Unconditional Authenticity and Privacy from an Arbitrarily Weak Secret

Renato Renner[1] Stefan Wolf[2]

[1] Department of Computer Science, ETH Zürich, Switzerland.
renner@inf.ethz.ch
[2] Département d'Informatique et R.O., Université de Montréal, Canada.
wolf@iro.umontreal.ca

Abstract. Unconditional cryptographic security cannot be generated simply from scratch, but must be based on some given primitive to start with (such as, most typically, a private key). Whether or not this implies that such a high level of security is necessarily impractical depends on how weak these basic primitives can be, and how realistic it is therefore to realize or find them in—classical or quantum—reality. A natural way of minimizing the required resources for information-theoretic security is to reduce the *length* of the private key. In this paper, we focus on the *level of its secrecy* instead and show that even if the communication channel is completely insecure, a shared string of which an arbitrarily large fraction is known to the adversary can be used for achieving fundamental cryptographic goals such as message authentication and encryption. More precisely, we give protocols—using such a weakly secret key—allowing for both the exchange of authenticated messages and the extraction of the key's entire amount of privacy into a shorter virtually secret key. Our schemes, which are highly interactive, show the power of two-way communication in this context: Under the given conditions, the same objectives cannot be achieved by one-way communication only.

Keywords. Information-theoretic security, authentication, privacy amplification, extractors, quantum key agreement.

1 Information-Theoretic Security and its Price

1.1 Unconditional Authentication and Privacy Amplification with an Arbitrarily Weak Key by Completely Insecure Communication

The main advantage of *information-theoretic*—as opposed to *computational*—cryptographic security is the fact that it can be based on a mathematical proof which does not depend on any assumption on the hardness of certain computational tasks nor on an adversary's computing power or memory space. An important *disadvantage* of such unconditional security, on the other hand, is often perceived to be its impracticality. At the origin of this belief stands Shannon's famous result [22] stating that the perfectly secret transmission of a message over a public channel requires a private key of, roughly speaking, the same length.

In the present paper, we take a step towards making unconditional security more practical by showing that such a private key can be generated, by communication over a *completely insecure channel*, from an *arbitrarily weakly secret key*. One of the main ingredients of our protocol is a new interactive method for unconditionally secure message authentication requiring only a weak secret key as well. No such method has previously been known which works when the adversary knows more than half of the partial secret.

The problem of extracting a highly secret from a longer, partly compromised key—so-called *privacy amplification* [4], [3]—has been studied intensively since it is the final step of any information-theoretic key-agreement protocol based on classical or quantum correlations (e.g., quantum key agreement [2]). It is a direct consequence of our result that the assumption—usually made in the context of privacy amplification—that the communication channel is authentic can simply be dropped: Privacy amplification by communication over a *completely insecure* channel is possible even with arbitrarily weakly secret strings, and the length of the extractable private key is asymptotically the same as in the case of an authenticated channel or, equivalently, an only passive adversary. Previous results were pointing into another direction [14], [24], [17].

1.2 Towards Making Unconditional Security Practical

The main motivation for this work is to relax the conditions under which unconditional cryptographic security can be achieved. Our results should be seen in the context of a number of more or less recent steps, taken by various authors, towards making unconditionally provable security more practical by reducing the requirements for achieving it. For instance, techniques and protocols have been proposed allowing for generating provably secret keys from noisy channels [26], [6], weakly correlated classical information by public [13], [3] and even unauthenticated [15], [16], [17] communication, or from quantum channels [2]; reducing the required key size for authentication [23], [10] as well as encryption [21]; basing cryptographic tasks on keys from weak random sources [7]; or realizing information-theoretic secrecy against memory-bounded yet otherwise unlimited adversaries [12], [1], [8], to mention just a few.

1.3 Determining the Cryptographic Value of an Arbitrarily Weak Secret: The Power of Interaction

In the setting where two parties initially share some key in order to achieve cryptographic goals, two natural quantities to be minimized are the *length* and the *level of privacy* of this key. In this paper we address the question what the cryptographic value of a shared string is about which the adversary has almost complete information if we also assume her to have perfect read and write access to the communication channel. Our results show that such a key is useful both for achieving authenticity (Section 2) and privacy (Section 3) in this scenario.

We consider the following model. Two parties Alice and Bob both know an n-bit string S about which an adversary Eve has some partial information U. We also assume her to be able to read, modify, or delete any message sent over the communication channel connecting the legitimate partners. Ultimately, the goal of Alice and Bob is the exchange of a message M in a both authentic and confidential way. We achieve this in two steps which are described and analyzed in detail in Sections 2 and 3, respectively. First, we give an interactive protocol that allows, using the weak secret S, for the authentication of short messages. Second, we use this authentication technique as a building block of a protocol for distilling, from S, a *highly* secret string S'' the length of which is, roughly, the *min-entropy*[3] of S given Eve's knowledge $U = u$; this string can then be used for private key cryptography, in particular encryption and message authentication.

We describe our results in more detail. Let $0 < t \leq 1$ be an arbitrary constant, and assume that tn is a lower bound on the min-entropy of (the n-bit string) S from Eve's viewpoint (i.e., conditioned on $U = u$). Then Protocol AUTH of Section 2 allows for authenticating an l-bit message M, where $l = tn/s$ for a security parameter s: The probability of a successful active attack is of order $2^{-\Omega(s)}$. The protocol is computationally very efficient and uses $\Theta(l)$ rounds of communication. Note that all previously described protocols—interactive or one-way—for authentication with a partially secret key work only under the assumption that the key is more than "half secret" [24], [17].

Protocol AUTH can be used as a building block of a protocol allowing for distilling a short but highly secret key S'' from the initial string S, using communication over a completely insecure channel only. Let again tn be the min-entropy of the shared n-bit string S in Eve's view. Then Protocol PA of Section 3 allows Alice and Bob to generate a common string S'' of length $(1-o(1))tn$ about which Eve only has an exponentially (in n) small amount of information. In contrast to privacy amplification over an authenticated channel, privacy amplification secure against active adversaries has so far been known possible only for keys offering a relatively high secrecy level initially (at least two thirds of the key should be unknown to the adversary), and the length of the extractable secret was only a small fraction of the key's entropy [14], [24], [17]. It was speculated that this might be the price which has to be paid for the missing authenticity of the channel. Protocol PA shows that this is not so: Privacy amplification secure against active adversaries is equally powerful as against passive adversaries with respect to the condition on the initial string as well as to the size of the extractable secret.

Our results can alternatively be interpreted as realizing encryption and authentication using *private keys generated by weak random sources*—instead of *highly compromised keys*—as studied in [18], [7]. In [18] it was shown that weakly random keys from certain sources with substantial min-entropy do not allow for

[3] The *min-entropy* $H_\infty(X)$ of a random variable X with range $\mathcal{X}$ is simply the negative logarithm of the maximal probability occurring: $H_\infty(X) := -\log(\max_{x\in\mathcal{X}} P_X(x))$. We have $0 \leq H_\infty(X) \leq H(X) \leq \log|\mathcal{X}|$ for all random variables X. All logarithms, here and in the rest of the paper, are binary.

information-theoretically secure (one-way) encryption; in [7], it was proven that a weakly random key allows for (one-way) authentication only if its min-entropy exceeds half its length. Therefore, the results of [18] and [7] suggest that not all private keys with substantial randomness—i.e., min-entropy—are useful for basic cryptographic tasks. This is true, however, only in the one-way communication model: Our results add to this picture by showing that if *two-way communication* is allowed (and perfect randomness is available locally), then keys from *all* sources with non-negligible min-entropy allow for *both* authentication and encryption.

In [10], it has been shown that the use of interaction in authentication allows for dramatically reducing the *length* of the used (private) key. Our results underline the power of two-way communication, suggested by that result, in this context: Interaction alternatively allows for strongly relaxing the condition on the *degree of privacy* of the used key.

2 Authentication with an Arbitrarily Weak Key

2.1 Intuition and Building Blocks

In standard (one-way) authentication, the message to be authenticated is sent together with a so-called *authenticator*, i.e., an additional string depending on that message and the secret key. These methods fail as soon as the adversary has substantial knowledge about the key (more precisely, half the knowledge in terms of min-entropy [7]) since, very roughly speaking, this knowledge could consist of the correct authenticator for one or several messages. A possible way of overcoming this problem is to use a challenge-response protocol. In [24] (see also [17]), for instance, it was proposed that one party, the *sender*, sends the message as a *challenge*, the reception of which is confirmed by the other, the *receiver*, by sending *back* an authenticator. In this case the person in the middle Eve *has* to find the correct authenticator of a message which is *not of her choice*, even in the case of a substitution attack. As shown in [24], [17], one advantage of this scheme is that the authenticator can be short and thus leaks only a small amount of additional information about the key. On the other hand, however, its security could be shown only under the assumption that the adversary knows less than half the key; the same condition that characterizes the possibility of one-way authentication [7]. The reason is the attack where Eve uses the receiver of the message as an oracle and gets the correct response to a challenge of her choice (where she can make this choice adaptively after having seen the challenge for which she has to generate the correct response). In summary, such an interactive authentication method, where *the challenge is identical with the message* to be authenticated, may be preferable to one-way authentication in certain cases [24], [17], but *cannot*, in order to resist adaptive substitution attacks, tolerate Eve to have more knowledge about the "private" key—namely roughly half of it—than simple one-way authentication.

In Section 2.2 we propose a new protocol solving this problem by, roughly speaking, *preventing adaptive substitution attacks completely.* The main idea is to encode the message differently: The message bits do not determine the challenge strings (which are just random), but rather *which of them will be answered*.

Let us first have a look at how the authenticator should depend on the key and the message. Since we want the adversary to be able to compute the correct authenticator to only very few messages unless she knows the entire key, a natural way is to interpret the key as a polynomial, and let the authenticator be its evaluation at a point determined by the challenge. (This idea was already used in previous protocols of this type [24], [17].) Lemma 1 states that when this function is used, then even an adversary who knows almost the entire key cannot correctly respond to a random challenge except with small probability. A similar result was shown in [24], [17] with respect to Rényi entropy H_2.

Lemma 1. *Let n, k, and a be positive integers such that $n = k \cdot a$ holds, and let, for $x \in \{0,1\}^k$, $f_x : \{0,1\}^n \to \{0,1\}^k$ be the function $f_x(s) := \sum_{i=0}^{a-1} s_i x^i$. Here, the strings $s_i \in \{0,1\}^k$ are defined by $s = (s_0, s_1, \ldots, s_{a-1})$, and the k-bit strings s_i, x, and $f_x(s)$ are interpreted as elements of $GF(2^k)$ with respect to a fixed basis of $GF(2^k)$ over $GF(2)$. Let now S be a random variable with range $\mathcal{S} \subseteq \{0,1\}^n$ and distribution P_S such that when given $x \in \{0,1\}^k$ chosen according the uniform distribution, the probability that $f_x(S)$ can be guessed correctly is α. Then we have $\max_{s\in\mathcal{S}} P_S(s) \geq (\alpha - a/2^k)^a$ or, equivalently, $\alpha \leq 2^{-H_\infty(S)/a} + a/2^k$.*

Proof. We can assume that the guessing strategy is deterministic, i.e., only depends on x. For $s \in \mathcal{S}$, let α_s be the number of x for which $f_x(s)$ is guessed correctly, divided by 2^k. Then we have $\alpha = E_S[\alpha_S]$. The probability that $f_{x_i}(S)$ is guessed correctly simultaneously for a randomly chosen $x_1, \ldots, x_a \in GF(2^k)$ is lower bounded by

$$E_S\left[\prod_{i=0}^{a-1}\left(\alpha_s - \frac{i}{2^k}\right)\right] \geq E_S\left[\left(\alpha_s - \frac{a}{2^k}\right)^a\right] \geq \left(E_S\left[\alpha_s - \frac{a}{2^k}\right]\right)^a = \left(\alpha - a/2^k\right)^a \ . \quad (1)$$

(The second inequality of (1) is Jensen's inequality [9].) Therefore, there must exist a particular a-tuple $x_1, \ldots, x_a$ such that the values $f_{x_i}(S)$ are simultaneously guessed correctly with probability at least $(\alpha - a/2^k)^a$. On the other hand, S is uniquely determined by these $f_{x_i}(S)$ since $f_x(s)$ is a polynomial in x of degree at most $a-1$ with coefficients $s_0, \ldots, s_{a-1}$. Hence there must exist a value $s \in \mathcal{S}$ with probability $P_S(s)$ at least (1), and this concludes the proof. □

During the execution of Protocol AUTH and Protocol PA, the adversary observes a number of messages that depend on the key and hence leak information about it. An important argument in the analysis of these protocols is an upper bound on the effect of such information on the min-entropy of the key (from the adversary's viewpoint). Roughly speaking, the min-entropy does not, except with small probability, decrease much more than by the number of physical bits observed. Results similar to Lemma 2 were proven in [5], [24], [17].

Lemma 2. *Let S, V, and W be discrete random variables with ranges $\mathcal{S}$, $\mathcal{V}$, and $\mathcal{W}$, respectively, such that S and V are independent, and let $b \geq 0$. Then* $\text{Prob}_{VW}[H_\infty(S|V=v, W=w) \geq H_\infty(S) - \log|\mathcal{W}| - b] \geq 1 - 2^{-b}$.

Proof. We have $\text{Prob}\,[P_{W|V}(w,v) < 2^{-b}/|\mathcal{W}|] < 2^{-b}$ (where $P_{W|V}$ stands for the conditional distribution of W given V), which implies that $P_{S|VW}(s,v,w) = P_{SVW}(s,v,w)/P_{VW}(v,w) = P_S(s) \cdot P_V(v) \cdot P_{W|SV}(w,s,v)/(P_V(v) \cdot P_{W|V}(w,v)) \leq P_S(s)/P_{W|V}(w,v) \leq P_S(s) \cdot |\mathcal{W}| \cdot 2^b$ holds with probability at least $1 - 2^{-b}$ over V and W. The statement now follows by maximizing over $s \in \mathcal{S}$ and by taking negative logarithms. □

Lemma 3 finally gives a bound on the min-entropy of substrings in terms of the min-entropy of the full string [14]. It follows from the fact that every r-bit string s' corresponds to exactly 2^{n-r} n-bit strings s.

Lemma 3. *Let S be a random variable with range $\mathcal{S} \in \{0,1\}^n$, and let S' be an r-bit substring of S. Then we have $H_\infty(S') \geq H_\infty(S) - (n-r)$.*

2.2 The Authentication Protocol and its Analysis

We now give Protocol AUTH. Let s be a string of length n and k a divisor of n. For $s \in \{0,1\}^n$ and $x \in \{0,1\}^k$, let $f_x(s)$ be defined as in Lemma 1. Finally, $b = (b_1, \ldots, b_r)$ are the bits to be authenticated; the bits are authenticated separately, one after another.

Protocol AUTH (Authentication)

Alice		**Bob**
$s \in \{0,1\}^n$, $k \| n$ $b = (b_1, \ldots, b_r) \in \{0,1\}^r$		$s \in \{0,1\}^n$, $k \| n$
for $i = 1, 2, \ldots, r$: $x_i \in_r \{0,1\}^k$		
	$\xrightarrow{x_i}$	
		compute $f_{x_i}(s)$ $y_i \in_r \{0,1\}^k$
	$\xleftarrow{f_{x_i}(s),\ y_i}$	

if $f_{x_i}(s)$ is incorrect:
 reject and abort
 if $b_i = 1$:
 compute $f_{y_i}(s)$		

	(b_i, a_i), where $a_i = \perp$ if $b_i = 0$, $a_i = f_{y_i}(s)$ if $b_i = 1$. $\longrightarrow$	
		if $f_{y_i}(s)$ is incorrect: reject and abort

Let us first discuss some properties of Protocol AUTH intuitively. Note first that since the values x_i and y_i, which are chosen randomly, are independent from the message bit b_i, a person-in-the-middle attack in which x_i or y_i is substituted by another value is of no use for changing a message bit: only the fact *whether* *a response was given or not* is important. If no such response is given by the legitimate party, then it is, according to Lemma 1, in any case difficult for the adversary to generate one, provided the min-entropy of the key is still large enough. This is why it is hard for the adversary to flip a bit from 0 to 1.

Since, on the other hand, an active adversary can simply delete a given response, it is trivial to flip a message bit 1 into a 0. Furthermore, message bits 0 can always be generated towards Bob without participation of Alice at all. We will take care of these problems later and transform the "semi-authentication protocol" into a complete authentication allowing for no undetected modifications of the message at all.

Let us, however, first make precise what Protocol AUTH, as given above, achieves. Lemma 4 states that the two above-mentioned types of undetected modifications of the message are the only ones possible unless Eve is able to generate a random challenge's response by herself. Note that statement 2 in Lemma 4 is a formalization of the fact that the string Bob receives can be obtained from the string sent by Alice by changing 1s into 0s and generating 0s from scratch.

Lemma 4. *Assume that Alice and Bob execute Protocol AUTH in the presence of an adversary Eve, that Alice has not aborted the protocol and has, so far, authenticated a string $b = (b_1, \ldots, b_j)$. Assume further that either Bob has rejected and aborted, or that in his view, a protocol round has just been completed (i.e., that the last message he received was a message bit together with the authentication string if the bit was* 1*) and that the string sent and authenticated up to this point is, still in Bob's view, $b' = (b'_1, \ldots, b'_{j'})$.*

Then if Eve has been passive*, Bob has not rejected and $b' = b$ holds. If on the other hand Eve is* active*, at least one of the following three statements is true.*

1. *Bob has rejected and aborted.*
2. *There exists an injective monotonically increasing function $g : \{1, \ldots, j\} \to \{1, \ldots, j'\}$ such that for all $1 \le i \le j'$, $b'_i = 1$ implies both $i \in \mathrm{Im}(g)$ and $b_{g^{-1}(i)} = 1$. (Note that this implies, in particular, $j' \ge j$ and $w_H(b') \le w_H(b)$, where w_H denotes the Hamming weight.)*
3. *Eve has successfully computed and sent $f_z(s)$ for a value $z \in \{0,1\}^k$ that she received from Alice or Bob* without *having received another message in-between. (In this case, we say that Eve was able to answer a random challenge* without help.*)*

Proof. Assume first that there is no active adversary, i.e., that no message sent has been modified or deleted. Then, clearly, Bob is accepting and $b' = b$ holds.

Let us suppose that Eve is (potentially) an active adversary. We prove the statement by induction over j. Let first $j = 0$, i.e., Alice has not sent (nor received) any message. Assume that Bob is accepting and has received the string $b' = (b'_1, \ldots, b'_{j'})$. We have to show $b'_i = 0$ for all $1 \le i \le j'$ unless Eve was able to answer a challenge without help. Assume $b'_i = 1$ for some i. Then Bob's challenge y_i must have been correctly answered (by $f_{y_i}(s)$) without Bob sending any other message between sending y_i and receiving the response. Since also Alice has not sent any message so far, Eve must have generated $f_{y_i}(s)$ without help.

Suppose now that the statement is true for $j \geq 0$; we prove its validity for $j+1$. Assume that a protocol round has been concluded in Bob's view, and that 1. and 3. are not true. Let $(b'_1, \ldots, b'_{j''})$ be the string authenticated so far in Bob's view. Just before Bob's receiving of x_{j+1} (or a possibly modified value x'_{j+1}) and sending of $f_{x_{j+1}}(s)$ (or of the value $f_{x'_{j+1}}(s)$ that will allow Eve to determine $f_{x_{j+1}}(s)$—and since 3. is wrong she *must* have received such a message), a protocol round had been concluded in his view and the message received up to that point was an initial substring of b', i.e., $(b'_1, \ldots, b'_{j'})$ for some $j' \leq j''$. At that point, Alice had authenticated the string $(b_1, \ldots, b_j)$. By the induction hypothesis, and since 1. and 3. do not hold, there exists $g : \{0,1\}^j \to \{0,1\}^{j'}$ with the required properties. For establishing the statement for $j+1$ (i.e., proving that g can be extended to $\{1, \ldots, j+1\}$), we have to show two facts. First, $j'' > j'$ must hold, and secondly, we must have $w_H((b_{j'+1}, \ldots, b_{j''})) \leq 1$, where equality implies $b_{j+1} = 1$.

Since Alice has received $f_{x_{j+1}}(s)$, and since this cannot have been generated by Eve without help, Bob must have sent at least one message after Alice's sending of x_{j+1}. Thus, because Bob is still accepting, we have $j'' > j'$. On the other hand, for every value $i \in \{j'+1, \ldots, j''\}$ with $b'_i = 1$, Bob must have received $f_{y_i}(s)$ correctly after his challenge y_i. Since Eve has not computed this value without help (3. is untrue), Alice must have sent a value $f_{y'_{j+1}}$ between Bob's sending and receiving of y_i and $f_{y_i}(s)$, respectively. This implies both $w_H((b_{j'+1}, \ldots, b_{j''})) \leq 1$—since Alice has sent at most one such value during what was a single protocol round in her view—and that in case of equality, Alice must have authenticated the bit $b_{j+1} = 1$ in the last step. This concludes the induction step and the proof. □

Clearly, Protocol AUTH cannot be used *directly* for the authentication of messages $(b_1, \ldots, b_r)$ by the following three reasons. First of all, it is, for an active adversary, easy to flip a bit from 1 to 0 without being detected, or to insert a 0 at any point. Secondly, Eve can block all messages sent after some point without Bob realizing that he only received part of the message. (In this case, Alice, but not Bob, would realize the attack, reject, and abort.) Finally, Bob can be used as an oracle for finding out the entire key: Eve simply impersonates Alice and authenticates a sufficient number of 0s to Bob.

The third problem can be solved by limiting the length of the message; this limit L must be chosen such that even $2L$ values $f_z(s)$ (where $2L$ different values for z can be chosen by Eve) do not reveal the entire key, but leave sufficient uncertainty in terms of min-entropy to guarantee the security of the protocol.

In order to get rid of the first two problems, we restrict the set of possible messages (of even length r): a string $b = (b_1, \ldots, b_r)$ is a valid message only if it is *balanced*, i.e., if half the bits are 0s and the other half are 1s, and if every initial substring $(b_1, \ldots, b_i)$, $i < r$, is *"underweight"*: the number of 1s is strictly less than $i/2$. If, given that the sent string b satisfies these conditions, Bob accepts the outcome[4] only if the received string b' is balanced, then he is prevented from erroneously accepting in case Eve performs one of the described attacks, and $b' = b$ must hold. Fortunately, the given restriction on the strings to be authenticated only reduces the effective message length insignificantly, as Lemma 5 shows. It follows from a well-known result on random walks (see for example [9]), and from Stirling's formula.

Lemma 5. *Let r be an even integer and let $z(r)$ be the number of r-bit strings $b = (b_1, \ldots, b_r)$ satisfying $w_H((b_1, \ldots, b_r)) = r/2$ and $w_H((b_1, \ldots, b_i)) < i/2$ for all $1 \le i < r$. Then we have $z(r) = \binom{r}{r/2}/(2(r-1)) = \Theta(2^r/r^{3/2})$, hence $\log(z(r)) = (1 - o(1))r$.*

Let us assume from now on that Protocol AUTH is used in the way described above: Bob rejects and aborts when given more than L challenges (where L is an additional protocol parameter to be properly chosen), and he accepts the outcome only if the received message is balanced. We will prove that with this modification, Protocol AUTH is a secure authentication protocol.

Note that since the protocol uses two-way communication, also Alice can detect an active attack, reject, and abort the protocol. Unfortunately, it is not possible to achieve agreement of Alice's and Bob's acceptance states in every case in the presence of an active adversary (who can, for instance, delete the final message sent). However, our protocol *does* achieve that whenever Bob accepts, then so does Alice, and Bob has received the correct string (except with small probability). This means that every active attack detected by Alice is automatically also perceived by Bob; the final decision whether the authentication succeeded is hence up to the receiver—just as in one-way authentication.

Theorem 1 makes the security of Protocol AUTH precise. It is only due to simplicity that the result is stated asymptotically. The protocol is useful already for short strings. The proof of Theorem 1 explicitly shows all the involved constants that are neglected in the asymptotic notation.

Theorem 1. *Let S be a random variable, with range $\mathcal{S} \subseteq \{0,1\}^n$, known to Alice and Bob, and let U summarize an adversary Eve's entire knowledge about S. Assume $H_\infty(S|U = u) \ge tn$ for the particular value $u \in \mathcal{U}$ known to Eve, where $0 < t \le 1$ is a constant. Let now $k < tn/7$ be of order[5] $k = \omega(\log n)$. Then, for some $l = (1 - o(1))(tn - k)/3k$, Protocol AUTH can be used to authenticate, by communication over a completely insecure channel, a message m of at most l bits sent from Alice to Bob. More precisely, the following holds: If Eve is* passive, *then Alice and Bob accept the outcome of the protocol and Bob receives the correct message. If Eve is* active, *then, with probability $1 - 2^{-\Omega(k)}$, either Bob rejects and aborts the protocol, or Alice and Bob both accept and Bob receives the correct message m.*

[4] We say that a party *accepts the outcome* of a protocol if he has not rejected and aborted and the execution of the protocol is, or could be, finished from his point of view.

[5] Here, $f = \omega(g)$ stands for $f/g \to \infty$.

Remark. Note that k can be freely chosen subject to the conditions $k < tn/7$ and $k = \omega(\log n)$. Since the success probability of an active attack is bounded as $2^{-\Omega(k)}$, choosing a greater k is more secure; on the other hand, the smaller k is, the longer can the authenticated message be.

Proof of Theorem 1. We first observe that we can assume n to be a multiple of k if we replace, at the same time, the entropy condition by $H_\infty(S|U = u) > tn - k > 6k$. The reason is that Alice and Bob can cut at most $k - 1$ bits at the end of S, reducing the min-entropy by less than k according to Lemma 3.

Let now L be the greatest even integer such that $L \leq (tn - k)/3k$ holds, and let $l := \lfloor \log z(L) \rfloor = (1 - o(1))(tn - k)/3k$. Here, z is the function defined in Lemma 5; the maximum message length l is hence chosen such that there exists a one-to-one mapping from $\{0, 1\}^l$ to the set of L-bit strings satisfying the conditions of Lemma 5. If the length of the *actual* message m to be authenticated is shorter, then the number r of bits b_i in Protocol AUTH is smaller as well. Let in the following $b = (b_1, \ldots, b_r)$ be the r-bit string (where $r \leq L$ holds) which satisfies the conditions of Lemma 5 and corresponds to the message m.

Assume that Alice and Bob execute Protocol AUTH with respect to the key S, the parameter k, maximal message length L, and the string b. Let first Eve be passive. Then, clearly, Alice and Bob accept the outcome of the protocol and Bob receives the correct message, as sent by Alice.

Let now Eve be a possibly active adversary. Since neither Alice nor Bob generate and send responses for more than L challenges during the execution of the protocol (L is the maximum possible length of the string b; if Bob, for instance, is challenged for more than L times, he will conclude that there is an active attack, reject, and abort), we have at every point in the protocol that, for any $a \geq 0$,

$$\text{Prob}_C[H_\infty(S|U = u, C = c) \geq tn - k - 2Lk - 2La] \geq 1 - 2L2^{-a} \qquad (2)$$

holds, where C stands for the collection of all messages sent by Alice and Bob so far. Inequality (2) follows from $2L$-fold application of Lemma 2: At most $2L$ times, Eve has observed a string $f_x(s)$ (or $f_y(s)$) where X and S are, given Eve's entire knowledge at this point, independent. (It is important to see that the latter is true even if x is chosen by Eve herself, depending on all her knowledge. When applying Lemma 2 here, the distribution P_S in the statement of the lemma has to be replaced by $P_{S|U=u,C'=c'}$, where $U = u$ and $C' = c'$ summarize this knowledge.)

Since the total number of challenges generated by Alice or Bob in Protocol AUTH with maximal message length L is also upper bounded by $2L$, the probability $\text{Prob}[\mathcal{A}]$ of the event $\mathcal{A}$ that Eve can correctly answer one of them without help is, according to Lemma 1, inequality (2), and the union bound, at most

$$\text{Prob}[\mathcal{A}] \;\leq\; 2L\left(2^{-(tn-k-2L(k+a))/(n/k)} + \frac{n/k}{2^k}\right) + 2L2^{-a}$$

for any $a \geq 0$; the choice $a := (tn-k)/(12L)$ leads to

$$\begin{aligned}\text{Prob}\,[\mathcal{A}] &\leq 2\,\frac{tn-k}{3k}\left(2^{-(tn-k)/(6n/k)} + \frac{n/k}{2^k} + 2^{-(tn-k)/(12L)}\right) \\ &= O((n/k)^2)\cdot 2^{-\Omega(k)} \;=\; 2^{-\Omega(k)}\;, \end{aligned} \tag{3}$$

where the first "equality" in (3) holds because of $k < tn/7$, and the second one since k is of order $\omega(\log n)$.

Let us assume that $\mathcal{A}$ does not occur, and that Bob accepts the outcome of the protocol. Let $(b_1,\ldots,b_j)$, for $j \leq r$, be the bits authenticated in Alice's view. According to Lemma 4, the string $b' = (b'_1,\ldots,b'_{j'})$ that Bob receives can be obtained from $(b_1,\ldots,b_j)$ by inserting 0s and flipping 1s to 0s. Then, $w_H(b') = j'/2$—a necessary condition for Bob to accept—implies $w_H((b_1,\ldots,b_j)) \geq j'/2 \geq j/2$. Since b satisfies the conditions of Lemma 5, we have $w_H((b_1,\ldots,b_j)) = j/2$ and $j = r$ (i.e., Alice has sent the entire message already and hence accepts the outcome), as well as $j' = j = r$ and $b' = (b'_1,\ldots,b'_{j'}) = (b_1,\ldots,b_j) = (b_1,\ldots,b_r) = b$: Bob receives the correct string b and message m. □

3 Confidentiality from an Arbitrarily Weak Key

3.1 Privacy Amplification and Extractors

Privacy amplification means extracting a weakly secret string's randomness as seen from the adversary's viewpoint, and has been shown possible under the assumption that either the communication channel is authentic [4], [3], [5], or that the initial key's privacy level is already high [14], [24], [25], [17]. Here, we show that these two restrictions can be dropped simultaneously.

It was shown in [3] that—in the authentic-channel model—*universal hashing* is a good technique for privacy amplification, allowing for extracting virtually all the so-called *Rényi entropy* into a highly secret key. Another randomness-extraction technique, which has attracted a lot of attention recently in the context of derandomization of probabilistic algorithms, are *extractors*, which allow, by using only very few additional truly random bits, for extracting a weakly random source's complete *min-entropy*. The fact that extractors can distill only the min-entropy—instead of the Rényi entropy, which is a priori up to two times greater—is insignificant according to a recent unpublished result [11] stating that for every distribution P, there exists a distribution P' such that the variational distance between P and P' (both with range $\mathcal{X}$), defined as $d(P,P') := (\sum_{x\in\mathcal{X}} |P(x) - P'(x)|)/2$, is small and $H_\infty(P') \approx H_2(P)$ holds. This means that universal hashing is ultimately nothing else than a particular extractor—one which is, however, very inefficient with respect to the number of additional randomness required.

Using extractors, we show in Sections 3.2 and 3.3 that privacy amplification over an unauthenticated public channel can be (almost) as powerful—both with respect to the conditions for the *possibility in principle* and to the *length* of the resulting key—as over an authentic channel. When used for privacy amplification over a public channel, extractors should be *strong*, meaning that the output's distribution is close to uniform *even when given the truly random bits*. The existence of such extractors, distilling virtually all the source's min-entropy, was proven in [19], [20]. Theorem 2 is a direct consequence of these results.

Theorem 2. *For integers $D \leq n$ and a real number ε of order $\Omega(2^{-n/\log n})$ there exist $r = O(((\log n)^2 + \log(1/\varepsilon)) \log D)$, $m = D - 2\log(1/\varepsilon) - O(1)$, and a function $E : \{0,1\}^n \times \{0,1\}^r \rightarrow \{0,1\}^m$ with the following property. If X is a random variable with range $\mathcal{X} \subseteq \{0,1\}^n$ and such that $H_\infty(X) \geq D$ holds, and if U_r stands for a random variable independent of X and with uniform distribution over $\{0,1\}^r$, then we have $d\left(P_{(E(X,U_r),U_r)}, P_{U_{m+r}}\right) \leq \varepsilon$, where $P_{U_{m+r}}$ is the uniform distribution over $\{0,1\}^{m+r}$.*

A function E having the properties given in Theorem 2 is called a *strong* (D,ε)-*extractor*.

Lemma 6 below justifies the use of strong extractors for privacy amplification (both in the passive- and active-adversary cases, and with respect to our new simplified notion of security of privacy amplification given below): The extractor's output is, with high probability, equal to a perfectly uniformly distributed "ideal" key independent of the random bits: The adversary has no information at all about this ideal key.

Lemma 6. *Let E be a function as in Theorem 2 with parameters n, D, r, m, and ε, let S be a random variable with $H_\infty(S) \geq D$, and let $R = U_r$ be the random variable corresponding to a uniformly distributed r-bit string independent of S. Let $S' := E(S,R)$. Then there exists a uniformly distributed m-bit random variable S'_{id} which is independent of R and such that* $\text{Prob}\,[S' = S'_{id}] \geq 1 - \varepsilon$ *holds.*

Proof. For every value r_0 that R can take, there exists a uniformly distributed string $S'_{id}(r_0)$ with $\text{Prob}\,[E(S,r_0) \neq S'_{id}(r_0)] = d(P_{E(S,r_0)}, P_{U_m})$ (again, P_{U_m} is the uniform distribution over $\{0,1\}^m$).

Let S'_{id} be the random variable defined by all the $S'_{id}(r_0)$. The statement now follows from $d\left(P_{(E(S,R),R)}, P_{U_{m+r}}\right) = E_R\left[d\left(P_{E(S,R)}, P_{U_m}\right)\right]$, which is true since R is uniform, and from $\text{Prob}\,[E(S,R) \neq S'_{id}] = E_R[\text{Prob}\,[E(S,R) \neq S'_{id}(R)]]$. □

3.2 The Idea Behind Protocol PA

Given Protocol AUTH of Section 2 and the extractors described in Section 3.1, it seems obvious how to achieve privacy amplification over an unauthenticated channel: Alice chooses the extractor's random input bits and sends them, using Protocol AUTH, to Bob. However, this solution has a conceptual error. Eve can, knowing the bits $b_1, \ldots, b_i$ already sent, perform active attacks—replacing x_{i+1} or y_{i+1} by values x'_{i+1} and y'_{i+1} of her choice and therefore learning $f_{x'_{i+1}}(S)$ and $f_{y'_{i+1}}(S)$—and obtain information about S that *depends on the bits* b_j; in other words, the extractor's second input would in this case not be independent of S from Eve's viewpoint, and Theorem 2 would not apply.

On the other hand, however, the string S *must* be used *both* for authentication *and* as the input for privacy amplification. This dilemma can be resolved by a two-step protocol: First, an extractor is used to generate, from S, a short key S' about which all the information—depending on the extractor bits or not—revealed during Protocol AUTH gives Eve less than half the total information (except with small probability). In a second step, the random bits actually used to apply privacy amplification on S are authenticated with the key S'. The crucial point here is that the information Eve learns about S during this second authentication is at most S', since "the rest" of S is not used at all. Such a two-step approach hence allows for controlling the information Eve obtains during the two authentication phases, even when she is carrying out active attacks.

For the second authentication, a key—namely S'—about which Eve knows less than half (in terms of min-entropy) can be used. In this case, a simpler and more efficient method than Protocol AUTH can be applied, namely *strongly universal (SU-) hashing* [23]. A result similar to Lemma 7, but with respect to Rényi entropy H_2, was proven in [14].

Lemma 7. *Let n be an even integer. Assume that Alice and Bob both know the value taken by a random variable S with range $\mathcal{S} \subseteq \{0,1\}^n$ and conditional min-entropy $H_\infty(S|U=u) \geq n/2+R$, where $U=u$ summarizes an adversary's entire knowledge about S. Assume further that Alice and Bob use S to authenticate an $n/2$-bit message M, where M and S are independent, given $U=u$, with strongly universal hashing, i.e., with the authenticator $A = MS_1 + S_2$, where S_1 and S_2 are the first and second halves of S, and where M, S_1, S_2, and A are interpreted as elements of $GF(2^{n/2})$ with respect to a fixed basis of $GF(2^{n/2})$ over $GF(2)$. Then Bob always accepts and receives the correct message if the adversary is passive, and in general we have with probability $1-2^{-\Omega(R)}$ that Bob either rejects or receives the correct message M, even if the adversary has full control over the communication channel.*

Proof. If Eve is passive, then, clearly, Bob accepts and receives the correct message. In general, the probability of a successful *impersonation* attack is at most the maximum probability of a subset of $2^{n/2}$ keys, i.e., at most $2^{n/2} \cdot 2^{-H_\infty(S|U=u)} \leq 2^{-R}$. Let now (m,a) be the correctly authenticated message observed by Eve in a *substitution attack*. According to Lemma 2, applied to the distribution $P_{S|U=u}$, this pair is with probability $\geq 1-2^{-R/2}$ such that $H_\infty(S|M=m, A=a, U=u) \geq R/2$ holds. (Note that M is independent of S, given $U=u$.) Because generating another correct pair (m',a'), for $m' \neq m$, is equivalent to guessing S and because of the union bound, the success probability of this attack is upper bounded by $2^{-R/2} + 2^{-R/2} = 2^{-(R/2-1)}$. □

3.3 Asymptotically Optimal Privacy Amplification by Insecure Communication

We are now ready to give Protocol PA, and to prove our second main result. In the following, s is the n-bit key known to Alice and Bob about which Eve's information $U = u$ is limited by $H_\infty(S|U = u) \geq tn$, where t is an arbitrary constant with $0 < t \leq 1$. E_1 and E_2 are suitably chosen extractors.

Protocol PA (Privacy Amplification)

Alice		**Bob**
$s \in \{0,1\}^n$		$s \in \{0,1\}^n$
$r \in_r \{0,1\}^{\Theta(\sqrt{n})}$	AUTH(r) with $k = \Theta(\sqrt{n}/\log n)$ $\longrightarrow$	
$s' = E_1(s, r)$		$s' = E_1(s, r)$
$\mathrm{len}(s') = l' = \Theta(n/\log n)$ $s' = (s'_1, s'_2)$ $r' \in_r \{0,1\}^{l'/2}$ $a = s'_1 r' + s'_2$		
	r', a $\longrightarrow$	

		if a is incorrect:
		reject
		if a is correct:

$s'' = E_2(s, r')$ $\mathrm{len}(s'') = (1 - o(1))tn$		$s'' = E_2(s, r')$

In the first protocol phase, Alice authenticates the string r with Protocol AUTH (with parameter k).

Before stating the main result, we make a few remarks on the security achieved by Protocol PA. Note first that privacy amplification cannot be *guaranteed* to work in every case if Eve is assumed to have full control over the communication channel; the best one can hope for is that a possible active attack is detected.

A natural security definition would require that, with high probability, any party accepting the outcome of the protocol indeed has a highly secret key, and that if both parties accept, their keys are identical. (As already mentioned in the context of Protocol AUTH, one cannot demand for similar acceptance decisions of Alice and Bob.) Our protocol achieves even more than that: *If* Bob accepts, *then* everything went well (in particular, Alice also accepts).

The condition concerning the privacy of the resulting key stated and proven below is, although equivalent to, somewhat simpler than security definitions used previously in this context: Instead of giving a condition on the Shannon entropy of the resulting key, we consider privacy amplification successful if the generated key is, except with small probability, *equal to a perfectly secret ideal key.* This allows for simplifying the security proofs.

Theorem 3. *Assume that Alice and Bob know an n-bit string S satisfying $H_\infty(S|U = u) \geq tn$ for arbitrary $0 < t \leq 1$, where $U = u$ summarizes an* *adversary Eve's entire knowledge about S. Then Protocol PA allows, for suitable choices of the parameters and the extractors E_1 and E_2, for privacy amplification by communication over a completely insecure channel, distilling an arbitrarily large fraction of the min-entropy of S, given $U = u$, into a virtually secret string. More precisely, Protocol PA satisfies the following two conditions.*

1. *If Eve is* passive, *then Alice and Bob accept the outcome of the protocol and end up with the same string S'' of length $l'' = (1 - o(1))tn$ with the following property: There exists a string S''_{id} of the same length such that for all possible protocol communications $C = c$, $P_{S''_{id}|C=c,U=u}$ is the uniform distribution over the set of l''-bit strings, and such that* $\text{Prob}\,[S'' = S''_{id}] = 1 - 2^{-\Omega(n/(\log n)^2)}$ *holds.*
2. *If Eve is* active, *then the probability that either Bob rejects (and Alice either rejects as well or accepts and has computed a key S'' satisfying 1.), or that both Alice and Bob accept and that all the conditions of 1. hold, is of order $1 - 2^{-\Omega(\sqrt{n}/\log n)}$.*

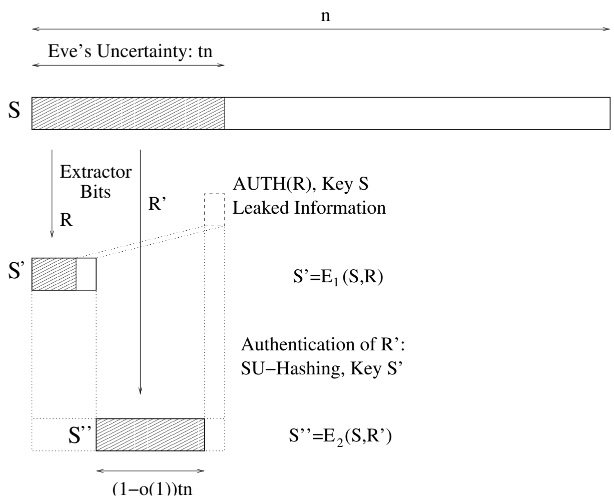

Fig. 1. Privacy amplification over an unauthenticated channel with Protocol PA. The privacy of S is extracted in two steps. The short key S' is more than half secret and used for authenticating the extractor bits for S''. This second key is the output of the protocol and highly secret although the information Eve obtains in the authentication depends on the extractor bits.

Proof. Note first that if Eve is passive, both parties accept the outcome and compute the same string S'', the secrecy of which, as stated in 1., follows from the subsequent analysis of the general case.

Let us hence assume that Eve is a possibly active adversary. Let $\varepsilon > 0$ be of order $2^{-\Theta(\sqrt{n}/\log n)}$. According to Theorem 2, there exist $r = \Theta(\sqrt{n})$, $k = \Theta(\sqrt{n}/\log n)$, and $l' = 7rk = o(n)$ as well as a strong extractor $E_1 : \{0,1\}^n \times \{0,1\}^r \to \{0,1\}^{l'}$ extracting l' bits out of S (distributed according to

$P_{S|U=u}$) with "error probability" ε. (Note that in fact, l' could be chosen much larger, namely almost tn, according to Theorem 2.)

The r randomly chosen bits R are now sent and authenticated using Protocol AUTH with parameter k. According to Theorem 1, the probability of a successful active attack to this authentication is $2^{-\Theta(k)} = 2^{-\Theta(\sqrt{n}/\log n)}$.

Let $S' = E_1(S, R)$ be the extractor's output. Because of Lemma 6, there exists an l'-bit string S'_{id} that is uniformly distributed conditioned on $U = u$, independent of the bits R (the second part of the extractor's input), and such that $\text{Prob}\,[S' = S'_{id}] \geq 1 - \varepsilon = 1 - 2^{-\Omega(\sqrt{n}/\log n)}$ holds.

Let C be all the messages sent by Alice and Bob during the execution of Protocol AUTH. Since every party sends at most $l = (1 + o(1))r$ messages (of length k) of the form $f_{x'_i}(s)$ or $f_{y'_i}(s)$ (and since the respective challenges x'_i and y'_i are, even if generated by Eve, independent of S, given Eve's knowledge about S at this moment), Lemma 2—applied $2l$ times—implies that

$$H_\infty(S'_{id}|C = c, U = u) \geq l' - 2lk - lk = l'/2 + \Omega(rk) \tag{4}$$

holds with probability at least $1 - 2l2^{-lk} = 1 - 2^{-\Omega(n/\log n)}$. The "equality" in (4) is true because l' has been defined to be equal to $7rk$, and because of $l = (1 + o(1))r$.

According to Lemma 7, the success probability of an active attack on the second authentication, using strongly universal hashing with the key S', is hence of order

$$2^{-\Omega(n/\log n)} + 2^{-\Omega(rk)} + 2^{-\Omega(\sqrt{n}/\log n)} = 2^{-\Omega(\sqrt{n}/\log n)}\ . \tag{5}$$

(The first term in (5) is the probability that (4) does not hold, the second term is the attack success probability if the key S'_{id} would be used and given that (4) holds, and the third term is the probability that the actually used key S' differs from S'_{id}. The bound (5) then follows from the union bound.)

Let us now look at the remaining min-entropy of S, given all the communication Eve has observed. Note first that the last authentication reveals information about S to Eve that depends on the random bits R' sent in this step. This dependence is a potential problem since R' must be chosen completely independently from S given Eve's knowledge and is, with respect to authentication with Protocol AUTH, the reason for the "two-step" nature of Protocol PA. However, under the (pessimistic) assumption that Eve learns the *entire* key S', she cannot obtain *any additional* information about S, in particular no information depending on R', since the rest of S is not used at all in this authentication. In other words, if we assume Bob to announce S' to Eve after the second authentication (what, of course, he does not actually have to do), then R' is independent of S given Eve's total knowledge.

We now have that $H_\infty(S|C = c, S' = s', U = u) \geq tn - \Theta(rk)$ holds with probability $1 - 2^{-\Omega(rk)} = 1 - 2^{-\Omega(n/\log n)}$, as above for S'. Because of Theorem 2, there exists a strong extractor $E_2 : \{0,1\}^n \times \{0,1\}^{r'} \to \{0,1\}^{l''}$ with parameters $r' \leq l'/2 = \Theta(n/\log n)$ (note that $l'/2$ is the possible message length in the last authentication), $\varepsilon' = 2^{-\Theta(n/(\log n)^2)}$, and $l'' = tn - \Theta(rk) - 2\log(1/\varepsilon') =$

$tn - o_1(n) - o_2(n) = (1-o(1))tn$. The extractor's output $S'' = E_2(S, R')$ satisfies, according to Lemma 6, the following condition. There exists an l''-bit string S''_{id} such that $P_{S''_{id}|C=c,U=u}$ is the uniform distribution (where C is the entire protocol communication) and $\text{Prob}\,[S'' = S''_{id}] \geq 1 - \varepsilon' = 1 - 2^{-\Omega(n/(\log n)^2)}$ holds. The final statement now follows from the union bound. □

4 Concluding Remarks

We have shown that two parties who are connected by a communication channel under full adversarial control and who share a key that is arbitrarily weakly secret can not only exchange authenticated messages, but also generate an unconditionally secret key. The given protocols for achieving this are computationally very efficient for the legitimate parties; they require two-way communication, where the number of rounds is of order $O(r)$ for the authentication protocol (if r is the length of the message to be authenticated) and $O(\sqrt{n})$ for privacy amplification of a weak n-bit secret. Clearly, the extracted highly secret key can then be used for all sorts of cryptographic tasks. The fact that unconditional security can be achieved even under assumptions as weak as that shows that this—most desirable—type of security might be more practical than generally assumed.

It is a natural question in this context whether such protocols can be given which even tolerate Alice's and Bob's initial strings to differ in a certain fraction of the positions (and how large this fraction can be). A positive answer to that would be useful in the context of quantum key agreement, for instance, since the usually-made assumption that the classical channel—used for the processing of the raw key—is authenticated, or that Alice and Bob share a short secret key already initially, could be dropped.

Acknowledgments

The authors thank Ueli Maurer for many interesting discussions, and Ronald Cramer as well as two anonymous reviewers for their helpful comments. The first author was supported by the Swiss National Science Foundation (SNF), and the second author by Canada's NSERC.

References

1. Y. Aumann, Y. Z. Ding, and M. O. Rabin, Everlasting security in the bounded storage model, *IEEE Trans. on Information Theory*, Vol. 48, pp. 1668–1680, 2002.
2. C. H. Bennett and G. Brassard, Quantum cryptography: public key distribution and coin tossing, *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing*, pp. 175–179, 1984.
3. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, Generalized privacy amplification, *IEEE Trans. on Information Theory*, Vol. 41, No. 6, pp. 1915–1923, 1995.

4. C. H. Bennett, G. Brassard, and J.-M. Robert, Privacy amplification by public discussion, *SIAM Journal on Computing*, Vol. 17, pp. 210–229, 1988.
5. C. Cachin, *Entropy measures and unconditional security in cryptography*, Ph. D. Thesis, ETH Zürich, Hartung-Gorre Verlag, Konstanz, 1997.
6. I. Csiszár and J. Körner, Broadcast channels with confidential messages, *IEEE Trans. on Information Theory*, Vol. 24, pp. 339–348, 1978.
7. Y. Dodis and J. Spencer, On the (non)universality of the one-time pad, *Proceedings of FOCS 2002*, 2002.
8. S. Dziembowski and U. M. Maurer, Tight security proofs for the bounded-storage model, *Proceedings of STOC 2002*, pp. 341–350, 2002.
9. W. Feller, *An introduction to probability theory and its applications*, 3rd edition, Vol. 1, Wiley International, 1968.
10. P. Gemmell and M. Naor, Codes for interactive authentication, *Advances in Cryptology - CRYPTO '93*, LNCS, Vol. 773, pp. 355–367, Springer-Verlag, 1993.
11. T. Holenstein, U. M. Maurer, and R. Renner, personal communication.
12. U. M. Maurer, Conditionally-perfect secrecy and a provably-secure randomized cipher, *Journal of Cryptology*, Vol. 5, No. 1, pp. 53–66, 1992.
13. U. M. Maurer, Secret key agreement by public discussion from common information, *IEEE Trans. on Information Theory*, Vol. 39, No. 3, pp. 733–742, 1993.
14. U. M. Maurer and S. Wolf, Privacy amplification secure against active adversaries, *Advances in Cryptology - CRYPTO '97*, LNCS, Vol. 1294, pp. 307–321, Springer-Verlag, 1997.
15. U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part I: Definitions and a completeness result, *IEEE Trans. on Information Theory*, Vol. 49, No. 4, pp. 822–831, 2003.
16. U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part II: The simulatability condition, *IEEE Trans. on Information Theory*, Vol. 49, No. 4, pp. 832–838, 2003.
17. U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part III: Privacy amplification, *IEEE Trans. on Information Theory*, Vol. 49, No. 4, pp. 839–851, 2003.
18. J. L. McInnes and B. Pinkas, On the impossibility of private key cryptography with weakly random keys, *Advances in Cryptology - CRYPTO '90*, LNCS, Vol. 537, pp. 421–436, Springer-Verlag, 1990.
19. R. Raz, O. Reingold, and S. Vadhan, Extracting all the randomness and reducing the error in Trevisan's extractors, *Proceedings of STOC '99*, pp. 149–158, 1999.
20. R. Raz, O. Reingold, and S. Vadhan, Error reduction for extractors, *Proceedings of FOCS '99*, pp. 191–201, 1999.
21. A. Russell and H. Wang, How to fool an unbounded adversary with a short key, *Advances in Cryptology - EUROCRYPT 2002*, LNCS, Vol. 2332, pp. 133–148, Springer-Verlag, 2002.
22. C. E. Shannon, Communication theory of secrecy systems, *Bell System Technical Journal*, Vol. 28, pp. 656–715, 1949.
23. D. R. Stinson, Universal hashing and authentication codes, *Advances in Cryptology - CRYPTO '91*, LNCS, Vol. 576, pp. 74–85, Springer-Verlag, 1992.
24. S. Wolf, Strong security against active attacks in information-theoretic secret-key agreement, *Advances in Cryptology - ASIACRYPT '98*, LNCS, Vol. 1514, pp. 405–419, Springer-Verlag, 1998.
25. S. Wolf, *Information-theoretically and computationally secure key agreement in cryptography*, ETH dissertation No. 13138, ETH Zürich, 1999.
26. A. D. Wyner, The wire-tap channel, *Bell System Technical Journal*, Vol. 54, No. 8, pp. 1355–1387, 1975.