hand, we have to give an upper bound on the probability that Eve correctly guesses $v = f_u(s_I||s_{II})$. As above, we conclude first that

$$H_2((S_I S_{II}) \mid H = h, A = a, Z = z) \geq (1 - 3\ell/2n)(t - 2/3)n \tag{7}$$

holds with probability at least $2^{-(3\ell/4n)(t-2/3)n+1}$. If (7) holds, then by Lemma 3, the probability of correctly guessing v is at most

$$2^{-(3\ell/4n)(t-2/3)n(1-3\ell/2n)} + 2n/(3\ell 2^\ell) \ .$$

Hence, by the union bound, the success probability of an active attack is upper bounded by

$$2^{-(t-2/3)n/2} + 3 \cdot 2^{-(t-2/3)n/4} + 3 \cdot 2^{-(3\ell/4n)(1-3\ell/2n)(t-2/3)n} + \frac{2n}{3\ell 2^\ell} \ .$$

□

Corollary 6 is an asymptotic version of Theorem 5 and follows directly from the latter.

Corollary 6: Let $2/3 < t < 1$ and $\Delta > 0$ be constants. Then Protocol UH is, for sufficiently large n and for an appropriate choice of the parameters, a robust $(n\,,\,\mathcal{D}_{n,2,tn}\,,\ (t - 2/3 - \Delta)n\,,\,2^{-\Omega(n)}\,,\,2^{-\Omega(n)})$-PA-protocol.

Note that the divisibility conditions required in Theorem 5 can be satisfied by appending a certain number of 0s at the end of the string. Then Theorem 5 can be applied for an appropriate choice of the parameters s and ℓ, both of order $\Theta(n)$, where $s \leq \Delta \cdot n$ holds.

III. Protocol EX Based on Extractors

One limitation of Protocol UH is due to the fact that the message to be transmitted and authenticated, i.e., the description of the function from the universal class, is as long as the string that finally forms the input to the hashing. As described in Section III-B, there exist, however, methods for privacy amplification or, more generally, for "distribution uniformizing," that are more efficient than universal hashing with respect to the required amount of random (message) bits, namely so-called *extractors.*

A. Message Authentication with a Partially Secret Key II: Short Messages and the Power of Feedback

The use of extractors for privacy amplification will allow for reducing the size of the message to be transmitted (and hence authenticated) to a small constant fraction of the length of the authentication key. In this case, a challenge-response authentication method, similar to the method described in Section II-B, can be used: The message is not authenticated by the sender, but reconfirmed by the receiver. Intuitively, the use of such an authentication method puts the adversary, who must answer a given challenge instead of authenticate an incorrect message m' of his own choice, into a much less comfortable position.

Lemma 5 states that the interactive authentication method is secure against an active attacker whose Rényi entropy exceeds half the length of the authentication key. Note that the new authentication method has an important advantage as compared to strongly universal hashing in the context of authentication with a partially secret key. When using the latter method, a correct message-authenticator pair reveals, roughly speaking, half the information about the key (namely a linear equation, two of which are sufficient to determine the key). Hence the key is "used up" after one application. With the interactive method, however, only an arbitrarily small constant fraction of information about the key is gained by an opponent observing a message (i.e., a challenge) together with its authenticator (i.e., the response). This implies that the same key can be used for secure authentication with a partially secret key many times.

Lemma 5: Let N and ℓ be integers such that 2ℓ divides N and $2^\ell \geq N/\ell$ holds, and let K be a random variable with range $\mathcal{K} \subseteq GF(2^N)$. Let further for any $d \in GF(2^\ell)$ the function $f_d : \{0,1\}^N \to \{0,1\}^\ell$ be defined as in Lemma 3. Assume that there exists a (possibly probabilistic) function s, mapping $GF(2^\ell)$ to $GF(2^\ell)$,

$$s : d \mapsto s(d) =: d' ,$$

such that $d' \neq d$ holds for all d, and such that given $f_{d'}(K)$, the value $f_d(K)$ can, for $d \in_R GF(2^\ell)$, be guessed correctly (with some strategy) with probability α, taken over the distribution of K, the choice of d, and the coin tosses of the guessing strategy. Then

$$H_2(K) \leq \frac{N}{2} - \frac{2N}{\ell} \cdot \log\left(\alpha - \frac{N/\ell}{2^\ell}\right)$$

or, equivalently,

$$\alpha \leq 2^{-(\ell/2N)\cdot(H_2(K)-N/2)} + \frac{N/\ell}{2^\ell}$$

holds.

Proof. Note first that, by the same arguments as used in the proof of Lemma 3, we can assume without loss of generality that the function $d'(d)$ and the strategy of guessing $f_d(k)$ from $f_{d'}(k)$ are deterministic, and conclude that there exist distinct elements $d_1, \ldots, d_{N/\ell}$ of $GF(2^\ell)$ such that $f_{d_i}(k)$ is guessed correctly from $f_{d'_i}(k)$, where $d'_i := d'(d_i)$, for all $i = 1, \ldots, N/\ell$ simultaneously with probability at least

$$\left(\alpha - \frac{N/\ell}{2^\ell}\right)^{N/\ell} .$$

Let $\mathcal{E}$ ($\subseteq \mathcal{K}$) be this event. We prove that $|\mathcal{E}| \leq \sqrt{|\mathcal{K}|}$.

By canceling $N/2\ell$ of the pairs (d_i, d'_i) and renumbering the remaining pairs, we can obtain $N/2\ell$ pairs (d_i, d'_i) with the property that $d_i \notin$

$\{d'_1, \ldots, d'_{i-1}\}$ holds for all $i = 1, \ldots, N/2\ell$. (In the worst case, all the pairs (d_i, d'_i) occur twice in different orderings. Then, every second pair (d_i, d'_i) must be canceled.)

The event $\mathcal{E}$ has the property that

$$f_{d'_1}(k) = f_{d'_1}(k^*) \implies f_{d_1}(k) = f_{d_1}(k^*)$$

for all $k, k^* \in \mathcal{E}$. Otherwise $f_{d_1}(k)$ could not be guessed correctly from $f_{d'_1}(k)$ for all $k \in \mathcal{E}$. Hence $\mathcal{E}$ must be contained in a set $\mathcal{E}_1$ of the form

$$\mathcal{E}_1 = \bigcup_{a \in GF(2^\ell)} \{k \,:\, f_{d_1}(k) = b(a) \text{ and } f_{d'_1}(k) = a\}$$

for some function $b(a)$. Analogously, $\mathcal{E}$ must also be contained in sets $\mathcal{E}_i$, $i = 2, \ldots, N/2\ell$, of the same form (with d_1 and d'_1 replaced by d_i and d'_i, respectively), hence

$$\mathcal{E} \subseteq \bigcap_{i=1}^{N/2\ell} \mathcal{E}_i \ . \tag{8}$$

We show that the cardinality of the set on the right-hand side of (8) is $\sqrt{|\mathcal{K}|}$. First, observe that every set of at most $N/\ell \,(\leq 2^\ell)$ functions f_{d_i} is, for pairwise distinct $d_i \in GF(2^\ell)$, linearly independent over $GF(2^\ell)$ (the Vandermonde determinant is non-zero in this case, as shown in the proof of Lemma 3). We define

$$r_l := \left| \bigcap_{i=1}^{l} \mathcal{E}_i \right| \ .$$

From the linear independence of $\{f_{d_1}, f_{d'_1}\}$, we first conclude that $r_1 = 2^{N-\ell}$. Furthermore, the linear independence of $f_{d_{l+1}}$ from the set

$$\{f_{d_1}, \ldots, f_{d_l}, f_{d'_1}, \ldots, f_{d'_{l+1}}\}$$

(because $d_{l+1} \notin \{d_1, \ldots, d_l, d'_1, \ldots, d'_{l+1}\}$ according to the choice of the pairs (d_i, d'_i)) implies

$$r_{l+1} = r_l / 2^\ell$$

for $l = 1, \ldots, N/2\ell - 1$. Note that this also holds if $d'_{l+1} = d_i$ or $d'_{l+1} = d'_i$ for some $i < l + 1$. We conclude that

$$|\mathcal{E}| \leq r_{N/2\ell} = 2^{N-(N/2\ell)\cdot\ell} = 2^{N/2} = \sqrt{|\mathcal{K}|} \ .$$

On the other hand, $\mathrm{P}[\mathcal{E}] = \sum_{k \in \mathcal{E}} P_K(k) \geq (\alpha - (N/\ell)/2^\ell)^{N/\ell}$ holds. In the case where P_K restricted to $\mathcal{E}$ is the uniform distribution (this case maximizes the Rényi entropy) with probability $(\alpha - (N/\ell)/2^\ell)^{N/\ell}/|\mathcal{E}|$ or greater, we have

$$\begin{aligned} \sum_{k \in \mathcal{K}} P_K(k)^2 &\geq \sum_{k \in \mathcal{E}} P_K(k)^2 \\ &\geq |\mathcal{E}| \cdot \frac{(\alpha - (N/\ell)/2^\ell)^{2N/\ell}}{|\mathcal{E}|^2} \\ &\geq \frac{(\alpha - (N/\ell)/2^\ell)^{2N/\ell}}{2^{N/2}} \ , \end{aligned}$$

and the claim follows when the negative logarithm is computed on both sides. □

B. Privacy Amplification with Extractors

For constructing a protocol allowing privacy amplification with shorter messages, we use a different technique for privacy amplification, based on *extractors*.

Roughly speaking, an extractor allows for efficiently isolating the randomness of some source into virtually random bits, using a small additional number of random bits as a catalyst, i.e., in such a way that these bits reappear as a part of the almost uniformly distributed output. Extractors are of great importance in theoretical computer science, where randomness is often regarded as a resource. They have been studied intensively in the past years by many authors. For an introduction to the subject and some constructions, see for example [17] or [18], and the references therein.

Recent results, described below, show that such functions allow, using only a small amount of true randomness, to distill (almost) the entire randomness, measured in terms of H_∞, of some string into an almost uniformly distributed string. A disadvantage of using extractors instead of universal hashing is that a string of length only roughly equal to the *min*-entropy instead of the generally greater *Rényi* entropy of the original random variable can be extracted. However, this drawback has no effect in connection with typical sequences, i.e., almost uniform distributions. (Note that for uniform distributions, all the introduced entropy measures are equal.)

Definition 4: A function $E : \{0,1\}^N \times \{0,1\}^d \to \{0,1\}^r$ is called a (δ', ε')-extractor if for any distribution P on $\{0,1\}^N$ with min-entropy $H_\infty(P) \geq \delta' N$, the variational distance[3] of the distribution of

$$[V, E(X, V)]$$

to the uniform distribution over $\{0,1\}^{d+r}$ is at most ε' when choosing X according to P and V independently according to the uniform distribution over $\{0,1\}^d$.

The following theorem was proved in [18]. It states that there exist extractors which distill virtually all the min-entropy out of a weakly random source, thereby requiring only a small (i.e., "poly-logarithmic") number of truly random bits. Note that Theorem 7 is formally slightly stronger than the corresponding theorem in [18] because it not only states that the length of the extractor output is roughly equal to the min-entropy of the source plus the number of random bits, but even that these bits reappear as a part of the output. Although it has not been explicitly stated, it is not difficult to see that the extractors described in [18] do have this property.

Theorem 7: **[18]** For every choice of the parameters N, $0 < \delta' < 1$, and $\varepsilon' > 0$, there exists a (δ', ε')-extractor

$$E : \{0,1\}^N \times \{0,1\}^d \longrightarrow \{0,1\}^{\delta' N - 2\log(1/\varepsilon') - O(1)} ,$$

where $d = O((\log(N/\varepsilon'))^2 \log(\delta' N))$.

[3]The *variational distance* of two distributions P_X and P_Y over the same range $\mathcal{X}$ is defined as $(\sum_{x \in \mathcal{X}} |P_X(x) - P_Y(x)|)/2$.

Corollary 8, which is a consequence of Theorem 7, is what we need for the analysis of Protocol EX. The statement of Corollary 8 is related to Theorem 4, where universal hashing is replaced by extractors, and min-entropy must be used instead of Rényi entropy (see Figure 4).

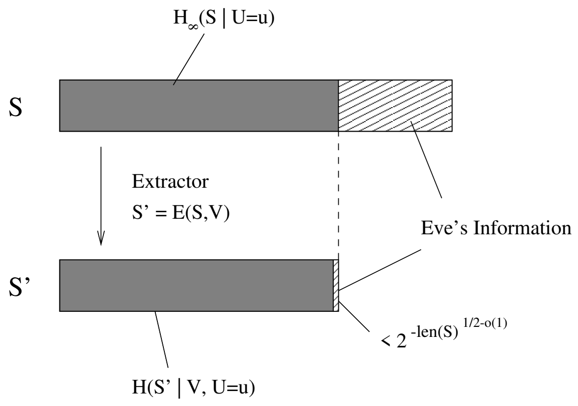

Fig. 4. Privacy Amplification with Extractors

Corollary 8: Let $\delta', \Delta_1, \Delta_2 > 0$ be constants. Then there exists, for all sufficiently large N, a function

$$E : \{0,1\}^N \times \{0,1\}^d \longrightarrow \{0,1\}^r ,$$

where $d \leq \Delta_1 N$ and $r \geq (\delta' - \Delta_2)N$, such that for all random variables T with $\mathcal{T} \subseteq \{0,1\}^N$ and

$$H_\infty(T) > \delta' N ,$$

we have

$$H(E(T,V)|V) \geq r - 2^{-N^{1/2-o(1)}} . \quad (9)$$

Lemma 6: Let Z be a random variable with range $\mathcal{Z} \subseteq \{0,1\}^k$. Then

$$H(Z) \geq k \cdot (1 - d(U_k, P_Z) - 2^{-k}) \quad (10)$$

holds, where U_k stands for the uniform distribution over $\{0,1\}^k$.

Proof. Let $d := d(U_k, Z)$. We can assume that $d < 1 - 2^{-k}$ holds because otherwise the inequality is trivially satisfied. The distribution P_Z of Z can be thought of as obtained from the uniform distribution U_k by increasing some of the probabilities (by total amount d) and decreasing some others (by the same total amount). The function

$$\frac{\mathrm{d}}{\mathrm{d}p}(-p \log p) = -\frac{\ln p + 1}{\ln 2}$$

is monotonically decreasing, hence increasing (or decreasing) a *smaller* probability increases (or decreases, respectively) the entropy more than modifying a greater probability by the same amount. Hence a distribution with distance d from U_k with minimal entropy can be obtained by adding d to one of the probabilities, and by reducing as many probabilities as possible to 0, leaving the other probabilities unchanged. One of the probabilities of the new distribution equals $2^{-k} + d$, $\lfloor 2^k d \rfloor$ probabilities are equal to 0, one probability equals $2^{-k}(2^k d - \lfloor 2^k d \rfloor)$ (if this is not 0), and $\lfloor 2^k(1-d) \rfloor - 1$ probabilities are unchanged and hence equal to 2^{-k}. Thus the entropy of the new random variable Z can be bounded from below by

$$\begin{aligned} H(Z) &\geq 2^{-k}(2^k d - \lfloor 2^k d \rfloor) \cdot k \\ &\quad + (\lfloor 2^k (1-d) \rfloor - 1) \cdot 2^{-k} \cdot k \\ &= k \cdot (1 - d - 2^{-k}) \ . \end{aligned}$$

□

Proof of Corollary 8. Let $\varepsilon'(N) := 2^{-\sqrt{N}/\log N}$. Then there exists N_0 such that for all $N \geq N_0$ we have a (δ', ε')-extractor E, mapping $\{0,1\}^{N+d}$ to $\{0,1\}^r$, where $d \leq \Delta_1 N$ (note that $d = O(N/\log N)$ holds for this choice of ε') and $r \geq (\delta' - \Delta_2)N$. By definition, this means that for a uniformly distributed d-bit string V and if $H_\infty(T) \geq \delta' N$, the distance of the distribution of $[V, E(T,V)]$ to the uniform distribution U_{d+r} over $\{0,1\}^{d+r}$ is at most $\varepsilon' = 2^{-\sqrt{N}/\log N}$. Because

$$d([V, E(T,V)], U_{d+r}) = \mathrm{E}_V[d(E(T,V), U_r)] \leq \varepsilon'$$

holds for uniformly distributed V, the distance of the distribution of $E(T,v)$ to the uniform distribution U_r (over $\{0,1\}^r$) is at most $\sqrt{\varepsilon'}$ with probability at least $1 - \sqrt{\varepsilon'}$ over v, i.e.,

$$P_V\left[d\left(E(T,V), U_r\right) \leq 2^{-\sqrt{N}/2\log N}\right] \geq 1 - 2^{-\sqrt{N}/2\log N} \ .$$

Inequality (9) now follows from Lemma 6. □

In analogy to Lemma 4, which gives an upper bound on the effect of side information on the Rényi entropy of a random variable (hereby linking information reconciliation and privacy amplification with universal hashing) we now need such a result with respect to min-entropy H_∞. Lemma 7 is an immediate consequence of Lemma 2.

Lemma 7: Let X and Q be random variables, and let $s > 0$. Then with probability at least $1 - 2^{-s}$ (taken over $q \in \mathcal{Q}$), we have

$$H_\infty(X) - H_\infty(X|Q=q) \leq \log|\mathcal{Q}| + s \ .$$

C. Protocol EX

In this section we present Protocol EX for privacy amplification secure against active adversaries. This protocol uses an extractor function for privacy amplification and the interactive challenge-response authentication methods. One important difference to Protocol UH is that a shorter string is used as the authentication key. This allows for extracting a significantly longer string in the case where Eve's information about the original string is small.

Let the binary string S (of length n) be composed by strings S_I and S_{II} of lengths n_1 and n_2, respectively. Assume that 2ℓ divides n_1 and let $d := n_1/\ell$. The substrings S_i, $i = 0, \ldots, d-1$, of S_I are all of length ℓ. The function E is an extractor to be specified below.

Protocol EX (Extractor)

Alice		Bob
$S = S_I \| S_{II}$		$S = S_I \| S_{II}$
$S_I = S_0 \| \cdots \| S_{d-1}$		$S_I = S_0 \| \cdots \| S_{d-1}$
$h \in_R GF(2^\ell)$		
	$\xrightarrow{h}$	
		$a := f_h(S_I)$
		$b \in_R GF(2^\ell)$
	$\xleftarrow{(a,b)}$	
if $a \neq f_h(S_I)$:		$S'_B := E(S_{II}, h)$
`stop`		
if $a = f_h(S_I)$:		
$c := f_b(S_I)$		
	$\xrightarrow{c}$	
`accept`		if $c = f_b(S_I)$: `accept`
$S'_A := E(S_{II}, h)$		if $c \neq f_b(S_I)$: `reject`

Theorem 9 below implies that Protocol EX can be much more effective that Protocol UH, in particular for strings with a high level of initial security. Note first that Protocol EX works under the same condition as Protocol UH: The Rényi entropy of the string, given Eve's knowledge, must be larger than two thirds of the length of the string. The length of the extractable key however can be equal to roughly

$$H_\infty(S) - 2(\mathrm{len}(S) - H_2(S))$$

instead of only the excess $H_2(S) - (2/3)\mathrm{len}(S)$ as for Protocol UH. This expression can be substantially greater, in particular if $H_2(S)$ is close to $\mathrm{len}(S)$. On the other hand, since $H_\infty(S)$ can be smaller (by a factor up to 2) than $H_2(S)$, Protocol EX can also be *less* effective than Protocol UH. An illustration of the statement of Theorem 9 is given in Figure 5.

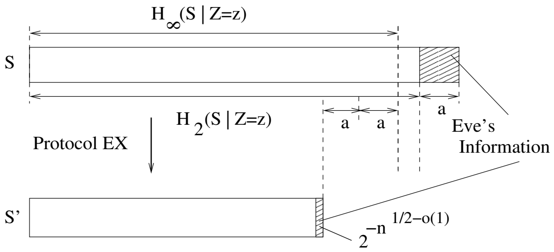

Fig. 5. Analysis of Protocol EX

Theorem 9: Let $0 < t' \leq t < 2/3$ and $\Delta > 0$ be constants. Then Protocol EX is, for sufficiently large n and for an appropriate choice of the parameters, a robust $(n\,,\, \mathcal{D}_{n,2,tn} \cap \mathcal{D}_{n,\infty,t'n}\,,\ (t' - 2(1-t) - \Delta)n\,,\, 2^{-n^{1/2-o(1)}}\,,\, 2^{-\Omega(n)})$-PA-protocol.

Proof. Let $n_1 := (2(1-t) + \Delta/3)n$, $\ell := (\Delta/6)n$, and let $z \in \mathcal{Z}$ be the particular value known to Eve. We can assume without loss of generality that n_1 and ℓ are integers and that 2ℓ divides n_1. (Otherwise, n_1 and ℓ can both be chosen smaller, subject to $n_1 = (2(1-t) + \Theta(1))n$ and $l = \Theta(n)$, respectively, such that the conditions are satisfied.) Let now S_I be the first n_1 and S_{II} be the remaining $n_2 := n - n_1$ bits of S.

Assume first that Eve is *passive*. We give a lower bound on the min-entropy of the string

S_{II} from Eve's point of view and given the entire communication C held over the public channel. Since this communication is, given S_I and $Z = z$, independent of S_{II}, we have

$$\begin{aligned} H_\infty(S_{II} \mid C = c, S_I = s_I, Z = z) \\ = H_\infty(S_{II} \mid S_I = s_I, Z = z) \\ \geq (t' - 2(1-t) - 2\Delta/3)n \quad (11) \end{aligned}$$

with probability $1 - 2^{-\Omega(n)}$. (Note that Alice and Bob could publish S_I at the end of the protocol, only helping a possible adversary.)

Let now E be the extractor function according to Corollary 8 with $d \leq \ell = \Theta(n)$,

$$\delta' = (t' - 2(1-t) - 2\Delta/3)n/n_2 \ ,$$

and

$$r \geq (\delta' - \Delta \cdot n/(3n_2))n_2 = (t' - 2(1-t) - \Delta)n \ .$$

For the choice $P_T = P_{S_{II}|C=c,Z=z}$ and $S' = E(S_{II}, V)$ (where V is composed by the first d bits of H in a fixed representation) we obtain, using (11) and $I(H; SZ) = 0$,

$$H(S' \mid C, Z = z) \geq r - 2^{-N^{1/2-o(1)}} \ .$$

We consider the case where Eve is an *active* adversary and give an upper bound on the probability of the event that Alice and Bob do not both reject and secret-key agreement has nevertheless not been successful. It is obvious that this can only occur if Eve can either guess $f_h(S)$ from some $f_{h'}(S)$ (where $h' \neq h$) or guess $f_b(S)$ correctly, where h and b are randomly chosen. The success probability δ of such an active attack is of order

$$\delta = 2^{-\Omega(n)} \ . \quad (12)$$

To see this, note first that

$$H_2(S_I \mid Z = z) \geq n_2/2 + \ell + \Delta/6$$

holds because of Lemma 1 and by the definitions of n_2 and ℓ. Then, the probabilities of the events that $f_h(S_I)$ is guessed correctly from $f_{h'}(S_I)$ (note that $\ell = \Theta(n_2)$), that

$$H_2(S_I \mid f_h(S_I) = f_h(s_I), Z = z) < n_2/2$$

holds (we call this event $\mathcal{E}$), and that $f_b(S_I)$ is guessed correctly, given that $\mathcal{E}$ does not occur, are all of order $2^{-\Omega(n)}$, because of Lemmas 5, 2, and 3, respectively. The bound (12) then follows from the union bound. □

Remark. Note that for the proof of Theorem 9, the following combination of the Protocols UH and EX is also sufficient: The key is partitioned as in Protocol EX, but the authentication techniques of Protocol UH are used for the shorter key, i.e., strongly universal hashing and challenge-response confirmation (Section II-B). However, Protocol EX has an important advantage as compared to this protocol, and to Protocol UH, which is not stated explicitly in Theorem 9: In case of failure because of a detected active substitution attack, Protocol EX can (with roughly the same parameters) be restarted again and again ($O(1)$ times) with the same key until secret-key agreement eventually succeeds. The reason for the possibility of such multiple trials is that the observation of a correctly authenticated message reveals—unlike in the case of authentication with strongly universal hashing—only a small fraction of the total information about the authentication key (see Section III-A).

IV. Discussion

We have described two protocols, Protocol UH and Protocol EX, for privacy amplification secure in the active-adversary model. Protocol UH is based on universal hashing and is successful as soon as the Rényi entropy of the partially secret key, from the adversary's viewpoint, exceeds two thirds of the length of the string. Then the protocol distills a string whose length is roughly equal to this excess. Protocol UH is computationally extremely simple and works for strings of any length. For sufficiently long strings, Protocol EX, based on extractor functions, can be used. The condition on the initial key is the same as for Protocol UH, but the extracted highly secret key is often longer. An additional advantage of Protocol EX is that failed privacy-amplification attempts do (almost) not use up the key: The procedure can be repeated many times with the same key.

In Figure 6, the required conditions on the partially secret key as well as the possible length of the resulting secret key when using Protocol UH or EX are illustrated and compared to the corresponding quantities in the case of privacy amplification (according to [1]) in the passive-adversary case. This representation therefore shows the price that must be paid for authentication in the context of privacy amplification.

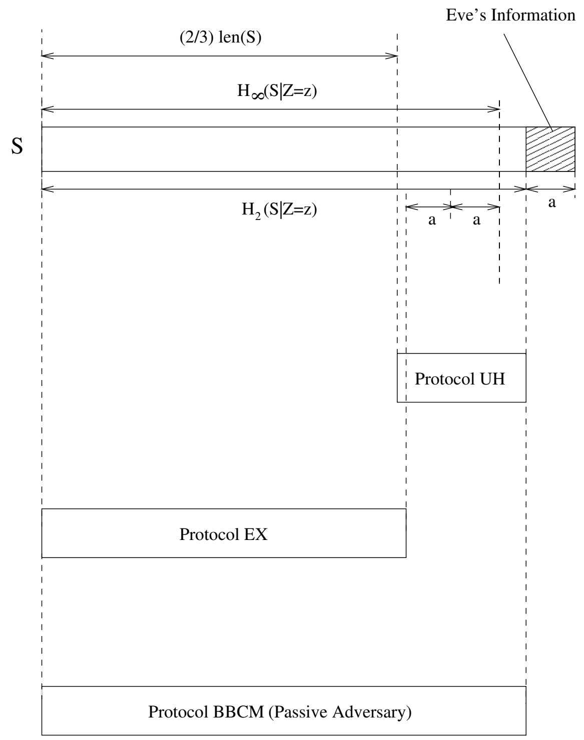

Fig. 6. Privacy Amplification Secure Against Active Adversaries

V. Concluding Remarks

In the general setting of secret-key agreement from correlated randomness by insecure and not authentic communication, we have analyzed the important special case of privacy amplification. Different problems arise here as compared to the independent-realizations model considered in [15]. Examples are the need for authentication with a partially secret key or hashing with an only small amount of joint randomness.

Our results are based on the combination of new message-authentication methods—that require interaction but only a possibly highly insecure key, and that can be used repeatedly with the same key—and a new technique for privacy amplification.

In analogy to the scenario where a random experiment is repeated many times [15], we found that privacy amplification secure against active adversaries is achievable, but only under certain conditions stronger than the ones for the passive-adversary case. For privacy amplification, however, we have not shown these conditions to be necessary, and we state as an open problem to prove or disprove their necessity.

In contrast to the independent-repetitions scenario, a certain price has to be paid for the channel's missing authenticity even if robust privacy amplification is possible in principle: the generated key is shorter. It is a challenging open problem to find protocols extracting the same amount of secrecy in the presence of active adversaries as is possible against only passive wire-tappers (or to prove that such protocols cannot exist).

References

[1] 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.

[2] C. H. Bennett, G. Brassard, and J.-M. Robert, Privacy amplification by public discussion, *SIAM Journal on Computing*, Vol. 17, pp. 210–229, 1988.

[3] C. Cachin, *Entropy measures and unconditional security in cryptography*, Ph. D. Thesis, ETH Zürich, Hartung-Gorre Verlag, Konstanz, 1997.

[4] C. Cachin and U. M. Maurer, Linking information reconciliation and privacy amplification, *Journal of Cryptology*, Vol. 10, No. 2, pp. 97–110, Springer-Verlag, 1997.

[5] J. L. Carter and M. N. Wegman, Universal classes of hash functions, *Journal of Computer and System Sciences*, Vol. 18, pp. 143–154, 1979.

[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*, Vol. 24, No. 3, pp. 339–348, 1978.

[8] P. Gemmell and M. Naor, Codes for interactive authentication *Advances in Cryptology - CRYPTO '93*, Lecture Notes in Computer Science, Vol. 773, Springer–Verlag, pp. 355–367, 1994.

[9] 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.

[10] 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.

[11] U. M. Maurer, Authentication theory and hypothesis testing, *IEEE Transactions on Information Theory*, Vol. 46, No. 4, pp. 1350–1356, 2000.

[12] 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.

[13] U. M. Maurer and S. Wolf, Unconditionally secure key agreement and the intrinsic conditional information, *IEEE Transactions on Information Theory*, Vol. 45, No. 2, pp. 499–514, 1999.

[14] U. M. Maurer and S. Wolf, Information-theoretic key agreement: from weak to strong secrecy for free, *Advances in Cryptology - EUROCRYPT 2000*, Lecture Notes in Computer Science, Vol. 1807, pp. 351–368, Springer-Verlag, 2000.

[15] U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part I: Definitions and a completeness result, *IEEE Transactions on Information Theory*, 2003.

[16] U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part II: The simulatability condition, *IEEE Transactions on Information Theory*, 2003.

[17] N. Nisan and D. Zuckerman, Randomness is linear in space, *Journal of Computer and System Sciences*, Vol. 52, No. 1, pp. 43–52, 1996.

[18] R. Raz, O. Reingold, and S. Vadhan, Extracting all the randomness and reducing the error in Trevisan's extractors, *Proceedings of STOC'99*, 1999.

[19] C. E. Shannon, Communication theory of secrecy systems, *Bell System Technical Journal*, Vol. 28, pp. 656–715, 1949.

[20] G. J. Simmons, A survey of information authentication, *Proceedings of the IEEE*, Vol. 76, pp. 603–620, 1988.

[21] 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.

[22] S. Wolf, Strong security against active attacks in information-theoretic secret-key agreement, *Advances in Cryptology - ASIACRYPT '98*, Lecture Notes in Computer Science, Vol. 1514, pp. 405–419, Springer-Verlag, 1998.

[23] S. Wolf, *Information-theoretically and computationally secure key agreement in cryptography*, ETH dissertation No. 13138, Swiss Federal Institute of Technology (ETH Zurich), 1999.

[24] A. D. Wyner, The wire-tap channel, *Bell System Technical Journal*, Vol. 54, No. 8, pp. 1355–1387, 1975.