Privacy Amplification
Secure Against Active Adversaries

Ueli Maurer Stefan Wolf

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

Abstract. Privacy amplification allows two parties Alice and Bob knowing a partially secret string S to extract, by communication over a public channel, a shorter, highly secret string S'. Bennett, Brassard, Crépeau, and Maurer showed that the length of S' can be almost equal to the conditional Rényi entropy of S given an opponent Eve's knowledge. All previous results on privacy amplification assumed that Eve has access to the public channel but is passive or, equivalently, that messages inserted by Eve can be detected by Alice and Bob. In this paper we consider privacy amplification secure even against active opponents. First it is analyzed under what conditions information-theoretically secure authentication is possible even though the common key is only partially secret. This result is used to prove that privacy amplification can be secure against an active opponent and that the size of S' can be almost equal to Eve's min-entropy about S minus $2n/3$ if S is an n-bit string. Moreover, it is shown that for sufficiently large n privacy amplification is possible when Eve's min-entropy about S exceeds only $n/2$ rather than $2n/3$.

Keywords: Privacy amplification, Secret-key agreement, Unconditional secrecy, Authentication codes, Information theory, Extractors.

1 Introduction and Preliminaries

Privacy amplification introduced by Bennett *et. al.* [2] is a technique for transforming a string that is only partially secret into a highly secret (but generally shorter) string. More precisely, two parties Alice and Bob who share a string S about which an opponent Eve has partial information agree, by communication over an insecure channel, on a string S' such that Eve's information about S' is negligible, i.e., such that $H(S'|U = u) \geq \log |\mathcal{S}'| - \varepsilon$ holds with very high probability for some small $\varepsilon > 0$, where the random variable U summarizes Eve's complete knowledge about S', and where u is the particular value known to Eve. (All the logarithms in this paper are to the base 2, unless otherwise stated.) Privacy amplification is an important sub-protocol in many information-theoretic protocols such as protocols in quantum cryptography and secret-key agreement by public discussion [8].

Before we formalize the main problem considered in this paper, we give some definitions and state previous results on privacy amplification.

1.1 Entropy Measures

We recall the definitions of some entropy measures we need in this paper. We assume that the reader is familiar with the basic information-theoretic concepts. For a good introduction, we refer to [4]. Let R be a discrete random variable with range $\mathcal{R}$. Then the *(Shannon) entropy* $H(R)$ is defined as

$$H(R) := -\sum_{r \in \mathcal{R}} P_R(r) \cdot \log(P_R(r)) \ .$$

The *Rényi entropy* $H_2(R)$ is defined as

$$H_2(R) := -\log \left(\sum_{r \in \mathcal{R}} P_R^2(r) \right) \ .$$

Finally, the *min-entropy* $H_\infty(R)$ is

$$H_\infty(R) := -\log \max_{r \in \mathcal{R}} (P_R(r)) \ .$$

It is not difficult to see that for any random variable R the entropy measures H, H_2, and H_∞ satisfy

$$H(R) \geq H_2(R) \geq H_\infty(R) \geq H_2(R)/2 \ .$$

Equality of the first three expressions holds if and only if R is uniformly distributed over some set, in which case this value is the logarithm of the cardinality of this set.

1.2 Universal and Strongly Universal Hashing

In the technique presented in this paper, hashing is used for two different purposes: universal hashing for privacy amplification and strongly universal hashing for authentication.

Definition 1. A class $\mathcal{F}$ of functions $\mathcal{A} \rightarrow \mathcal{B}$ is called *universal*$_2$ (or simply *universal*) if, for any x_1, x_2 in $\mathcal{A}$ with $x_1 \neq x_2$, the probability that $f(x_1) = f(x_2)$ is at most $1/|\mathcal{B}|$ when f is chosen from $\mathcal{F}$ according to the uniform distribution.

The following is a well-known example of such a class of hash functions $\{0,1\}^n \rightarrow \{0,1\}^r$ containing 2^n distinct functions. Let $b \in GF(2^n)$, and interpret $x \in \mathcal{A} = \{0,1\}^n$ also as an element of $GF(2^n)$. Consider the function f_b assigning to the argument x the first r bits of the element $b \cdot x$ of $GF(2^n)$. The set of these functions f_b for $b \in GF(2^n)$ is a universal class of functions for $1 \leq r \leq n$.

Definition 2. Let $\varepsilon > 0$. A class $\mathcal{H}$ of (hash) functions $\mathcal{A} \rightarrow \mathcal{B}$ is called *ε-almost-strongly-universal*$_2$ (or ε-ASU$_2$ for short) if the following two conditions are satisfied:

1. For every $a \in \mathcal{A}$ and $b \in \mathcal{B}$, the number of functions $h \in \mathcal{H}$ with $h(a) = b$ is $|\mathcal{H}|/|\mathcal{B}|$.
2. For every distinct $a_1, a_2 \in \mathcal{A}$ and for every $b_1, b_2 \in \mathcal{B}$, the number of hash functions $h \in \mathcal{H}$ for which both $h(a_1) = b_1$ and $h(a_2) = b_2$ hold is at most $\varepsilon \cdot |\mathcal{H}|/|\mathcal{B}|$.

An $(1/|\mathcal{B}|)$-ASU$_2$ class is also called *strongly-universal$_2$* (or SU$_2$).

Some constructions of ε-ASU$_2$ classes are described in [12], and lower bounds on the size of such classes are proved. An SU$_2$ class of functions mapping n-bit strings to n-bit strings can be constructed similarly to the universal class described above: the class $\mathcal{H} = \{h_{ab} : (a, b) \in (GF(2^n))^2\}$, where $h_{ab}(x) := a \cdot x + b$, is an SU$_2$ class of hash functions $\{0,1\}^n \rightarrow \{0,1\}^n$ with 2^{2n} elements. It is shown in [12] that ε-ASU$_2$ classes can be obtained which are close to strongly-universal, but substantially smaller.

1.3 Privacy Amplification by Authenticated Public Discussion

Bennett *et. al.* [1] analyzed the privacy amplification technique of [2] under the assumption that the two parties Alice and Bob are connected by an authentic (but otherwise insecure) channel, or equivalently, that the opponent is not able to insert or modify messages without being detected. The idea of this technique is to take a hash value of the string S as the highly secret key. More precisely, Alice chooses a hash function h at random from a universal class and sends this function to Bob. Then they both compute $S' := h(S)$.

It was shown that the amount of almost secret key that can be extracted is at least equal to the conditional Rényi entropy H_2 of S, given Eve's knowledge $U = u$. This fact is an immediate consequence of the following result of [1] which states that if a random variable X is used as the argument of universal hashing, where the output Y is an r-bit string, and r is equal to $H_2(X)$ minus a security parameter, then the resulting string Y has almost maximal Shannon entropy r, given the hash function (which is chosen uniformly from the universal class).

Theorem 3. [1] *Let X be a random variable with probability distribution P_X and Rényi entropy $H_2(X)$, and let G be the random variable corresponding to the random choice (with uniform distribution) of a member of a universal class of hash functions mapping $\mathcal{X}$ to r-bit strings, and let $Y = G(X)$. Then*

$$r \geq H(Y|G) \geq H_2(Y|G) \geq r - \frac{2^{r-H_2(X)}}{\ln 2} \ .$$

Of course the theorem also holds when all the probabilities are conditioned on a particular event (e.g., $U = u$).

1.4 Privacy Amplification by NOT Authenticated Public Discussion

In this paper we consider the generalized problem of privacy amplification when dropping the condition that the channel connecting the two parties Alice and Bob be authentic, i.e., privacy amplification secure even against active adversaries who are able to insert or modify messages.

Privacy amplification is often used as the final phase of unconditional secret-key agreement. In [6], it was investigated under what conditions secret-key agreement by not authenticated public discussion is possible when the parties Alice, Bob, and Eve have access to random variables X, Y, and Z, respectively (the "initialization phase"). Several impossibility results were shown, whereas a positive result was derived in [6] only for the special case where the information that the parties obtain consists of many independent repetitions of a random experiment. Privacy amplification, which was not treated in [6], corresponds to the situation where $X = Y$, and where the random experiment is not repeated.

We make precise what we mean by a protocol for privacy amplification by communication over a non-authentic insecure channel. Assume that two parties Alice and Bob both know a random variable S, for example an n-bit string, and that the adversary Eve has some information about S. Let again the random variable U summarize Eve's entire information about the random variable S. In the following, all the results are stated for some particular value $u \in \mathcal{U}$ (where $\mathcal{U}$ is the range of the random variable U), i.e., for a fixed event $U = u$, and hence all the probabilities are conditioned on $U = u$. The type of the opponent's information about S is not necessarily precisely specified, i.e., $P_{S|U=u}$ is not assumed to be known. However, the amount of information is limited in some way, for example in terms of the conditional min-entropy.

Formally, a protocol for privacy amplification consists of two phases. During the first phase (the communication phase), Alice and Bob exchange messages $C_1, C_2, \ldots$ over some channel (where Alice sends the messages $C_1, C_3, \ldots$, and Bob sends $C_2, C_4, \ldots$). Each of these messages can depend on the sender's knowledge when sending the message and some random bits. In the second phase, both parties decide whether they accept or reject the outcome of the protocol. In case of acceptance, Alice and Bob compute strings S'_A and S'_B, respectively. (Note that it is not required that Alice and Bob are synchronized in the sense that they both either reject or accept. This would be impossible to achieve in the presence of an active adversary, who could for instance delete all messages from Alice to Bob after Alice has accepted.) Definition 4 defines security of such a protocol.

Definition 4. A protocol is called an $(n, l, n', \varepsilon, \delta)$-*protocol for privacy amplification over an insecure and non-authentic channel* if it is a protocol for privacy amplification with the following properties. If there exists a random variable S with $|\mathcal{S}| \leq 2^n$ (i.e., we can assume that $\mathcal{S} \subset \{0,1\}^n$) that is known to Alice and Bob, and such that given Eve's entire knowledge $U = u$ about S, the conditional min-entropy of S is at least l, i.e.,

$$H_\infty(S|U = u) \geq l \ ,$$

then the protocol satisfies the following conditions. In the case of a *passive* (only wire-tapping) adversary, Alice and Bob always accept at the end of the protocol and obtain a common n'-bit string S' ($= S'_A = S'_B$) such that Eve's knowledge about S' is virtually 0 or, more precisely,

$$H(S'|SC) = 0 \ ,$$

and

$$H(S'|C, U = u) \geq n' - \varepsilon \ , \tag{1}$$

where C summarizes the entire communication $(C_1, C_2, \ldots)$ between Alice and Bob. In the case of an *active* adversary, with probability at least $1 - \delta$ one of the following conditions must be satisfied: either the adversary's presence is detected by at least one of the two parties (who hence rejects), or Alice and Bob both accept and successfully agree on a common string S' ($= S'_A = S'_B$) satisfying (1).

This definition can be generalized to different ways of limiting Eve's knowledge about S, for example in terms of the Rényi entropy instead of the min-entropy.

1.5 Outline

The rest of this paper is organized as follows. In Section 2 we investigate the general problem of information-theoretically secure message authentication under the (weakened) condition that two parties share a partially (rather than completely) secret key. In Section 3 we show a first result concerning privacy amplification. It states that privacy amplification (by communication over a non-authentic channel) is possible if Eve's min-entropy about S exceeds two thirds of the length n of the string, and the maximal length of the generated highly secret string is roughly $H_\infty(S|U = u) - 2n/3$. In Section 4 it is demonstrated that this result is not optimal: it is sufficient that Eve's min-entropy about S is greater than *half* of the length of the string (where the length of the extracted highly secret string is a constant fraction of $H_\infty(S|U = u) - n/2$) if the string is sufficiently long. Section 5 provides evidence that some of the results of Sections 2, 3, and 4 are optimal, and Section 6 states some open problems.

2 Unconditionally-Secure Authentication with a Partially Secret Key

All previous results on unconditionally-secure authentication require a key that is completely secret, i.e., the opponent's a priori probability distribution of the key is uniform. In this section we consider authentication where the opponent is allowed to have some partial information about the key.

There exists a variety of constructive results as well as impossibility results on information-theoretically secure authentication (see for example [11], [7], or [12]). 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 [7].)

One possibility for realizing information-theoretically secure authentication is by using strongly-universal (or almost-strongly-universal) classes of hash functions (see for example [12]). The secret key then determines a hash function of the class, and the message is authenticated by appending its hash value. The authentication code corresponding to an ε-ASU$_2$ class of hash function satisfies

$$p_{imp} = 1/|\mathcal{B}| \qquad \text{and} \qquad p_{sub} \leq \varepsilon \ .$$

There are also different ways to realize authentication codes than with strongly universal hashing. One example is given in [5], where a construction is described with a smaller amount of secret key, but which requires more communication.

Let us now investigate the general scenario in which the key is not entirely secret, i.e., where the opponent Eve has a certain amount of information about the key. We first prove a bound on the information that is gained by Eve when observing a correctly authenticated message. The following lemma states that the min-entropy of the key, given Eve's information $U = u$, decreases by more than the length of the authenticator only with exponentially small probability. (A related result for different entropy measures is proved in [3].) For simplicity, the condition $U = u$ is omitted in the lemma and the proof. Of course the analogous result holds also when all the probabilities are conditioned on $U = u$.

Lemma 5. *Let S, X, and Y be arbitrary discrete random variables (with ranges $\mathcal{S}$, $\mathcal{X}$, and $\mathcal{Y}$, respectively) such that S and X are independent (i.e., $P_{SX} = P_S \cdot P_X$). Then for all real numbers $\ell > 0$*

$$H_\infty(S) - H_\infty(S|X = x, Y = y) \leq \log|\mathcal{Y}| + \ell$$

holds with probability greater than $1 - 2^{-\ell}$ or, more precisely,

$$P_{XY}\left[\{(x,y) \in \mathcal{X} \times \mathcal{Y} \ : \ H_\infty(S) - H_\infty(S|X = x, Y = y) > \log|\mathcal{Y}| + \ell\}\right] < 2^{-\ell} \ .$$

Proof. Let $p_0 := 2^{-\ell}/|\mathcal{Y}|$. Then we have for all $x \in \mathcal{X}$

$$P_{Y|X=x}[\{y \ : \ P_{Y|X=x} < p_0\}] < 2^{-\ell} \ ,$$

and hence

$$P_{XY}\left[\{(x,y) \in \mathcal{X} \times \mathcal{Y} \ : \ P_{Y|X=x}(y) < p_0\}\right] < 2^{-\ell} \ .$$

This inequality implies that

$$\begin{aligned} P_{S|XY}(s,x,y) &= \frac{P_{SXY}(s,x,y)}{P_{XY}(x,y)} = \frac{P_S(s) \cdot P_X(x) \cdot P_{Y|SX}(y,s,x)}{P_X(x) \cdot P_{Y|X}(y,x)} \\ &\leq \frac{P_S(s)}{P_{Y|X}(y,x)} \leq \frac{P_S(s)}{p_0} = P_S(s) \cdot |\mathcal{Y}| \cdot 2^{\ell} \ , \end{aligned}$$

holds with probability greater than $1-2^{-\ell}$ (over values x and y). The statement of the lemma follows by maximizing over all $s \in \mathcal{S}$, and by taking negative logarithms. □

We will show in Section 5 that the bounds of this lemma (and hence also those of the following theorem) are almost tight.

We can now prove a result concerning authentication with a partially secret key which states that information-theoretically secure authentication is possible under the sole condition that no conditional probability of a certain key, given Eve's information, exceeds a bound which is roughly $1/\sqrt{|\mathcal{S}|}$.

Theorem 6. *Assume that two parties Alice and Bob have access to a random variable S, which is a binary string of length n (n even), and that S is used as the key in the authentication scheme based on strongly-universal hashing described in Section 2. Assume further that an adversary Eve knows a random variable U, jointly distributed with S according to some probability distribution, and that Eve has no further information about S. Let*

$$H_\infty(S|U=u) \geq \left(\frac{1}{2}+t\right) \cdot n$$

for a particular realization u of U, and let $\mathcal{D}$ be the event that Eve can either insert a message (successful impersonation attack) or modify a message sent by Alice or Bob (successful substitution attack) without being detected. Then for every strategy, the conditional probability of $\mathcal{D}$, given $U=u$, can be upper bounded as follows:

$$P(\mathcal{D}|U=u) \leq 2^{-(tn/2-1)} \qquad (2)$$

holds under the condition that the correctly authenticated message observed by Eve is independent of S, given $U=u$.

Remark. Note that in Theorem 6 it need not be assumed that the message observed by Eve be independent of S (but independent of S *given* $U=u$). For example (2) holds also when the message is selected by Eve herself.

Proof. First we prove an upper bound on the success probability p_{imp} of the impersonation attack. For every possible message $x \in GF(2^{n/2})$ and for every authenticator $y \in GF(2^{n/2})$ there exist exactly $2^{n/2}$ possible keys such that the authentication is correct. The probability of such a set of keys, given $U=u$, is upper bounded by

$$2^{n/2} \cdot 2^{-H_\infty(S|U=u)} \leq 2^{n/2} \cdot 2^{-(1/2+t)n} = 2^{-tn} ,$$

and hence

$$p_{imp} \leq 2^{-tn} .$$

In a substitution attack the adversary sees a message-authentication pair $(X,Y) \in GF(2^{n/2})^2$, where X is independent of S given $U=u$. According to

Lemma 5 (applied to distributions conditioned on $U = u$), we have for every $r > 0$ that

$$H_\infty(S|X = x, Y = y, U = u) \geq \frac{n}{2} + tn - \log|\mathcal{Y}| - rn = (t - r)n$$

holds with probability greater than $1 - 2^{-rn}$. A successful substitution attack immediately yields the key $S = (A, B)$ because the equations $Y = AX + B$ (from the observed message) and $Y' = AX' + B$ (from the modified message) uniquely determine the key (and can efficiently be solved). Hence the success probability of such an attack can be upper bounded as follows:

$$p_{sub} < 2^{-rn} \cdot 1 + (1 - 2^{-rn}) \cdot 2^{-(t-r)n} < 2^{-rn} + 2^{-(t-r)n} \ . \qquad (3)$$

The reason for this is that with probability greater than $1 - 2^{-rn}$, the maximal probability of a particular key is at most $2^{-(t-r)n}$. Inequality (3) is true for every $r > 0$; the choice $r = t/2$ gives

$$p_{sub} \leq 2 \cdot 2^{-tn/2} = 2^{-(tn/2-1)} \ .$$

The probability $P(\mathcal{D}|U = u)$ is equal to the maximum of p_{imp} and p_{sub}, given $U = u$. This concludes the proof. □

3 Privacy Amplification with Universal Hashing

The results in this and the next section are of the following type: If the min-entropy of a partially secret string S of length n, given the opponent's knowledge, is greater than a certain fraction of n, then Alice and Bob can, by communication over a non-authentic and insecure channel, agree on a common string about which Eve has virtually no information. The maximal length of the resulting highly secret string depends on Eve's knowledge about S and the security conditions. The idea is to use the partially secret string in a first step to authenticate a message containing the description of a function from a suitable class of hash functions. In the second step, this hash function is used for privacy amplification, and the string is used again as the input to this hash function. There are two possibilities to proceed: one can divide the string into two parts and use the first part for authentication and the remaining part as the argument for the final privacy amplification. The second possibility is to use the whole string for both authentication and as argument for privacy amplification. The disadvantage of the second possibility is that the authenticator gives Eve information about the argument of the hashing. A drawback of the first method is that Eve's information about S could be about either string (in fact about both, see below). However, the following lemma implies a tight bound on Eve's information about substrings.

Lemma 7. *Let $S = (S_1, S_2, \ldots, S_n)$ be a random variable consisting of n binary random variables. For any k-tuple $\underline{i} = (i_1, i_2 \ldots, i_k)$, where $1 \leq i_1 < i_2 < \ldots < i_k \leq n$, let $S_{\underline{i}}$ be the string $(S_{i_1}, S_{i_2}, \ldots, S_{i_k})$. Then*

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

Proof. A string $(s_{i_1}, s_{i_2}, \ldots, s_{i_k})$ corresponds to exactly 2^{n-k} strings $(s_1, \ldots, s_n)$. Hence the maximal probability of such a k-bit string is at most 2^{n-k} times the maximal probability of a string in S, i.e., $H_\infty(S) - H_\infty(S_{\underline{i}}) \leq n - k$. □

Remark. Note that when the string S is split into two parts S_l and S_r, then the bounds of Lemma 7 applied to S_l and S_r are tight simultaneously. For example let $s = (s_l, s_r)$ be a particular n-bit string, and let s_l and s_r be the first and second half 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 but counter to intuition, Eve's information about S in terms of min-entropy appears entirely in both substrings S_l and S_r.

The following theorem states that if Eve's knowledge (in terms of H_∞) is less than one third of the length of the entire string (this is an intuitive, but somewhat imprecise description of $H_\infty(S|U = u) > 2n/3$), then privacy amplification by not authenticated public discussion is possible using two thirds of the string to authenticate a hash function from a universal class, and the remaining third as the input to the hash function. We can assume that the length of the string is divisible by 3 (otherwise Alice and Bob discard one or two bits).

Theorem 8. *For every n (multiple of 3) and for all positive numbers $t \leq 1/3$ and r such that $(t - r)n$ is a positive integer, there exists a*

$$\left(n\, ,\, (2/3 + t)n\, ,\, (t - r)n\, ,\, 2^{-rn}/\ln 2\, ,\, 2^{-(tn/2-1)}\right) \ \text{- } \textit{protocol}$$

for privacy amplification over an insecure and non-authentic channel.

Proof. Let $n = 3k$, and let S be the random variable known to Alice and Bob where $\mathcal{S} \subset \{0,1\}^n$. Let further U be the opponent Eve's information about S, and let finally

$$H_\infty(S|U = u) \geq \left(\frac{2}{3} + t\right) \cdot n$$

for a particular $u \in \mathcal{U}$. We denote by S_1 the string consisting of the first $2k = 2n/3$ bits of S (more precisely, S_1 is interpreted as a pair (A, B) of elements of $GF(2^k)$), and by S_2 the remaining k bits (i.e., $S_2 \in GF(2^k)$). The idea of the protocol is to use S_1 for authenticating an element of the universal class of hash functions described in Section 1, and S_2 as the input to this function. According to Lemma 7 applied to conditional distributions (with respect to $U = u$),

$$H_\infty(S_1|U = u) \geq \left(\frac{1}{3} + t\right) \cdot n$$

and

$$H_\infty(S_2|U=u) \geq tn \ .$$

Alice randomly chooses an element X of $GF(2^k)$, which she sends to Bob together with the authenticator $Y = AX + B$ (see Section 2). According to Theorem 6, the probability $P(\mathcal{D}|U=u)$ of undetected modification is bounded by

$$P(\mathcal{D}|U=u) \leq 2^{-(tn/2-1)} \ .$$

Let the hash function be specified by X (see Section 1). The argument S_2 of the hash function satisfies

$$H_2(S_2|U=u) \geq H_\infty(S_2|U=u) \geq tn \ .$$

Let S' be the first $(t-r)n$ bits of $S_2 \cdot X$ (where the product is taken in $GF(2^k)$). Then

$$H(S'|XY, U=u) \geq (t-r)n - \frac{2^{-rn}}{\ln 2}$$

follows from Theorem 3. □

It is not difficult to verify that the use of authentication codes based on the ε-ASU_2 classes of hash functions explicitly given in [12] do not lead to a better result than stated in Theorem 8. For a more detailed discussion of the optimality of our results, see Section 5.

4 Privacy Amplification with Extractors

It appears that the condition in Theorem 8 on Eve's min-entropy about S can be weakened if the description of the hash function is shorter. Extractors are a method for extracting all or part of the min-entropy of a random source into an almost uniformly distributed string by requiring only a small amount of truly random bits. By using extractors instead of universal hashing for privacy amplification, we show that privacy amplification can be secure against an active opponent, provided his min-entropy about S exceeds half of the length of the string. The length of the resulting secret string can be a constant fraction of Eve's min-entropy about S minus $n/2$.

In [10], extractors are defined as follows (for an introduction to the theory of extractors, see for example [9] or [10]).

Definition 9. **[10]** A function $E : \{0,1\}^n \times \{0,1\}^w \to \{0,1\}^{n'}$ 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 distance of the distribution of $[V, E(X,V)]$ to the uniform distribution of $\{0,1\}^{w+n'}$ is at most ε' when choosing X according to P and V according to the uniform distribution in $\{0,1\}^w$. The distance between two distributions P and P' on a set $\mathcal{X}$ is defined as

$$d(P, P') := \frac{1}{2} \sum_{x \in \mathcal{X}} |P(x) - P'(x)| \ .$$

Various possible constructions of extractors have been described. The following theorem of [10] states that it is possible to extract a constant fraction of the min-entropy of a given source where the number of required random bits is polynomial in the logarithm of the length of the string and in $\log(1/\varepsilon')$.

Theorem 10. [10] *For any parameters $\delta = \delta(n)$ and $\varepsilon' = \varepsilon'(n)$ with $1/n \leq \delta \leq 1/2$ and $2^{-\delta n} \leq \varepsilon' \leq 1/n$, there exists an efficiently computable (δ, ε')-extractor $E: \{0,1\}^n \times \{0,1\}^w \to \{0,1\}^{n'}$, where $w = O(\log(1/\varepsilon') \cdot (\log n)^2 \cdot (\log(1/\delta))/\delta)$ and $n' = \Omega(\delta^2 n / \log(1/\delta))$.*

We also need the following lemma, which states that a random variable whose distribution is close to uniform (in terms of the distance d) has a Shannon entropy close to maximal.

Lemma 11. *Let Z be a random variable with range $\mathcal{Z} \subset \{0,1\}^k$. Then*

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

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

Proof. Let $d := d(U_k, Z)$. We can assume that $d < 1 - 2^{-k}$ 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 [decreasing] a *smaller* probability increases [decreases] 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

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

□

For certain values of d equality can hold in the above inequality. In particular $H(Z) \geq k \cdot (1 - d(U_k, P_Z))$ is false in general: $H(Z) = 0$ is possible when $d(U_k, P_Z) = 1 - 2^{-k} < 1$.

Theorem 12 below states that if Eve's min-entropy about S is greater than half of the length of S, then a constant (where this constant is not explicitly specified) fraction of the difference of this entropy and half of the length of S (plus a security parameter) can be extracted by privacy amplification using public discussion over a non-authentic channel, provided that S is sufficiently long. In contrast to the proof of Theorem 8, the entire string S is used twice here: once for authentication, and once as the input of the extractor.

Theorem 12. *Let t and r be positive numbers such that $r < t < 1/2$. There exists a constant c with the following property. Let $\varepsilon'(n)$ be a function such that*

$$\lim_{n\to\infty} n \cdot \sqrt{\varepsilon'(n)} = 0 \tag{4}$$

and

$$\varepsilon'(n) = 2^{-o(n)/(\log n)^2} \tag{5}$$

(i.e., $\log(1/\varepsilon'(n)) \cdot (\log n)^2/n \to 0$ for $n \to \infty$). Then there exists a bound n_0 such that for all $n \geq n_0$ there exists an $n' \geq c(t-r)n$ and an

$$\left(n\,,\,(1/2+t)n\,,\,n'\,,\,n'\cdot\left(2\sqrt{\varepsilon'(n)}+2^{-rn}+2^{-n'}\right)\,,\,2^{-(tn/2-1)}\right)\text{ - protocol}$$

for privacy amplification over an insecure and non-authentic channel.

Remark. The function $\varepsilon'(n)$ is directly related to the tolerable amount of information that Eve obtains about the key S' as a function of the length n of the string S. Possible functions $\varepsilon'(n)$ satisfying both (4) and (5) are $\varepsilon'(n) = 2^{-n/(\log n)^3}$, $\varepsilon'(n) = 2^{-n^\alpha}$ for any $0 < \alpha < 1$, or $\varepsilon'(n) = 1/(n^2(\log n)^2)$. The choice of a more restrictive $\varepsilon'(n)$ with respect to Eve's knowledge increases the bound n_0.

Proof. The number w of random bits required as the second part of the input for a (δ, ε')-extractor according to Theorem 10, where $\delta = t - r$ is constant, is

$$w = O(\log(1/\varepsilon') \cdot (\log n)^2)\ , \tag{6}$$

and the length n' of the output is $\Omega((t-r)n)$, i.e., $n' \geq c(t-r)n$ for some constant c.

Because of (5) and (6) there exists an n_0 (depending on $\varepsilon'(n)$) such that $n \geq n_0$ implies $w \leq n/2$. Let $n \geq n_0$ (and we can assume that n is even). The message sent from Alice to Bob is a random element $X \in GF(2^{n/2})$ (of which the first w bits are used as the second input V to the extractor) and is authenticated by $Y = S_1 \cdot X + S_2 \in GF(2^{n/2})$, i.e., the authentication scheme based on strongly universal hashing (see Section 2) is used with S as partially secret key, and where S_1 and S_2 are the first and second half of S, interpreted as elements of $GF(2^{n/2})$. According to Theorem 6, Eve's probability of undetected modification satisfies

$$P(\mathcal{D}|U = u) \leq 2^{-(tn/2-1)}\ .$$

Lemma 5 implies that

$$H_\infty(S|X = x, Y = y, U = u) \geq (t-r)n \tag{7}$$

holds with probability greater than $1 - 2^{-rn}$. We can assume that $(t-r)n$ is an integer. If (7) holds, the extractor's output satisfies

$$d([V, E(S,V)], U_{w+n'}) \leq \varepsilon'(n)\ . \tag{8}$$

Here and below the random variable S is meant to be distributed according to $P_{S|U=u}$, i.e., Eve's point of view is taken. It is easy to see that the distance in (8) is the expected value of the distances $d(E(S,V),U_{n'})$, where V is chosen at random from $\{0,1\}^w$. We conclude that for every K,

$$P_V\left[d(E(S,V),U_{n'}) \leq K \cdot \varepsilon'(n)\right] \geq 1 - \frac{1}{K} \ , \tag{9}$$

where V is uniformly distributed in $\{0,1\}^w$. From (7) and (9), with the special choice $K = \sqrt{1/\varepsilon'(n)}$, we obtain that

$$d(E(S,V),U_{n'}) \leq \sqrt{\varepsilon'(n)}$$

holds with probability at least $1 - \sqrt{\varepsilon'(n)} - 2^{-rn}$. With Lemma 11, this leads to

$$\begin{aligned} H(S'|XY, U=u) &\geq (1 - \sqrt{\varepsilon'(n)} - 2^{-rn}) \cdot (1 - \sqrt{\varepsilon'(n)} - 2^{-n'}) \cdot n' \\ &\geq n' - n' \cdot (2\sqrt{\varepsilon'(n)} + 2^{-rn} + 2^{-n'}) \ . \end{aligned}$$

□

5 Optimality Considerations

This section provides evidence that the result of Section 2 (and the condition on Eve's knowledge in Theorem 12) is optimal: if Eve's min-entropy about S is less than half of the length of the string, then no non-trivial upper bound on the probability of undetected modification $P(\mathcal{D}|U=u)$ can be shown. This fact also implies that no better result than that of Section 3 can hold if one splits the string into two parts, one of which is used for authentication and the other for privacy amplification.

We will show that the bound given in Lemma 5 is tight (Lemma 13) and that this implies that when using the authentication code based on an ε-ASU$_2$ class of hash functions, a substantially better result than Theorem 6 cannot be derived. The omission of the additional random variable X in Lemma 13 is for simplicity. It is obvious that the same tightness result also holds in the situation of Lemma 5.

Lemma 13. *For every integer $k > 0$ and for every number $\ell \geq 0$ there exist random variables S and Y (with ranges $\mathcal{S}$ and $\mathcal{Y}$) such that $|\mathcal{Y}| = k$, and such that*

$$H_\infty(S|Y=y) = H_\infty(S) - \log|\mathcal{Y}| - \ell \tag{10}$$

holds with probability

$$\left(1 - \frac{1}{|\mathcal{Y}|}\right) \cdot 2^{-\ell} \ , \tag{11}$$

and even with probability 1 *in the case $\ell = 0$.*

Proof. Let $R := 2^\ell k$ and $\mathcal{S} := \{s_1, s_2, \ldots, s_R\}$ with $P_S(s_i) = 1/R$ for $1 \leq i \leq R$. Let further

$$P_{Y|S}(y_i, s_i) = 1$$

for $1 \leq i \leq k-1$, and

$$P_{Y|S}(y_k, s_i) = 1$$

for $i \geq k$. Then we have $H_\infty(S) = \log k + \ell = \log|\mathcal{Y}| + \ell$, and $H_\infty(S|Y = y_i) = 0$ for $1 \leq i \leq |\mathcal{Y}| - 1$. Hence (10) holds with the probability given in (11). In the case $\ell = 0$, let $Y = S$, and the result follows immediately. □

Let us now assume that an ε-ASU$_2$ class $\mathcal{H}$ of hash functions mapping $\mathcal{X}$ to $\mathcal{Y}$ (where $|\mathcal{Y}| \leq |\mathcal{X}|$ and $1/\varepsilon \leq |\mathcal{Y}|$) is used for authentication. We show that when Eve observes a correctly authenticated message, then the min-entropy of the key must be reduced by at least half of the key size to obtain a lower bound for the min-entropy of the correct authenticator of a different message. This implies the optimality of our results in the earlier sections when using this authentication method.

According to Lemma 13 we must assume that the min-entropy of the key, given Eve's information, is decreased by at least $\log|\mathcal{Y}|$ when Eve observes a correctly authenticated message x. On the other hand, given an arbitrary additional message-authenticator pair (x', y') (with $x' \neq x$), it is possible that $\varepsilon \cdot |\mathcal{H}|/|\mathcal{Y}|$ keys are compatible with both pairs. Hence the conditional min-entropy of the correct authenticator for a given message x' (this min-entropy is directly linked with the substitution attack success probability) can, in the worst case, be smaller than the min-entropy of the key by $\log|\mathcal{H}| - \log|\mathcal{Y}| - \log(1/\varepsilon)$. Both reductions of the initial min-entropy together are, in the worst case, of size $R := \log|\mathcal{H}| - \log(1/\varepsilon)$. Because of

$$|\mathcal{H}| \geq \frac{|\mathcal{Y}| - 1}{\varepsilon} \geq \frac{1}{\varepsilon^2}(1 - \varepsilon)$$

(see Theorem 4.2 in [12]), we have

$$\log\left(\frac{1}{\varepsilon}\right) \leq \frac{1}{2} \cdot \log|\mathcal{H}| - \frac{1}{2} \cdot \log(1 - \varepsilon) \leq \frac{1}{2} \cdot \log|\mathcal{H}| + \frac{\varepsilon}{2\ln 2}$$

and

$$R \geq \frac{1}{2} \cdot \log|\mathcal{H}| - \frac{\varepsilon}{2\ln 2} \ . \tag{12}$$

The lower bound in (12) is almost $(\log|\mathcal{H}|)/2$. Hence these worst-case estimates suggest that the result of Theorem 6 and the condition in Theorem 12 are optimal.

6 Open Problems

It is conceivable that stronger results than those of Theorems 6, 8, and 12 can be shown under certain additional conditions on Eve's information. We state as an open problem to find such conditions, as well as the question whether the results of the previous sections can be improved by using different authentication protocols (e.g., [5]), or even a completely different type of protocol for privacy amplification by not authenticated public discussion. Finally, are there different kinds of scenarios, besides the situations of independent repetitions of a random experiment [6] and of privacy amplification, for which a positive result can be proved for secret-key agreement by not authenticated public discussion?

Acknowledgments

We would like to thank Christian Cachin for interesting discussions on the subject of this paper.

References

1. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, Generalized privacy amplification, *IEEE Transactions on Information Theory*, Vol. 41, Nr. 6, 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, Smooth entropy and Rényi entropy, *Advances in Cryptology - EUROCRYPT '97*, Lecture Notes in Computer Science, Vol. 1233, pp. 193-208, Springer-Verlag, 1997.
4. T. M. Cover and J. A. Thomas, *Elements of information theory*, Wiley Series in Telecommunications, 1992.
5. P. Gemmell and M. Naor, Codes for interactive authentication, *Advances in Cryptology - CRYPTO '93*, Lecture Notes in Computer Science, Vol. 773, pp. 355-367, Springer-Verlag, 1993.
6. U. 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.
7. U. M. Maurer, A unified and generalized treatment of authentication theory, *Proceedings 13th Symp. on Theoretical Aspects of Computer Science - STACS '96*, Lecture Notes in Computer Science, Vol. 1046, pp. 387-398, Springer-Verlag, 1996.
8. 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.
9. N. Nisan, Extracting randomness: how and why - a survey, preprint, 1996.
10. N. Nisan and D. Zuckerman, Randomness is linear in space, *Journal of Computer and System Sciences*, Vol. 52, No. 1, pp. 43-52, 1996.
11. G. J. Simmons, A survey of information authentication, *Proc. of the IEEE*, Vol. 76, pp. 603-620, 1988.
12. 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.