

Secret-Key Agreement over Unauthenticated Public Channels – Part I: Definitions and a Completeness Result

Ueli Maurer, *Fellow, IEEE*

Stefan Wolf

Abstract— This is the first part of a three-part paper on secret-key agreement secure against active adversaries. In all three parts, we address the question whether two parties, knowing some correlated pieces of information X and Y , respectively, can generate a string S about which an adversary, knowing some information Z and having read and write access to the communication channel used by the legitimate partners, is almost completely ignorant. Whether such key agreement is possible, and if yes at which rate, is an inherent property of the joint probability distribution P_{XYZ} . In this part, we first prove a number of general impossibility results. We then consider the important special case where the legitimate partners as well as the adversary have access to the outcomes of many independent repetitions of a fixed tripartite random experiment. In this case, the result characterizing the possibility of secret-key agreement secure against active adversaries is of all-or-nothing nature: Either a secret key can be generated at the same rate as in the (well-studied) passive-adversary case, or such secret-key agreement is completely impossible. The exact condition characterizing the two cases is presented.

Keywords. Cryptography, unconditional security, secret-key agreement, authentication, typical sequences.

I. INTRODUCTION AND PRELIMINARIES

A. Motivation and Outline

One of the fundamental problems in cryptography is the generation of a secret key by two parties, Alice and Bob, not sharing such a key initially, in the presence of an adversary Eve who has access to the communication channel connecting Alice and Bob. Several scenarios, which differ in their assumptions about Eve's capabilities and possibly about the intractability of certain computational problems, have been considered in the literature.

Public-key cryptography introduced by Diffie and Hellman [8] solves this problem under the two assumptions that (1) Eve is unable to solve a certain computational problem (such as factoring integers or computing discrete logarithms in a given finite group) in feasible time, and (2) that Eve has only passive (read) access to the communication channel between Alice and Bob, i.e., that the communication between Alice and Bob is authenticated. The purpose of this paper is to investigate the described key-distribution problem when neither of these assumptions is made: We assume the presence of adversaries with infinite computing power and complete control over the

communication channel connecting Alice and Bob. Several impossibility results are proved, and some scenarios in which secret-key agreement secure against active adversaries *is* possible are characterized. First of all, secret-key agreement can be possible in this scenario only if Alice and Bob (but possibly also Eve) have correlated information to start with. More formally, while Alice and Bob share no secret key initially, they know some random variables X and Y , respectively, whereas the random variable Z is known to Eve. The joint probability distribution is denoted by P_{XYZ} .

One can have different opinions about whether it is reasonable to assume that a specific computational problem is difficult. Furthermore, since quantum computation has been invented as a (at least for now) theoretical model of computation, it is not completely clear whether intractability assumptions in the Turing machine model of computation are still adequate. There also exist different opinions about whether certain methods of authentication, like speaker identification on a voice channel, are strong enough to support the second assumption above. It is not a goal of this paper to discuss these issues, but we believe that avoiding both assumptions is an interesting and fundamental problem.

The outline of this paper is as follows. In Section I-B we describe different models of information-theoretic key agreement. In Sections I-C and I-D, two important techniques required later are described, namely typical sequences and almost strongly universal hashing for unconditionally secure authentication. In Section II, precise definitions of the security of secret-key agreement with respect to passive and active adversaries are given. Section III contains pessimistic impossibility results. In particular, the so-called *simulatability condition* is defined which leads to a general result of this kind. In the special case, studied in Section IV, where the parties have access to the outcomes of an independently repeated fixed random experiment, this condition is even shown to separate the complementary cases where *no* key agreement is possible at all, and where it is not only possible in principle but *at the same rate* as in the presence of only passive wire-tappers.

B. Secret-Key Agreement from Common Information by Public Discussion: Security Against Passive and Active Adversaries

Shannon’s pessimistic result on information-theoretic security states that if an adversary has perfect access to the ciphertext when a classical symmetric cryptosystem is used, then perfect secrecy can only be achieved if the parties share a secret key which is, roughly speaking, as long as the message to be secretly communicated [19]. Motivated by this, many scenarios have been considered in which the adversary’s information is in some sense limited. Such models can be based on noisy channels or on the laws of quantum mechanics. Examples of the former are the models by Wyner [22] or Csiszár and Körner [7], where the only assumption made is the presence of noisy communication channels from Alice to Bob and to an adversary Eve.

Maurer [12], and subsequently Ahlswede and Csiszár [1], described a more natural setting in which insecure communication between Alice and Bob is possible in both directions and not even regarded as a resource. In this interactive scenario, the parties Alice and Bob as well as the adversary Eve receive correlated pieces of information X , Y , and Z , respectively, i.e., realizations of random variables that are distributed according to a given joint distribution P_{XYZ} . (See Figure 1.) It was shown in [12] that Shannon’s pessimistic results on unconditional secrecy carry over to the interactive model.

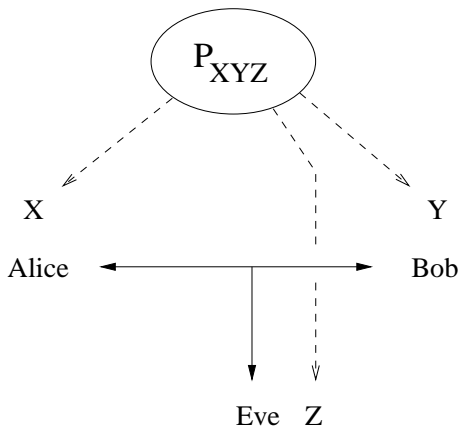


Fig. 1. Secret-Key Agreement by Public Discussion from Common Information

The case where the public communication channel is authentic (or equivalently, where the adversary is only a passive wire-tapper) was extensively studied [12], [15], [21], [11], [10]. In this paper we deal with the case where even this assumption is dropped, i.e., we assume that the adversary has full control over the communication channel. (See Figure 2.)

The main result of this (first part of the three-part) paper is a complete characterization—in the important special case where the parties’ knowledge stems from a large number of independent repetitions of a random experiment characterized by P_{XYZ} —of the possibility of secret-key




Fig. 2. Unconditional Security Against Active Opponents

agreement in terms of the passive-adversary case and of a specific property of P_{XYZ} , called non-simulatability. The result is of all-or-nothing nature: Either secret-key agreement against active adversaries is possible at the same rate as against only passive wire-tappers, or not possible at all.

C. Typical Sequences

Some of the proofs in this paper are based on arguments using typical sequences. Intuitively, a sequence of independent realizations of a random variable is called *typical* if the actual rate of occurrences of every specific outcome symbol is close to the probability of this symbol.

Definition 1: Let X be a random variable with distribution P_X and range \mathcal{X} , let $n > 0$ be an integer, and let $\gamma > 0$. A sequence $x^n = (x_1, x_2, \dots, x_n) \in \mathcal{X}^n$ is called (*strongly*) γ -*typical* if for all $a \in \mathcal{X}$, the actual number $N(a, x^n)$ of appearances of a in x^n satisfies

$$\left| \frac{N(a, x^n)}{n} - P_X(a) \right| \leq \frac{\gamma}{|\mathcal{X}|}.$$

It is a consequence of the law of large numbers that for every $\gamma > 0$, sufficiently long sequences of independent realizations of a random variable are typical with overwhelming probability. This fact is a very powerful tool in many contexts since it allows for reducing the proof of a statement about general distributions to the much simpler case of almost uniform distributions.

Theorem 1: [6], [4] Let $X^n = X_1 X_2 \dots X_n$ be a sequence of n independent realizations of the random variable X with distribution P_X and range $|\mathcal{X}|$, and let $0 < \gamma \leq 1/2$. Then

$$\text{Prob}[X^n \text{ is strongly } \gamma\text{-typical}] \geq 1 - (n+1)^{|\mathcal{X}|} \cdot 2^{-\frac{n}{2 \ln 2} \cdot \frac{\gamma^2}{|\mathcal{X}|^2}} = 1 - 2^{-\Omega(n\gamma^2)}.$$

D. Unconditionally Secure Authentication

Strongly universal classes of hash functions have been shown useful for unconditionally secure message authentication. The idea is that the secret key K determines a specific function h_K from the class, and that the authenticator of a message M is given by $h_K(M)$. However, the disadvantage when such hash functions are used is that the length of the required secret key cannot be much smaller than the length of the message to be authenticated. So-called ε -almost strongly universal classes of hash functions lead to a higher success probability of a substitution attack¹, but can be considerably smaller than strongly universal classes. Hence the required secret key can be shorter.

Definition 2: [20] Let $\varepsilon > 0$. A class H of functions from A to B is called ε -almost strongly universal (ε -ASU for short) if the following two conditions are satisfied.

1. For every $x \in A$ and $y \in B$, we have

$$|\{h \in H : h(x) = y\}| = \frac{|H|}{|B|},$$

2. for every $x_1, x_2 \in A$, $x_1 \neq x_2$, and for every $y_1, y_2 \in B$, we have

$$|\{h \in H : h(x_1) = y_1, h(x_2) = y_2\}| \leq \varepsilon \cdot \frac{|H|}{|B|}.$$

A class is called *strongly universal* if it is $1/|B|$ -ASU.

When using an ε -ASU class of functions for authentication as described above, then the success probability p_{imp} of an impersonation attack is $1/|B|$, whereas a substitution attack can be successful with probability $p_{sub} \leq \varepsilon$.

Theorem 2 shows that if p_{sub} can be tolerated to be greater than the optimal value $1/|B|$ by some factor i , then the length of the secret key can be made smaller by a factor of roughly $2^i/i$.

Theorem 2: [20] Let q be a prime power and let $i \geq 1$ be an integer. Then there exists an $((i+1)/q)$ -ASU class of q^{i+2} hash functions from A to B , where $|A| = q^{2i}$ and $|B| = q$.

II. DEFINITIONS

A. Secret and Authenticated Channels

Definition 3: Let U_b and U_a summarize an adversary Eve's entire knowledge before and after a message M is sent over the public discussion channel. This message is called *secret* (with respect to Eve) if

$$I(M; U_a | U_b) = 0.$$

A message M is called *authenticated* if

$$M = \overline{M},$$

¹In an *impersonation attack* to an authentication scheme, Eve tries to generate a correctly authenticated message without having seen such a message beforehand. In a *substitution attack* on the other hand, Eve waits until she sees a correctly authenticated message sent over the channel and tries to replace it by another correctly authenticated message.

where M stands for the message as sent by the sender and \overline{M} is the message as received by the receiver.

B. Definition of Key-Agreement Protocols

Definition 4: A protocol for secret-key agreement from common information consists of two phases: a communication phase and a key-generation phase.

During the *communication phase*, Alice and Bob exchange messages C_1, C_2, C_3, \dots over the public channel. It is assumed here that Alice sends $(C_1, C_3, \dots, C_{2k+1}, \dots)$, whereas Bob sends $(C_2, C_4, \dots, C_{2k}, \dots)$. A message C_i sent at some point during the protocol can (only) depend on the sender's knowledge when sending the message. More precisely, we have for odd i

$$H(C_i | X C_1 \overline{C_2} \overline{C_3} \dots \overline{C_{i-1}}) = 0, \quad (1)$$

whereas for even i

$$H(C_i | Y \overline{C_1} \overline{C_2} \overline{C_3} \dots \overline{C_{i-1}}) = 0 \quad (2)$$

holds. Here C_j stands for the j -th message as it is actually sent, and $\overline{C_j}$ for the same message as it is received (i.e., possibly modified by an active adversary if the communication channel is not assumed to be authentic). If t messages are exchanged in total during the communication phase, we denote by $C := (C_1, \dots, C_t)$ the list of all messages sent by the legitimate partners. The protocol is called *one-way-transmission protocol* if messages are sent only in one direction, i.e., if $C = (C_1)$ or $C = (C_2)$ holds.

In the subsequent *key-generation phase* of the protocol, Alice and Bob both decide (independently) whether they accept or reject the outcome of the protocol. In case of acceptance, a party computes a binary string S_A or S_B , respectively. More precisely, three different types of outcomes are possible for each party: to accept the outcome of the protocol and to compute a string S_A or S_B with

$$H(S_A | CX) = 0 \quad (3)$$

or

$$H(S_B | CY) = 0, \quad (4)$$

respectively, to reject but to compute a key S_A or S_B nevertheless, or to reject without computing a key.

Remark. We have assumed in Definition 4 that the protocol messages, the acceptance decisions, and the generated keys are determined by the knowledge of the corresponding parties and do not depend on additional random bits. This assumption does not restrict the generality because possibly required randomness can be assumed to be part of the random variables X and Y . Since adding independent randomness to X and Y does not change information-theoretic quantities such as $I(X; Y)$ or $I(X; Y | Z)$, all the results below also hold for probabilistic protocols.

Definition 5: Assume first that all the messages sent over the public channel are authentic but not secret. Let r be an integer, and let $\varepsilon > 0$. A $(P_{XYZ}, r, \varepsilon)$ -protocol has the following property.² First, it is required that both Alice

²Here and in the rest of this chapter, all probabilities are taken over the random variables X , Y , and Z with joint distribution P_{XYZ} .

and Bob accept the outcome of the protocol (and hence compute r -bit strings S_A and S_B , respectively). Furthermore, there must exist a perfectly uniform r -bit string S (i.e., $H(S) = r$) such that

$$\text{Prob}[S_A = S_B = S] \geq 1 - \varepsilon \quad (5)$$

and

$$H(S|ZC) \geq r - \varepsilon \quad (6)$$

hold.³

Assume now that the messages exchanged are neither authenticated nor secret. Let r be an integer, and let $\varepsilon, \delta > 0$. A *robust* ($P_{XYZ}, r, \varepsilon, \delta$)-protocol has the following properties.

1. *Correctness and Privacy.* If Eve is a passive wire-tapper, then the probability that both Alice and Bob accept at the end of the protocol and that secret-key agreement has been successful must be at least $1 - \delta$. Here, secret-key agreement is called *successful* if there exists a perfectly uniformly distributed r -bit string S such that (5) and (6) hold.

2. *Robustness.* For every possible strategy of Eve, the probability that either *both* Alice and Bob reject the outcome of the protocol, or secret-key agreement has been successful, must be at least $1 - \delta$.

Remark. Note that we do not require that Alice and Bob simultaneously accept the outcome in case of successful key agreement in the presence of an active adversary. The reason is that such a perfect synchronization cannot be achieved (see Section III-C).

III. IMPOSSIBILITY RESULTS

A. An Upper Bound on the Key Size

For the authenticated-communication case (equivalently, if the adversary is only passive), upper bounds on the size of the generated key have been proven in [12] and [15].

Definition 6: [15] For a distribution P_{XYZ} , the *intrinsic conditional mutual information between X and Y when given Z* , denoted by $I(X; Y \downarrow Z)$, is

$$I(X; Y \downarrow Z) :=$$

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

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

Theorem 3 below leads to an upper bound on the size of a secret key that can be generated by a secret-key agreement protocol. The proof of Theorem 3 is subdivided into a few steps stated as lemmas. We first define what it means that a random variable C can be generated from X and Y by a deterministic protocol.

Definition 7: Let P_{XY} be the joint distribution of two discrete random variables X and Y . Then a random variable C , jointly distributed with X and Y according to

³Definition 5 is slightly stronger than earlier definitions, e.g., as given in [12]. The new, more natural definitions, however, can be shown to be equivalent to the previous ones [16].

P_{XYC} , can be generated by (deterministic) communication from X and Y if there exist $t \in \mathbf{N}$ and random variables C_1, C_2, \dots, C_t , distributed according to $P_{XYC_1C_2 \dots C_t}$, with the following properties.

1. For all odd i , $1 \leq i \leq t$, we have

$$H(C_i | XC_1 \dots C_{i-1}) = 0,$$

whereas for even i ,

$$H(C_i | YC_1 \dots C_{i-1}) = 0$$

holds;

2. there exists a bijection φ between the ranges of C and $[C_1, C_2, \dots, C_t]$ such that

$$\text{Prob}[\varphi(C) = [C_1, C_2, \dots, C_t]] = 1$$

holds.

Theorem 3: Let X, Y, Z, C, S_A, S_B , and S be discrete random variables such that C can be generated by deterministic communication from X and Y , and with $H(S_A|XC) = H(S_B|YC) = 0$. Then we have

$$H(S) \leq H(S|S_A) + H(S|S_B) + I(X; Y \downarrow Z) + I(S_A S_B; ZC). \quad (7)$$

As a preparation for the proof of Theorem 3, we first show that Alice and Bob cannot, by public communication, increase the mutual information shared between them (and conditioned on the adversary's knowledge).

Lemma 1: Let P_{XY} be the joint distribution of two random variables X and Y , and let C be generated by deterministic communication from X and Y . Then we have for all conditional distributions $P_{Z|XY}$

$$I(XC; YC | ZC) = I(X; Y | ZC) \leq I(X; Y | Z).$$

Remark. The first equality holds for arbitrary random variables X, Y, Z , and C . It may be somewhat surprising that the condition given in Lemma 1 is *not sufficient* for C being a (deterministic) communication, as the following example shows. Let $\mathcal{X} = \mathcal{Y} = \{1, 2, 3\}$, $P_{XY}(i, j) = 1/9$ for all $(i, j) \in \{1, 2, 3\}^2$, and

$$\begin{aligned} P_{C|XY}(1, 1, 1) &= P_{C|XY}(1, 1, 2) = \\ P_{C|XY}(2, 1, 3) &= P_{C|XY}(2, 2, 3) = \\ P_{C|XY}(3, 3, 3) &= P_{C|XY}(3, 3, 2) = \\ P_{C|XY}(4, 3, 1) &= P_{C|XY}(4, 2, 1) = \\ P_{C|XY}(5, 2, 2) &= 1. \end{aligned}$$

The random variable C , as a function of (X, Y) , is represented in the following table.

We have, for all $P_{Z|XY}$,

$$0 = I(X; Y | ZC) \leq I(X; Y | Z),$$

although C (which depends deterministically on (X, Y)) cannot be generated by deterministic communication from X and Y .

C		X		
		1	2	3
Y	1	1	1	2
	2	4	5	2
	3	4	3	3

Fig. 3. The function $C(X, Y)$

Proof of Lemma 1. Clearly, we can show the statement for the random variable $[C_1, \dots, C_i]$ (as in Definition 7) instead. We show that for all i ,

$$I(X; Y|ZC_1 \cdots C_i) \leq I(X; Y|ZC_1 \cdots C_{i-1})$$

holds. This implies the statement. Let i be odd, i.e., $H(C_i|XC_1 \cdots C_{i-1}) = 0$. (The proof for even i is analogous, where X and Y have to be interchanged.) Then

$$\begin{aligned} I(X; Y|ZC_1 \cdots C_i) &= H(Y|ZC_1 \cdots C_i) - \underbrace{H(Y|XZC_1 \cdots C_i)}_{=H(Y|XZC_1 \cdots C_{i-1})} \\ &\leq H(Y|ZC_1 \cdots C_{i-1}) - H(Y|XZC_1 \cdots C_{i-1}) \\ &= I(X; Y|ZC_1 \cdots C_{i-1}). \end{aligned}$$

□

Lemma 2: Let A , B , and C be arbitrary discrete random variables. Then

$$H(A) \leq H(A|B) + H(A|C) + I(B; C).$$

Proof. We have $H(A|B) + H(A|C) + I(B; C) - H(A) = I(B; C|A) + H(A|BC) \geq 0$. □

Lemma 3: Let A , B , and C be arbitrary discrete random variables. Then

$$I(A; B) \leq I(A; B|C) + I(AB; C).$$

Proof. We have $I(A; B|C) + I(AB; C) - I(A; B) = I(A; C|B) + I(B; C|A) \geq 0$. □

Proof of Theorem 3. Let $P_{\bar{Z}|Z}$ be an arbitrary discrete conditional distribution, and let \bar{Z} be generated by sending Z over this channel. Then we have

$$\begin{aligned} H(S) &\leq H(S|S_A) + H(S|S_B) + I(S_A; S_B) \\ &\leq H(S|S_A) + H(S|S_B) \\ &\quad + I(S_A; S_B|\bar{Z}C) + I(S_A S_B; \bar{Z}C) \\ &\leq H(S|S_A) + H(S|S_B) \\ &\quad + I(XC; YC|\bar{Z}C) + I(S_A S_B; ZC) \\ &\leq H(S|S_A) + H(S|S_B) \\ &\quad + I(X; Y|\bar{Z}) + I(S_A S_B; ZC). \end{aligned}$$

The four inequalities hold because of Lemma 2, Lemma 3, the data-processing lemma [6], and Lemma 1, respectively.

Since $P_{\bar{Z}|Z}$ was an arbitrary discrete channel, we have

$$\begin{aligned} H(S) &\leq H(S|S_A) + H(S|S_B) \\ &\quad + I(X; Y|Z) + I(S_A S_B; ZC), \end{aligned}$$

and this concludes the proof. □

Corollary 4: Assume that a $(P_{XYZ}, r, \varepsilon)$ -protocol exists. Then

$$r < \frac{I(X; Y|Z) + 3h(\varepsilon) + \varepsilon}{1 - 4\varepsilon}.$$

Lemma 4: Let A and B be discrete random variables, let \mathcal{E} be an event, and let $\bar{\mathcal{E}}$ be the complementary event of \mathcal{E} . Then we have

$$\begin{aligned} I(A; B) &\leq h(\text{Prob}[\mathcal{E}]) + \text{Prob}[\mathcal{E}] \cdot I(A; B|\mathcal{E}) \\ &\quad + (1 - \text{Prob}[\mathcal{E}]) \cdot I(A; B|\bar{\mathcal{E}}). \end{aligned}$$

Proof. Let C be the random variable indicating whether the event \mathcal{E} occurs ($C = 1$) or not ($C = 0$). Then the statement to be proven translates to

$$\begin{aligned} I(A; B) &\leq H(C) + P_C(0) \cdot I(A; B|C = 0) \\ &\quad + P_C(1) \cdot I(A; B|C = 1) \\ &= H(C) + I(A; B|C). \end{aligned}$$

This is true because of $H(C) + I(A; B|C) - I(A; B) = I(A; C|B) + I(B; C|A) + H(C|AB) \geq 0$. □

Proof of Corollary 4. From Theorem 3 and Lemma 4, we can conclude that

$$\begin{aligned} r &= H(S) \\ &\leq H(S|S_A) + H(S|S_B) + I(X; Y|Z) \\ &\quad + I(S_A S_B; ZC) \\ &\leq 2(h(\varepsilon) + \varepsilon r) + I(X; Y|Z) + h(\varepsilon) \\ &\quad + (1 - \varepsilon)\varepsilon + \varepsilon \cdot 2r \\ &< I(X; Y|Z) + 3h(\varepsilon) + 4\varepsilon r + \varepsilon. \end{aligned}$$

□

B. Key Agreement Without Joint Randomness

In this section we consider the special case where no joint randomness is given to the involved parties or, equivalently, where X and Y are independent. The results below demonstrate an interesting difference between computational and information-theoretic cryptography. In both models a secret channel from Alice to Bob can be transformed into an authenticated channel from Bob to Alice. This is achieved by Alice sending a secret key to Bob and Bob using the key in a message authentication technique for authenticating a message to be sent to Alice.

In sharp contrast, only the computational model allows for transforming an authenticated channel from Alice to Bob into a secret channel from Bob to Alice. This is achieved by Alice sending her public key for a public-key cryptosystem to Bob who uses it to encrypt the message to be sent secretly to Alice. The security of public-key cryptosystems is inherently bound to be computational rather

than information-theoretic. (In fact, this follows from Theorem 5 below.) It is not surprising that in the information-theoretic model, when Alice and Bob have no common information initially, authenticated channels are of no use, in contrast to secret channels.

Theorem 5: Let P_{XYZ} be such that X and Y are independent, let $\varepsilon > 0$ and $r > (3h(\varepsilon) + \varepsilon)/(1 - 4\varepsilon)$. Then there exists no $(P_{XYZ}, r, \varepsilon)$ -protocol. Moreover, there exists no robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol for any $\delta < 1/2$.

Proof. The first part follows from Corollary 4 and from $I(X; Y \downarrow Z) = 0$.

To prove the second part, note that from Bob's point of view, Alice has no advantage compared to Eve. If Eve performs the same protocol as Alice would, pretending to be Alice, Bob accepts with the same probability as he would accept a protocol execution with Alice, which is at least $1 - \delta$ according to Definition 5. The condition $1 - \delta \leq \delta$ is not satisfied for any $\delta < 1/2$. \square

Theorem 5 is pessimistic: it states that information-theoretically secure secret-key agreement against active or passive adversaries is impossible to achieve when the channel between Alice and Bob is completely insecure. However, if Alice and Bob have correlated information initially (not necessarily a secret key, but merely bit strings that are somehow correlated)—information about which also Eve has partial knowledge—then secret-key agreement can be possible, as was shown in [12] for the passive-adversary model, and as we will see in Section IV for the case of active adversaries.

As mentioned, the first statement of Theorem 5 is in sharp contrast to the public-key-cryptographic scenario. The following well-known result is an observation following from Theorem 5.

Corollary 6: A public-key cryptosystem cannot be information-theoretically secure.

Theorem 7: Let P_{XYZ} be such that X and Y are independent. Assume that one secret (but not necessarily authenticated) message can be sent from Alice to Bob. Then, for any r and $\delta > 0$, a robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol exists if, in addition, either an authenticated message can be sent from Alice to Bob or a secret message can be sent from Bob to Alice.

Proof. Note that when a message from Alice to Bob is simultaneously secret and authenticated, then Alice can simply send a secret key as the message. When two messages can be sent from Alice to Bob, one secret and one authenticated, then Alice can send a random n -bit string R to Bob⁴ ($n \geq -2 \log \delta$) over the secret channel and the description of a function f in a universal class hash functions from $\{0, 1\}^n$ to $\{0, 1\}^n$ [5] over the authenticated channel, together with the first $n/2$ bits of $f(R)$. The other half of $f(R)$ is kept by Alice and Bob as their secret key. If Eve's capability to interfere with the secret channel is limited to sending fraudulent messages (but she is assumed to be unable to modify a message sent from Alice to Bob), then

⁴All logarithms in this paper are binary.

no universal hash function is needed; it could instead be replaced by the identity function.

The proof for the case of a secret channel from Bob to Alice is based on the following protocol. Bob (secretly) sends Alice a random string U of sufficient length ($\Omega(\log(1/\delta))$). Then they use the above protocol, where the authenticated channel is obtained by Alice by using an authentication scheme [9] with U as the secret key. \square

C. Perfect Synchronization is Impossible

Of course it would be most desirable to use robust protocols for which, with high probability, Alice and Bob either both accept (and secret-key agreement is successful) or both reject. However, the following theorem states that such a synchronization cannot be achieved. Hence the given definition of robustness against active attacks appears to be the strongest achievable.

Theorem 8: Assume that a robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol exists satisfying even the modified robustness property that with probability at least $1 - \delta$, either both Alice and Bob reject, or both parties accept and secret-key agreement has been successful. Then suitable strings can be computed even without communication, i.e., there exist two functions f and g , mapping \mathcal{X} and \mathcal{Y} to $\{0, 1\}^r$, respectively, such that $S_A := f(X)$ and $S_B := g(Y)$ satisfy (5) and (6), for some S , with probability at least $1 - 3\delta$.

Proof. We assume that such a protocol exists although suitable S_A and S_B cannot be directly obtained from X and Y , respectively, with probability $1 - 3\delta$. Let us consider a particular execution of the protocol (i.e., fixed values of X , Y , and Z), where we assume first that Eve is passive. Let C_1, \dots, C_t be the messages sent during this execution. Let for all $2i + 1 < t$ $\text{state}_A(2i + 1)$ be the state (**accept** or **reject**) of Alice after sending the message C_{2i+1} , whereas for $2j < t$, $\text{state}_B(2j)$ denotes Bob's state after sending C_{2j} . More precisely, these are the final states of the parties, provided that no further messages are received. Let analogously $\text{state}_A(t)$ and $\text{state}_B(t)$ be the final states of Alice and Bob in this execution. Because Eve is passive, we must have that $\text{state}_A(t) = \text{state}_B(t) = \text{accept}$ with probability at least $1 - \delta$ by the protocol definition. Because suitable S_A and S_B cannot be computed directly from X and Y , respectively, except with probability smaller than $1 - 3\delta$, at least one message *must* be sent in the protocol or in other words, $\text{state}_B(0) = \text{reject}$, with probability greater than $1 - \delta - (1 - 3\delta) = 2\delta$.

From $\text{state}_B(0) = \text{reject}$ and $\text{state}_B(t) = \text{accept}$, an event that occurs with probability greater than

$$1 - (1 - 2\delta) - \delta = \delta$$

(by the union bound for the complementary event), we conclude that there exists j such that $\text{state}_B(2j) = \text{reject}$ and $\text{state}_B(2j + 2) = \text{accept}$ (or $\text{state}_B(t - 1) = \text{reject}$ and $\text{state}_B(t) = \text{accept}$). Hence we have either $\text{state}_B(2j) \neq \text{state}_A(2j + 1)$ or $\text{state}_B(2j + 2) \neq \text{state}_A(2j + 1)$. If Eve blocks every message sent over the

channel after C_{2j} —or C_{2j+1} in the second case—then Alice and Bob end up in opposite states. (In the case where $\text{state}_B(t-1) = \text{reject}$ and $\text{state}_B(t) = \text{accept}$, Eve can block the last message C_t .) We conclude that Eve can achieve disagreement in Alice’s and Bob’s acceptance states with probability greater than δ , which is a contradiction to the protocol definition. \square

D. The Necessity of Two-Way Communication

In many cases, secret-key agreement protocols secure only against passive adversaries can be one-way. An example is secret-key agreement in cases where Alice and Bob have an initial advantage over Eve, in particular privacy amplification [3], [2], [18]. In this section, however, we show that the communication in protocols secure against active adversaries must be two-way. This is not surprising since the party sending the first message must be protected from accepting when this message is deleted or modified.

Theorem 9: Assume that a robust $(P_{XYZ}, r, \varepsilon, \delta)$ -one-way-transmission protocol exists. Then suitable strings can be computed even without communication with probability at least $1 - 2\delta$.

Proof. Note first that because Alice never receives a message, her decision is independent of Eve’s strategy. The probability that Alice does not accept at the end is at most δ , even if Eve deletes the message sent. On the other hand, assume that the probability that key agreement is successful without any communication is smaller than $1 - 2\delta$. Then the probability that Alice accepts and no secret key can be computed nevertheless is greater than δ (by the union bound), which is a contradiction to the protocol definition. \square

E. The Simulatability Condition

Secret-key agreement secure against active adversaries can only be possible if Alice and Bob have some advantage over Eve in terms of the distribution P_{XYZ} . More precisely, this advantage must be such that Eve cannot generate from Z a random variable \bar{X} which Bob, knowing Y , is unable to distinguish from X (and vice versa).

Definition 8: [13] Let X, Y , and Z be random variables. Then X is simulatable by Z with respect to Y , denoted

$$\text{sim}_Y(Z \rightarrow X),$$

if there exists a conditional distribution $P_{\bar{X}|Z}$ such that $P_{\bar{X}Y} = P_{XY}$, where

$$P_{\bar{X}Y}(x, y) = \sum_{z \in \mathcal{Z}} P_{YZ}(y, z) \cdot P_{\bar{X}|Z}(x, z).$$

Another way of stating that $\text{sim}_Y(Z \rightarrow X)$ holds is that there exists a random variable \bar{X} such that $I(\bar{X}; Z|XY) = 0$, i.e., $XY \rightarrow Z \rightarrow \bar{X}$ is a Markov chain, and $P_{\bar{X}Y} = P_{XY}$ holds. In the second part [17] of this three-part paper we describe simple criteria for non-simulatability in terms of the probabilities $P_{XYZ}(x, y, z)$. The following theorem states that a robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol can only exist

if both X and Y are not simulatable by Z with respect to each other. In the scenario where the parties obtain the outcomes of repeated realizations of a fixed random experiment, this condition is also sufficient for the possibility of key agreement (see Section IV).

Theorem 10: Let X, Y , and Z be random variables with distribution P_{XYZ} , and let

$$r > \frac{h(2\delta + \varepsilon) - h(\varepsilon) + (1 - 2\delta)\varepsilon}{1 - 4\delta}.$$

Then if either $\text{sim}_Y(Z \rightarrow X)$ or $\text{sim}_X(Z \rightarrow Y)$ holds, there exists no robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol.

Proof. We assume that $\text{sim}_Y(Z \rightarrow X)$ holds, and that a robust $(P_{XYZ}, r, \varepsilon, \delta)$ -protocol exists nevertheless. Because Bob cannot distinguish between Alice and an impersonating Eve, the probability that he accepts and that there was no interaction with Alice is at least $1 - \delta$ (i.e., the same as for the passive-adversary case). On the other hand, the probability that Bob accepts and that key agreement was not successful must be upper bounded by δ . Hence, by the union bound, the probability that Bob accepts and that secret-key agreement has been successful even without interaction with Alice, is at least $1 - 2\delta$. We conclude that Alice (hence also Eve, who can simulate Alice towards Bob) and Bob can compute strings S_A, \bar{S}_A , and S_B from X, Z , and Y , respectively, such that

$$\text{Prob}[\bar{S}_A \neq S_B] \leq 2\delta + \varepsilon,$$

hence

$$H(S_B|\bar{S}_A) \leq h(2\delta + \varepsilon) + (2\delta + \varepsilon)r \quad (8)$$

by Fano’s inequality (see [6]). On the other hand, we have

$$\begin{aligned} H(S_B|\bar{S}_A) &\geq H(S_B|Z) \\ &\geq H(S|ZC) + h(\varepsilon) + \varepsilon r \\ &\geq (1 - 2\delta)(r - \varepsilon) + h(\varepsilon) + \varepsilon r \end{aligned} \quad (9)$$

by the protocol definition. It is not difficult to verify that the inequalities (8) and (9) together imply

$$r \leq \frac{h(2\delta + \varepsilon) - h(\varepsilon) + (1 - 2\delta)\varepsilon}{1 - 4\delta},$$

and this contradiction concludes the proof. \square

IV. THE ROBUST SECRET-KEY RATE AND A COMPLETENESS RESULT

In the following, we denote by P_{XYZ}^n the distribution over $\mathcal{X}^n \times \mathcal{Y}^n \times \mathcal{Z}^n$ that corresponds to n independent realizations of the random experiment characterized by P_{XYZ} .

Definition 9: The *secret-key rate* $S(X; Y||Z)$ is the least upper bound of the set of numbers $R \geq 0$ with the property that for all $\varepsilon > 0$, and for sufficiently large n , there exists a $(P_{XYZ}^n, \lfloor (R - \varepsilon)n \rfloor, \varepsilon)$ -protocol.

The *robust secret-key rate* $S^*(X; Y||Z)$ is the least upper bound of the set of numbers $R \geq 0$ with the property that for all $\varepsilon, \delta > 0$, and for sufficiently large n , there exists a robust $(P_{XYZ}^n, \lfloor (R - \varepsilon)n \rfloor, \varepsilon, \delta)$ -protocol.

Clearly,

$$S^*(X; Y|Z) \leq S(X; Y|Z)$$

holds for all distributions P_{XYZ} . Theorem 11 expresses $S^*(X; Y|Z)$ in terms of $S(X; Y|Z)$ and P_{XYZ} , and corrects a result of [13], where the same criterion was given in connection with the weaker protocol definition. The result of Theorem 11 is of all-or-nothing nature: $S^*(X; Y|Z)$ is either equal to $S(X; Y|Z)$ or to 0, depending on whether either one of X or Y is simulatable by Z .

Theorem 11: Let P_{XYZ} be a distribution with $S(X; Y|Z) > 0$. Then $S^*(X; Y|Z) = 0$ holds if either $\text{sim}_Y(Z \rightarrow X)$ or $\text{sim}_X(Z \rightarrow Y)$ is true. Otherwise, we have $S^*(X; Y|Z) = S(X; Y|Z)$.

Before proving Theorem 11 we show two lemmas. Lemma 5 states that whenever there exists a (passive-adversary) key-agreement protocol achieving some key-generation rate, then there exists even a protocol with a constant number of communication rounds and linear message length.

Lemma 5: Let P_{XYZ} be a distribution, and let $S(X; Y|Z)$ be its secret-key rate. Then there exists, for every $\varepsilon > 0$ and for sufficiently large $N \geq N_0(\varepsilon)$, a protocol for key agreement, with respect to the parameter N , that requires only a constant number of communication rounds, and where all the messages sent are of length $O(N)$.

Proof. The protocol achieving the stated key-generation rate and communication complexities was described in [16]. A detailed analysis is given there. The idea is as follows. The low-communication-complexity protocol consists of two phases. In a first phase, the (general) key-agreement protocol is repeated many times independently for some fixed N' . In the second phase, information reconciliation (i.e., error correction) and privacy amplification are applied to the concatenation of the generated keys. The key arguments of the analysis are based on typical sequences. \square

Lemma 6 shows how blocks of realizations of the random variables X and Y can directly be used for identification and message authentication if Eve cannot, given Z , simulate X with respect to Y , or Y with respect to X .

Lemma 6: Let P_{XYZ} be the joint distribution of three random variables X , Y , and Z , with the property that $\text{sim}_X(Z \rightarrow Y)$ does *not* hold. Then there exists, for every $\delta > 0$, an integer $N_0(\delta)$ of order $O(\log(1/\delta))$ such that for all $N \geq N_0(\delta)$ there exists a function

$$f : \mathcal{X}^N \times \mathcal{Y}^N \longrightarrow \{0, 1\}$$

with the following two properties.

1. $\text{Prob}[f(X^N, Y^N) = 1] \geq 1 - \delta$,
2. for every random variable W^N with range \mathcal{Y}^N (not necessarily consisting of N independent repetitions of a random variable with range \mathcal{Y}) for which

$$X^N Y^N \longrightarrow Z^N \longrightarrow W^N$$

is a Markov chain,

$$\text{Prob}[f(X^N, W^N) = 0] \geq 1 - \delta$$

holds.

Proof. Let $\delta > 0$. For every real number $\gamma > 0$ and integer $N \geq 1$, we define the function $f_{N,\gamma} : \mathcal{X}^N \times \mathcal{Y}^N \longrightarrow \{0, 1\}$ as follows:

$$f_{N,\gamma}(x^N, y^N) = \begin{cases} 0 & \text{if } (x^N, y^N) \text{ is not} \\ & \gamma\text{-typical} \\ 1 & \text{otherwise.} \end{cases}$$

We conclude from Theorem 1 that there exists $N_0(\delta) = \mathcal{O}(\log(1/\delta))$ such that for all $N \geq N_0(\delta)$,

$$\text{Prob}[f_{N,\gamma}(X^N, Y^N) = 1] \geq 1 - \delta$$

holds.

We prove the second part of the statement by contradiction. We assume that there exists $\delta > 0$ such that for arbitrarily large N and for all $\gamma > 0$, the function $f_{N,\gamma}$ does *not* satisfy the second condition. Thus there exist sequences

$$N_n \rightarrow \infty, \quad \gamma_n \rightarrow 0 \quad (\gamma_n > 0), \quad \text{and} \quad W^{N_n},$$

where W^{N_n} is a random variable with range \mathcal{Y}^{N_n} (but again, which does, in contrast to X^{N_n} , Y^{N_n} , and Z^{N_n} , not necessarily consist of N_n independent realizations of a random variable with range \mathcal{Y}), such that

$$X^{N_n} Y^{N_n} \longrightarrow Z^{N_n} \longrightarrow W^{N_n} \quad (10)$$

is a Markov chain and

$$\begin{aligned} \text{Prob}[X^{N_n} W^{N_n} \text{ is a } \gamma_n\text{-typical } P_{XY}\text{-sequence}] \\ > \delta \end{aligned} \quad (11)$$

holds. Because (10) is a Markov chain, and because of Theorem 1, N_n and γ_n can be chosen such that for all n , the following two conditions are simultaneously satisfied with probability at least $1 - \delta/2$. First, $X^{N_n} Z^{N_n}$ is a γ_n -typical P_{XZ} -sequence. Secondly, let for $b \in \mathcal{Y}$ and $c \in \mathcal{Z}$

$$F_n(b, c) := \frac{N((c, b), (z^{N_n}, w^{N_n}))}{N(c, z^{N_n})}.$$

Note that F_n is a conditional probability distribution on the set $\mathcal{Y} \times \mathcal{Z}$. Then the second condition is that for all $(a, b, c) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$,

$$\left| \frac{N((a, c, b), (x^{N_n}, z^{N_n}, w^{N_n}))}{N((a, c), (x^{N_n}, z^{N_n}))} - F_n(b, c) \right| < \gamma_n$$

holds. By the union bound, all the above events together occur with probability at least $\delta/2$ (> 0). In this case we have

$$\begin{aligned} \frac{N((a, b), (x^{N_n}, w^{N_n}))}{N_n} &= \\ \sum_{c \in \mathcal{Z}} \frac{N((a, c), (x^{N_n}, z^{N_n}))}{N_n} & \cdot \\ \frac{N((a, c, b), (x^{N_n}, z^{N_n}, w^{N_n}))}{N((a, c), (x^{N_n}, z^{N_n}))} & , \end{aligned}$$

and thus

$$\left| \sum_{c \in \mathcal{Z}} P_{XZ}(a, c) \cdot F_n(b, c) - P_{XY}(a, b) \right| \leq \frac{\gamma_n}{|\mathcal{X}| \cdot |\mathcal{Z}|} + \gamma_n + \frac{\gamma_n^2}{|\mathcal{X}| \cdot |\mathcal{Z}|} + \frac{\gamma_n}{|\mathcal{X}| \cdot |\mathcal{Y}|}. \quad (12)$$

Since the set of conditional probability distributions on the finite set $\mathcal{Y} \times \mathcal{Z}$ is compact in $\mathbf{R}^{(|\mathcal{Y}|-1) \cdot |\mathcal{Z}|}$, the sequence $F_n(\cdot, \cdot)$ in this set has a convergent subsequence $F_{n'}$. Let

$$P_{\overline{Y}|Z} := \lim_{n' \rightarrow \infty} F_{n'}.$$

It follows from (12) that for all $(a, b) \in \mathcal{X} \times \mathcal{Y}$,

$$\sum_{c \in \mathcal{Z}} P_{XZ}(a, c) \cdot P_{\overline{Y}|Z}(b, c) - P_{XY}(a, b) = 0$$

holds, hence we have $P_{X\overline{Y}} = P_{XY}$ and $\text{sim}_X(Z \rightarrow Y)$. This contradiction concludes the proof. \square

Proof of Theorem 11. First of all, it is a consequence of Theorem 10 that $S^*(X; Y|Z) = 0$ holds if X or Y is simulatable.

The idea of transforming a protocol secure against passive adversaries, achieving a key-generation rate arbitrarily close to $S(X; Y|Z)$, into a protocol secure against *active* adversaries, is as follows. The new protocol consists of three parts.

First, Alice and Bob generate a short highly secret key by applying the passive-adversary protocol and authenticating the messages sent by certain blocks (determined by the bits of the message) of independent realizations of their random variables X and Y , respectively.

The second part consists of the generation of a longer key by the same protocol. This time however, the messages sent are authenticated by ε -almost strongly universal hashing, using the key generated in the first protocol phase as the (almost) secret key.

Finally, a message, containing a block of realizations of the random variable known to that party, is sent by the receiver of the final message of the regular protocol. The purpose of this message is to prevent the sender of the final protocol message from erroneously accepting, i.e., although Eve has modified the final message. The receiver of this “control message” does not accept before correctly receiving the message.

We analyze the protocol steps in more detail. Let $\varepsilon, \delta > 0$, and let $R = S(X; Y|Z)$. We show that for all sufficiently large N , there exists a robust $(P_{XYZ}^N, \lfloor (R - \varepsilon)N \rfloor, \varepsilon, \delta)$ -protocol.

Let $\varepsilon' > 0$ be a parameter to be determined later. By the definition of $S(X; Y|Z)$, there exists n_0 and for all $n \geq n_0$ a $(P_{XYZ}^n, \lfloor (R - \varepsilon')n \rfloor, \varepsilon')$ -protocol (secure against *passive* adversaries). Because of Lemma 5, we can even assume that this protocol has a constant number r of rounds (i.e., messages sent in the protocol) for arbitrary n , and that the length of the messages is linear in n .

Let $\delta' := \delta/3$. Alice and Bob now carry out the $(P_{XYZ}^n, \lfloor (R - \varepsilon')n \rfloor, \varepsilon')$ -protocol, where each message bit sent is authenticated by the following method based on Lemma 6. Depending on whether the bit is 0 or 1, the authenticator consists, for some k , of the block $[X_{k+1}, X_{k+2}, \dots, X_{k+l}]$ or $[X_{k+l+1}, X_{k+2}, \dots, X_{k+2l}]$ of realizations of X (or Y , respectively), of a certain length l which is chosen such that the probability of a successful active attack on one of the messages and the probability that an authenticated message is erroneously rejected are upper bounded by δ' .

Let K be the length of the generated key. Then the number of realizations of the distribution P_{XYZ} required to generate this key is of order

$$O(K/R + K \cdot r \cdot \log(rK/\delta')). \quad (13)$$

In the second phase of the active-adversary protocol, the (passive-adversary) $(P_{XYZ}^N, \lfloor (R - \varepsilon/2)N \rfloor, \varepsilon/2)$ -protocol is carried out, where this time, the messages sent are authenticated by ε -ASU hashing. Let q be a prime power such that

$$\frac{q}{\lfloor \log q \rfloor} \geq \frac{2r}{\delta'} \quad (14)$$

holds. Note that the adversary has at most ε' bits of Shannon information about the K -bit key generated during the first protocol phase. Clearly, ε' can be chosen such that the success probability of an active attack is increased by a factor of at most 2 due to Eve’s small partial knowledge about the key. Since the length of the $O(1)$ messages to be authenticated is of order $O(N)$, we have

$$2^i \log q = O(N)$$

for the parameters i and q , and the choice $i := \lfloor \log q \rfloor$ leads to

$$O(1) \cdot (\log q)^2 = O((\log N)^2).$$

According to inequality (14), the success probability of an active attack for this second protocol phase is upper bounded by δ' .

Finally, after the last message is sent in the above protocol, the receiver of this message sends a block of length $O(\log(1/\delta'))$ of realizations of his random variable such that Eve’s chance of successfully faking such a message is at most δ' . Both parties do not accept the outcome before sending and correctly receiving, respectively, this final message.

The required number \overline{N} of realizations for the generation of a key of length $\lfloor (R - \varepsilon/2)N \rfloor$ is of order

$$\begin{aligned} \overline{N} &= O((\log N)^2 \log(\log N/\delta')) \\ &\quad + N + O(\log(1/\delta')) \end{aligned} \quad (15)$$

according to (13). The three summands in (15) correspond to the required numbers of realizations of P_{XYZ} in the three protocol phases as described above. Because of (15), the achievable key-generation rate can, for sufficiently large N , be made greater than $R - \varepsilon$. By construction, the

described protocol is a robust $(P_{XYZ}^N, [(R - \varepsilon)N], \varepsilon, \delta)$ -protocol. Hence

$$S^*(X; Y || Z) = S(X; Y || Z)$$

holds. \square

Example 1: Let us discuss the well-studied special case where the random variables X , Y , and Z are noisy versions of a symmetric binary random variable R (e.g., random bits, for instance broadcast by a satellite, that are received by Alice, Bob, and Eve over binary-symmetric channels with error probabilities $(\leq 1/2)$ α , β , and ε , respectively; see Figure 3).




Fig. 4. The ‘‘Satellite Scenario’’

For the case of only passive adversaries, it was shown in [14], [15] that $S(X; Y || Z) > 0$ holds whenever $\alpha, \beta < 1/2$ and $\varepsilon > 0$ hold. On the other hand, it is easy to see that neither X nor Y are simulatable by Z only if ε is greater than both α and β :

$$S^*(X; Y || Z) > 0 \iff \varepsilon > \max\{\alpha, \beta\} .$$

Thus, secret-key agreement against *active* adversaries is only possible if Alice’s and Bob’s channels are both less noisy than Eve’s channel.

Example 1 suggests that ‘‘non-simulatability’’ implies

$$I(X; Y) > \max\{I(X; Z), I(Y; Z)\} ,$$

i.e., that key agreement is possible with the (non-interactive) error-correction and privacy-amplification protocol phases only [12], [2], [16]. Example 2, however, shows that this is not true in general.

Example 2: Consider the following distribution P_{XYZ} . Let for $\alpha < 1/2$

$$P_{XY}(0, 0) = P_{XY}(1, 1) = \frac{1 - \alpha}{2} ,$$

$$P_{XY}(0, 1) = P_{XY}(1, 0) = \frac{\alpha}{2} .$$

The random variable Z on the other hand is generated by sending the pair $[X, Y]$ over an erasure channel with erasure probability $1 - r$. It is easy to see that if

$$1 - h(\alpha) \leq r < 1 - 2\alpha$$

holds, then Eve cannot simulate X with respect to Y nor Y with respect to X although her amount of Shannon information about X as well as about Y exceeds $I(X; Y)$.

We have studied the problem of unconditionally secure key agreement by communication over a completely insecure channel between parties having access to the outcomes of a certain random experiment. An important special case is when this experiment is repeated independently many times. The main result states that in this scenario, key agreement is either possible at the same rate as against only passive adversaries or not possible at all. In other words, giving the adversary complete access to the communication channel—instead of only read access—does either not reduce at all the possibility of secret-key agreement, or completely destroys it. Which of the two possibilities holds in a specific setting depends on a property, called *simulatability*, of the probability distribution P_{XYZ} modeling the parties’ initial information.

REFERENCES

- [1] R. Ahlswede and I. Csiszár, Common randomness in information theory and cryptography – Part I: secret sharing, *IEEE Transactions on Information Theory*, Vol. 39, No. 4, pp. 1121–1132, 1993.
- [2] C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer, Generalized privacy amplification, *IEEE Transactions on Information Theory*, Vol. 41, Nr. 6, pp. 1915–1923, 1995.
- [3] C. H. Bennett, G. Brassard, and J.-M. Robert, Privacy amplification by public discussion, *SIAM Journal on Computing*, Vol. 17, pp. 210–229, 1988.
- [4] R. E. Blahut, *Principles and practice of information theory*, Addison-Wesley Publishing Company, 1988.
- [5] J. L. Carter and M. N. Wegman, Universal classes of hash functions, *Journal of Computer and System Sciences*, Vol. 18, pp. 143–154, 1979.
- [6] T. M. Cover and J. A. Thomas, *Elements of information theory*, Wiley Series in Telecommunications, 1992.
- [7] I. Csiszár and J. Körner, Broadcast channels with confidential messages, *IEEE Transactions on Information Theory*, Vol. 24, No. 3, pp. 339–348, 1978.
- [8] W. Diffie and M. E. Hellman, New directions in cryptography, *IEEE Transactions on Information Theory*, Vol. IT-22, pp. 644–654, 1976.
- [9] P. Gemmel and M. Naor, Codes for interactive authentication, *Advances in Cryptology - Proceedings of Crypto ’93*, Lecture Notes in Computer Science, Vol. 773, Springer-Verlag, Berlin, pp. 355–367, 1994.
- [10] N. Gisin, R. Renner, and S. Wolf, Linking classical and quantum key agreement: is there a classical analog to bound entanglement?, *Algorithmica*, Vol. 34, pp. 389–412, 2002.
- [11] N. Gisin and S. Wolf, Linking classical and quantum key agreement: is there ‘‘bound information’’?, *Proceedings of CRYPTO 2000*, Lecture Notes in Computer Science, Vol. 1880, pp. 482–500, Springer-Verlag, 2000.
- [12] 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.
- [13] U. M. Maurer, Information-theoretically secure secret-key agreement by NOT authenticated public discussion, *Advances in Cryptology - EUROCRYPT ’97*, Lecture Notes in Computer Science, Vol. 1233, pp. 209–225, Springer-Verlag, 1997.
- [14] U. M. Maurer and S. Wolf, Towards characterizing when information-theoretic secret key agreement is possible, *Advances in Cryptology - ASIACRYPT ’96*, Lecture Notes in Computer Science, Vol. 1163, pp. 196–209, Springer-Verlag, 1996.
- [15] U. M. Maurer and S. Wolf, Unconditionally secure key agreement and the intrinsic conditional information, *IEEE Transactions on Information Theory*, Vol. 45, No. 2, pp. 499–514, 1999.
- [16] U. M. Maurer and S. Wolf, Information-theoretic key agreement: from weak to strong secrecy for free, *Advances in Cryptology - EUROCRYPT 2000*, Lecture Notes in Computer Science, Vol. 1807, pp. 351–368, Springer-Verlag, 2000.

- [17] U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part II: The simulatability condition, *IEEE Transactions on Information Theory*, 2003.
- [18] U. M. Maurer and S. Wolf, Secret-key agreement over unauthenticated public channels – Part III: Privacy amplification, *IEEE Transactions on Information Theory*, 2003.
- [19] C. E. Shannon, Communication theory of secrecy systems, *Bell System Technical Journal*, Vol. 28, pp. 656–715, 1949.
- [20] 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.
- [21] S. Wolf, *Information-theoretically and computationally secure key agreement in cryptography*, ETH dissertation No. 13138, Swiss Federal Institute of Technology (ETH Zurich), 1999.
- [22] A. D. Wyner, The wire-tap channel, *Bell System Technical Journal*, Vol. 54, No. 8, pp. 1355–1387, 1975.