Information-Theoretically Secure Secret-Key Agreement by NOT Authenticated Public Discussion[1]

Ueli Maurer

Department of Computer Science
ETH Zurich
CH-8092 Zurich, Switzerland
maurer@inf.ethz.ch

Abstract. All information-theoretically secure key agreement protocols (e.g. based on quantum cryptography or on noisy channels) described in the literature are secure only against passive adversaries in the sense that they assume the existence of an authenticated public channel. The goal of this paper is to investigate information-theoretic security even against active adversaries with complete control over the communication channel connecting the two parties who want to agree on a secret key. Several impossibility results are proved and some scenarios are characterized in which secret-key agreement secure against active adversaries is possible. In particular, when each of the parties, including the adversary, can observe a sequence of random variables that are correlated between the parties, the rate at which key agreement against active adversaries is possible is characterized completely: it is either 0 or equal to the rate achievable against passive adversaries, and the condition for distinguishing between the two cases is given.

1 Introduction

One of the fundamental problems in cryptography is the generation of a shared secret key by two parties, Alice and Bob, not sharing a secret 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 [9] (see also [20]) 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 certain finite group) in feasible time, and

[1] To appear in *Advances in Cryptology - EUROCRYPT '97*, W. Fumy (Ed.), Lecture Notes in Computer Science, Berlin: Springer-Verlag, Konstanz, Germany, May 1997. This work is supported in part by the Swiss National Science Foundation, grant no. 20-42105.94.

(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 consider 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. Secret-key agreement can be possible in this scenario only if Alice and Bob (but possibly also Eve) have correlated information. More formally, while Alice and Bob share no secret key initially, they know some random variables X and Y, respectively, jointly distributed with a random variable Z known to Eve. The joint probability distribution is denoted 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 research topic.

There exists a substantial body of results on secret-key agreement by public discussion secure against adversaries with infinite computing power (see Section 2.3 for a brief summary), but they all depend in a crucial manner on the assumption that eavesdroppers are *passive* and hence the communication between Alice and Bob can be assumed to be authenticated. Of course, as is pointed out in these papers, the authenticity can be guaranteed, even when the channel is completely insecure, when Alice and Bob initially share a secret key that is used for authentication purposes (see Section 2.2). Hence these results can be interpreted as providing information-theoretically secure protocols for expanding a short initially shared secret key to an arbitrarily long secret key.

This paper characterizes scenarios in which secret-key agreement against active adversaries is possible and shows that for an important class of scenarios of correlated random variables available to Alice, Bob and Eve, active adversaries are not more powerful than passive ones.

2 Key-agreement protocols

2.1 Scenarios and definitions

We now formalize key-agreement protocols; the security of such protocols will be defined later.

Definition 1. A *key-agreement protocol* consists of three phases:

- a (possibly missing) initialization phase[1] in which Alice, Bob and an adversary Eve receive random variables X, Y and Z, respectively, which are jointly distributed according to some probability distribution P_{XYZ}.
- During the communication phase Alice and Bob alternate sending each other messages $C_1, C_2, \ldots$ where we assume that Alice sends messages $C_1, C_3, C_5, \ldots$ and Bob sends messages $C_2, C_4, C_6, \ldots$ Each message depends possibly on the sender's entire view of the protocol at the time it is sent and possibly on privately generated random bits. Let t be the total number of messages and let $C^t = [C_1, \ldots, C_t]$ denote the set of exchanged messages.
- Finally, Alice and Bob each either accepts or rejects the protocol execution, depending on whether they believe to be able to generate a secret key. If Alice accepts, she generates a key S depending on her view of the protocol. Similarly, if Bob accepts, he generates a key S' depending on his view of the protocol.

In general, the channel connecting Alice and Bob is completely insecure, i.e. Eve can see every message C_i and replace it by an arbitrary message $\tilde{C}_i$ of her choice. She need not keep Alice and Bob synchronized and she can impersonate either party by fraudulently initiating a protocol execution.

For stating impossibility results in the strongest possible form, we also consider protocols in which certain messages can be sent in a secret or authenticated manner (by appropriate means not specified by the protocol).

Definition 2. If a message C_i is *secret* (by the protocol specification), Eve learns nothing about it except that it exists[2]. However, she may replace such a message by a different message. If a message C_i is *authenticated* (by the protocol specification), then the receiver will always (with probability 1) detect any modification to the message due to Eve, but Eve sees the message.

Considering a passive adversary is equivalent to assuming the entire communication to be authenticated. The above definition can be made information-theoretically precise.

If two parties share a secret key, they can use the one-time pad encryption to transmit a message in perfect secrecy over a completely insecure channel. They can also use part of the secret key for authenticating messages (see Section 2.2).

[1] The initialization phase summarizes the parties' entire initial information, for instance the history of previous executions of protocols, the information resulting from quantum transmissions (like in quantum cryptography [2]), or information received from other sources like a satellite broadcasting random bits (see Section 4.3) or the signal of a deep-space radio source. When the initialization phase is missing, this means that Alice's and Bob's complete knowledge at the beginning of the protocol is assumed to be statistically independent.

[2] It is possible that she later obtains information about C_i because subsequent messages depend on C_i, but Eve never learns anything about C_i not provided by subsequent messages. This will be formalized in the full paper.

However, in contrast to perfect secrecy, perfect authenticity cannot be achieved even if a secret key of arbitrary fixed size is used because an adversary can always guess the key with non-zero probability of success. Authenticity and confidentiality are dual security properties, and the duality can be shown in various ways (e.g., see [16]).

All the protocol steps proposed in this paper are polynomial-time computable, but there may generally be steps in subprotocols taken from the literature that are not known to be computable in polynomial time. However, for every protocol resulting in Alice and Bob sharing a secret key mentioned here, there also exist efficient protocols for generating a secret key (which may be somewhat shorter).

In general, the distribution P_{XYZ} may be under Eve's partial control and may only partly be known to Alice and Bob. Two examples are the privacy amplification scenario [3] mentioned in Section 2.3, and quantum cryptography, where both Bob's and Eve's distributions depend on the type of measurement performed by Eve on the photons sent by Alice. In this paper we assume that P_{XYZ} is known to all parties.

In the sequel we assume without loss of generality that S and S' are binary strings of length $|S| = |S'| = k$. Clearly, the goal of a protocol is that S and S' agree with very high probability and that Eve has very little information about S. An adversary can of course block the communication between Alice and Bob completely by replacing all messages by empty messages, thus preventing any secret-key agreement. The goal of the design of a protocol can thus only be to generate a (hopefully large amount of) secret key when Eve is passive, but to detect any tampering with very high probability. However, even when Eve's strategy is active, it is allowed that she goes undetected if the secret key shared by Alice and Bob at the end of the protocol nevertheless is secret. In other words, Alice and Bob should not primarily be interested in catching an active cheater but in making sure that whenever they believe (or at least one of them believes) to have agreed on a secret key, then this is indeed the case with very high probability.

Definition 3. A key-agreement protocol with $|S| = k$ is (ϵ, δ)-*secure* if, for every passive eavesdropping strategy,

$$\begin{aligned} P[S \neq S'] &\leq \epsilon, \\ I(S; C^t Z) &\leq \epsilon, \\ \text{and} \qquad H(S) &\geq k - \epsilon, \end{aligned}$$

and if for every active adverse strategy, with probability at least $1 - \delta$, either Eve is caught by at least Alice or Bob (i.e. they do not both accept) or they successfully generate a secret key S (and S') satisfying the above conditions.

Note that one cannot require both Alice and Bob to reject. Eve could delete the last message from Alice to Bob (or vice versa) that would make Bob accept after Alice has accepted. (Byzantine agreement is impossible between two players in the presence of an active adversary.)

Here $H(S)$ denotes the entropy[3] of S and $I(S;C^tZ) = H(S) - H(S|C^tZ)$ denotes the information about S given by Eve's total observation (consisting of C^t and Z). The condition $H(S) \geq k - \epsilon$ implies that S is virtually uniformly distributed and together with the condition $I(S;C^tZ) \leq \epsilon$ it implies $H(S|C^tZ) \geq k - 2\epsilon$ and hence that S is also virtually uniformly distributed from Eve's point of view, i.e., given Eve's total information. Such a uniformity constraint could alternatively be defined in terms of any reasonable constraint on the deviation of a distribution from the uniform distribution, without changing the results of this paper.

2.2 Unconditionally secure message authentication

Adversaries with complete control over the communication channel have previously been considered in message authentication scenarios where, unlike in this paper, a secret key is shared initially by Alice and Bob about which Eve is assumed to have no information *a priori*.

Unconditionally secure message authentication based on a shared secret key was first considered in [11] and later in a large number of papers (e.g. [22], [23]). One of the most recent papers on this topic is by Gemmell and Naor [10] who proved the surprising result that interactive protocols for authenticating an n-bit message are more efficient in terms of the length of the secret key required to restrict an adversary's cheating probability to at most p. In particular, they proposed a one-round protocol using only $\log n - 2\log p$ bits of secret key and showed that this can be reduced to $\log^{(k)} n - 5\log p$ in a k-round protocol. We will make use of these results.

2.3 Review of the literature

In this section some of the results on secret-key agreement by perfectly authenticated public discussion are reviewed. Shannon's [21] famous result on perfect secrecy, stating that a cipher can achieve perfect secrecy only if the entropy of the secret key is at least as large as the entropy of the plaintext, can be considered as a special case (for 1-round protocols) of Theorem 1 below. Although Wyner's wire-tap channel scenario [25] and Csiszár and Körner's generalization [8] thereof do not include a public channel between Alice and Bob, they should nevertheless be mentioned here. In those scenarios, Alice can send information over a so-called broadcast channel where Bob and Eve can receive different outputs of the channel. Secret information transmission (and hence secret-key agreement) was shown to be possible if and only if Eve's channel is noisier than Bob's channel [8], an assumption that is generally unrealistic.

In the scenario considered in quantum cryptography (see [2] and references therein), Alice can send polarized light pulses of very low intensity to Bob over

[3] $H(S) = -\sum_{s:\, P_S(s)>0} P_S(s) \log_2 P_S(s)$. See [6] for an introduction to the basic concepts of information theory.

some channel (e.g. an optical fiber) controlled by Eve. The use of this quantum communication results in Alice, Bob, and Eve possessing correlated strings. By subsequent discussion over the authenticated public channel, Alice and Bob manage to generate a secret key about which Eve has arbitrarily little information.

Another special case of key agreement protocols secure against passive adversaries is privacy amplification introduced in [4] and generalized in [3]. Privacy amplification is a protocol step that would typically be used as the last step in a practical key agreement protocol, but it can itself be described in the framework of key agreement protocols. Here Alice and Bob are assumed to know a string W (i.e. $X = Y = W$) about which Eve has some partial information. The protocol of [3] is secure even when Eve specifies an arbitrary probability distribution P_{ZW} unknown to Alice and Bob, subject to the only constraint that a bound on the second order Rényi entropy of W, given Eve particular value z of Z, is known to Alice and Bob. In the privacy amplification literature only passive adversaries have been considered. It is proved in [19] that privacy amplification secure against active adversaries is possible when the adversary's min-entropy about the string is more than half its length.

3 The case of no common initial information

In this section we characterize to what extent secret and/or authenticated communication between Alice and Bob can help them to agree on a secret key. These results 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 techniques (see Section 2.2) for authenticating a message to be sent to Alice.

In sharp contrast, only the computational model allows to transform 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. (Actually, this follows from Theorem 1 below.) See also [16] for a discussion of the described and other security transformations. It is hence 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 1. *Consider key agreement protocols without initialization phase which allow some of the exchanged messages to be either secret or authenticated. For* $\epsilon \leq 1 - 3/(|S| + 2)$ *there exists no such protocol that is* (ϵ, δ)*-secure, even when all messages are authenticated (or, equivalently, when Eve is passive.) Moreover, even if all messages from Alice to Bob are secret and all messages from Bob to*

Alice are authenticated, there exists no such protocol that is (ϵ, δ)-secure against active adversaries for any $\delta < 1$.

Proof. To prove the first part we make use of Theorem 1 of [14] which implies that

$$H(S) \leq H(S|S') + I(S; C^t) \tag{1}$$

for all such protocols. Note that the random variables X, Y do not exist in our context and hence $I(X;Y) = 0$ in Theorem 1 of [14] . Fano's Lemma (see [6]) states that the error probability p of guessing a random variable U when given a correlated random variable U' satisfies

$$H(U|U') \leq h(p) + p \log_2(|\mathcal{U}| - 1),$$

where $\mathcal{U}$ is the set of possible values that U can take on[4]. Therefore the condition $P[S \neq S'] \leq \epsilon$ implies

$$H(S|S') < h(\epsilon) + \epsilon k$$

which together with inequality (1) and the second and third conditions of Definition 3 gives

$$k - \epsilon \leq H(S) < h(\epsilon) + \epsilon k + \epsilon.$$

Using $h(\epsilon) \leq 1$, this implies $k - 1 \leq \epsilon(k+2)$ and hence $\epsilon > 1 - 3/(k+2)$.

To prove the second part, notice that from Bob's point of view, Alice has no advantage compared to Eve. When Eve performs the same protocol as Alice would, pretenting to be Alice, Bob accepts with the same probability as he would accept a protocol execution with Alice which according to the definition is 1. □

Note again that the first statement of the theorem is in sharp contrast to the public-key cryptographic scenario where, under a suitable intractability assumption, secret-key agreement secure against computationally bounded adversaries is possible when a single authenticated message in each direction can be sent. A public-key cryptosystem can be interpreted [16] as a means for transforming an authenticated channel into a secret channel in the other direction. The following well-known result is an observation following from Theorem 1.

Corollary 2. *A public-key cryptosystem can be computationally secure but not information-theoretically (i.e. unconditionally) secure.*

Theorem 3. *Assume that one secret (but not necessarily authenticated) message can be sent from Alice to Bob. Then, for any $\delta > 0$, key agreement $(0, \delta)$-secure against active adversaries is possible 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.*

[4] $h(p) = -p \log_2 p - (1-p) \log_2(1-p)$ denotes the binary entropy function which measures the entropy of a binary random variable that takes on the two values with probabilities p and $1-p$.

Proof. Note that when the same message from Alice to Bob is both 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 Bob ($n \geq -2\log_2 \delta$) over the secret channel and the description of a function f in a universal class hash functions from $\{0,1\}^n$ to $\{0,1\}^n$ [7] 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 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 ($\Omega(\log \delta)$). Then they use the above protocol where the authenticated channel is obtained by Alice by using an authentication scheme [10] using R as the secret key. □

Theorem 1 is pessimistic: it demonstrates 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 possibly only two bitstrings that are somehow correlated), about which also Eve has partial knowledge, then secret-key agreement can be possible.

In the following we consider such scenarios. One of our general goals is to achieve secret-key agreement under mild conditions on such an initialization phase, for instance conditions that can be argued to occur (or can be made to occur) in a realistic communications scenario.

4 Protocols with initialization phase

4.1 Impossibility results

The following theorem on authenticated public discussion follows from Corollary 1 in [14]. Recall from Section 2 that X, Y, and Z are the random variables obtained by Alice, Bob, and Eve, respectively, during the initialization phase.

Theorem 4. *For every probability distribution P_{XYZ}, a key agreement protocol that is (ϵ, δ)-secure against passive (or active) adversaries satisfies*

$$H(S) \;\leq\; \min[I(X;Y), I(X;Y|Z)] + h(\epsilon) + \epsilon(k+1).$$

In particular, for $\epsilon = 0$, we have $H(S) \;\leq\; \min[I(X;Y), I(X;Y|Z)]$.

Note that by definition, $I(X;Y) = H(X) - H(X|Y)$ and $I(X;Y|Z) = H(X|Z) - H(X|YZ)$ and that $I(X;Y|Z) \geq I(X;Y)$ is possible. It will be demonstrated in the following section that this theorem is not as pessimistic as it looks at first sight.

Theorem 4 states that secret-key agreement is possible and only if Y gives a substantial amount of information about X, both when Z is given or when it is not. In other words, X and Y must be correlated, and this correlation must to some extent be independent of Z. The bound $\min[I(X;Y), I(X;Y|Z)]$ can be replaced by the stronger bound derived in [18], called the intrinsic mutual information between X and Y given Z. It is the minimum of $I(X;Y|Z')$ over conditional probability distributions $P_{Z'|Z}$.

Definition 4. We call the distribution P_{XYZ} *X-simulatable by Eve* if Eve can generate from Z a random variable $\tilde{X}$ such that the pairs $[X,Y]$ and $[\tilde{X},Y]$ have the same distribution, i.e. if there exists a conditional probability distribution $P_{\tilde{X}|Z}$ such that

$$P_{\tilde{X}Y}(x,y) = P_{XY}(x,y)$$

for all x and y, where $P_{\tilde{X}Y}$ is the marginal distribution of $P_{X\tilde{X}YZ} = P_{XYZ} \cdot P_{\tilde{X}|Z}$, i.e.,

$$P_{\tilde{X}Y}(x,y) = \sum_{x'} \sum_{z} P_{XYZ}(x',y,z) \cdot P_{\tilde{X}|Z}(x,z).$$

Similarly, the distribution P_{XYZ} is called *Y-simulatable by Eve* if the symmetric condition with respect to Bob, with X replaced by Y and $\tilde{X}$ replaced by $\tilde{Y}$, is satisfied.

More intuitively, P_{XYZ} is X-simulatable by Eve if she can send Z through a (simulated) channel (characterized by $P_{\tilde{X}|Z}$) whose output $\tilde{X}$ has the same *joint* distribution with Y as X. (An example of such a distribution is given in Section 4.3.) Therefore, when P_{XYZ} is X-simulatable by Eve, then there is no way Bob can distinguish between a correct message sent by Alice and an appropriately generated fraudulent message sent by Eve. Similarly, when P_{XYZ} is Y-simulatable by Eve, then there is no way Alice can distinguish between a correct message sent by Bob or a fraudulent message sent by Eve. We obtain the following generalization of Theorem 1.

Theorem 5. *When P_{XYZ} is X-simulatable (or Y-simulatable) by Eve, then no key agreement protocol can be (ϵ,δ)-secure against active adversaries for any ϵ and $\delta < 1$, even if all messages from Alice to Bob (Bob to Alice) are perfectly secret and all messages from Bob to Alice (Alice to Bob) are authenticated.*

4.2 Independent repetition of a random experiment

In order to be able to derive interesting results on secret-key agreement against active or passive adversaries, we must consider specific types of probability distributions of the random variables given to Alice, Bob, and Eve.

One natural assumption is that the random experiment generating the triple $[X,Y,Z]$ is repeated many times independently. Hence we assume that Alice, Bob and Eve receive strings $X^n = [X_1,\ldots,X_n]$, $Y^n = [Y_1,\ldots,Y_n]$, and $Z^n =$

$[Z_1, \ldots, Z_n]$, respectively, where

$$P_{X^nY^nZ^n}(x_1, \ldots, x_n, y_1, \ldots, y_n, z_1, \ldots, z_n) \;=\; \prod_{i=1}^{n} P_{XYZ}(x_i, y_i, z_i).$$

Note that we have changed the notation here and for the rest of the paper: P_{XYZ} now denotes the distribution of one of several random experiments while it previously denoted the distribution of the overall experiment.

This particular scenario is motivated by the well-known models for discrete memoryless sources and channels of communication theory. Many concrete practical scenarios can be modeled in this way, for instance the one discussed below in which Alice, Bob, and Eve receive noisy versions of a random string broadcast by a satellite or of the signal emitted by a deep space radio source.

For such a scenario of independent repetitions of a random experiment, the quantity that is of most interest is the maximal rate at which Alice and Bob can generate secret key bits, where rate is to be understood per execution of the random experiment generating a triple $[X, Y, Z]$.

Definition 5. The *secret key rate of P_{XYZ} for passive adversaries*, denoted $S(P_{XYZ})$, is the maximum rate at which Alice and Bob can agree on a secret key S while keeping a passive adversary's information about S arbitrarily small. More formally, it is the maximal R such that for all $\epsilon > 0$, for all $R' < R$, and for all sufficiently large n there exists a protocol with $|S| = \lfloor R'n \rfloor$ that is $(\epsilon, 0)$-secure against passive adversaries[5]. The *secret key rate of P_{XYZ} for active adversaries*, denoted $S^*(P_{XYZ})$, is defined in the same way, except that the adversary is allowed to be active, and for any given $\delta > 0$, (ϵ, δ)-security is required instead of $(\epsilon, 0)$-security.

The first part of this definition is given in [15] as a considerably strengthened definition of that given in [14] , and the second part is new. In particular, in [14] it was only required that the *rate* at which Eve obtains information, $I(S; C^t Z^n)/n$ be arbitrarily small for large n, and proving results for the much stronger definition involves some technical steps, including privacy amplification [3]. The following result was proved in [15] (and in [14] using the weaker definition).

Theorem 6. *$S(P_{XYZ})$ is lower and upper bounded by*

$$\max[0,\, I(Y;X) - I(Z;X),\, I(X;Y) - I(Z;Y)] \;\leq\; S(P_{XYZ})$$

and

$$S(P_{XYZ}) \;\leq\; \min[I(X;Y),\, I(X;Y|Z)].$$

The lower bound is not tight in general. In particular, for the binary scenario discussed in Section 4.3, if Eve's channels is less noisy than both Alice's and Bob's channel, the lower bound vanishes while the secret-key rate is actually strictly positive.

[5] For the case of passive adversaries, $\delta = 0$ can trivially be achieved.

We are primarily interested in investigating the relation between $S(P_{XYZ})$ and $S^*(P_{XYZ})$, i.e., the power of authenticated versus non-authenticated communication. Quite surprisingly, it turns out that $S^*(P_{XYZ}) = 0$ or $S^*(P_{XYZ}) = S(P_{XYZ})$. However, before treating the general case, we consider the case of binary symmetric random variables which is of particular interest.

4.3 The binary case

In this section we consider the natural special case where the random variables known to Alice, Bob and Eve are noisy versions of a random string (e.g. broadcast by a satellite) received over binary symmetric channels C_A, C_B and C_E with bit error probabilities ϵ_A, ϵ_B and ϵ_E, respectively (see Figure 1). Without loss of generality we assume that these channels are independent because any scenario of dependent channels can be transformed [14] into an equivalent scenario of independent channels (with different bit error probabilities). In other words, when U denotes the random bit generated by the source ($P_U(0) = P_U(1) = 1/2$), we have

$$P_{XYZ|U} = P_{X|U} \cdot P_{Y|U} \cdot P_{Z|U}$$

where $P_{X|U}(x,r) = 1 - \epsilon_A$ if $x = u$ and ϵ_A else, $P_{Y|U}(y,r) = 1 - \epsilon_B$ if $y = u$ and ϵ_B else and $P_{Z|U}(z,r) = 1 - \epsilon_E$ if $z = u$ and ϵ_E else.

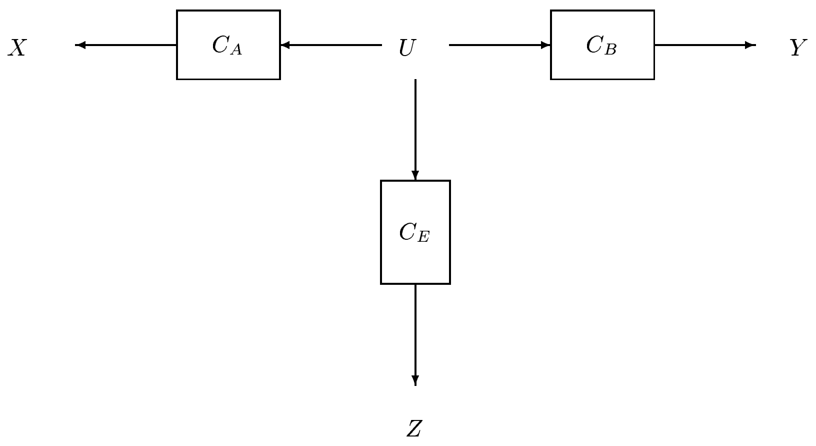

Fig. 1. The scenario of three independent channels

It is easy to verify that P_{XYZ} is X-simulatable by Eve if and only if $\epsilon_E \leq \epsilon_A$ and it is Y-simulatable by Eve if and only if $\epsilon_E \leq \epsilon_B$. Such a simulation can be achieved by Eve by sending Z through an additional (simulated) binary symmetric channel of appropriate bit error probability. Therefore, when either $\epsilon_E \leq \epsilon_B$ or $\epsilon_E \leq \epsilon_A$ in the described scenario, then $S^*(P_{XYZ}) = 0$ by Theorem 5.

Let

$$\epsilon_{AB} = \epsilon_A + \epsilon_B - \epsilon_A\epsilon_B$$

be the bit error probability between corresponding bits of Alice's and Bob's strings, and let similarly

$$\epsilon_{AE} = \epsilon_A + \epsilon_E - \epsilon_A\epsilon_E$$

and

$$\epsilon_{BE} = \epsilon_B + \epsilon_E - \epsilon_B\epsilon_E$$

be the bit error probabilities between corresponding bits of Alice's and Eve's and between Bob's and Eve's strings, respectively.

Assuming that Alice and Bob share no secret key initially, authentication for messages transmitted from Alice to Bob can nevertheless be achieved when Eve's channel is noisier than Alice's channel ($\epsilon_E > \epsilon_A$). This implies that $\epsilon_{BE} > \epsilon_{AB}$, i.e. that Alice's bits agree with Bob's bits with higher probability than Eve's bits agree with Bob's bits.

To demonstrate this fact, consider the following (very wasteful) authentication method.[6] A more efficient scheme will be considered below. In order to authenticate a single bit ($k = 1$) sent from Alice to Bob, Alice appends a substring of X^n of length l. The two substrings of X^n appended to authenticate a 0 or a 1 are disjoint. For instance, a 0 or a 1 is authenticated by appending (for some q) the string $[X_q, \ldots, X_{q+l-1}]$ or $[X_{q+m}, \ldots, X_{q+2l-1}]$, respectively, as the authenticator, and these $m = 2l$ bits of X^n are never used again for any other purpose. Bob expects to receive as an authenticator either a version of $[Y_q, \ldots, Y_{q+l-1}]$ or of $[Y_{q+m}, \ldots, Y_{q+2l-1}]$ with a fraction of close to ϵ_{AB} bit errors. Informally, Bob hence accepts the received bit if and only if the fraction of bits in the authenticator that agree with his noisy version of the authenticator ($[Y_q, \ldots, Y_{q+l-1}]$ or $[Y_{q+m}, \ldots, Y_{q+2l-1}]$) is not much smaller than $1 - \epsilon_{AB}$. It is easy to see that for any fixed $\epsilon_{BE} > \epsilon_{AB}$, the probability that Eve can successfully deceive Bob is exponentially small in l.

The described scheme is quite inefficient in terms of the number of bits used from the sequence. A much better approach is described in the proof of the following theorem.

Theorem 7. *When $\epsilon_{BE} > \epsilon_{AB}$ in the described binary scenario, a k-bit message sent from Alice to Bob can be authenticated by an l-bit authenticator with $l = 2k$ using $m = 4k$ bits of the random string X^n and achieving an arbitrarily small deception probability for sufficiently large k.*

Proof sketch. A scheme for authenticating a k-bit message sent from Alice to Bob using m bits of X^n (e.g. $[X_q, \ldots, X_{q+m-1}]$ for some q) can be derived as follows. Every message is authenticated by appending a particular subset of bits in $[X_q, \ldots, X_{q+m-1}]$. These subsets should be sufficiently disjoint to avoid that

[6] In the following we consider schemes for authenticating a k-bit message by an l-bit authenticator using $m > l$ bits of the common sequence.

such an authenticator can be guessed by Eve from an observed one. Bob checks whether his version of the authenticator (i.e. his subset of $[Y_q, \ldots, Y_{q+m-1}]$) agrees with the received authenticator on a fraction roughly $1 - \epsilon_{AB}$ of the bits, as expected when Alice sends the authenticator. Security requires that given one of these sets, it should be impossible for Eve to approximate a different authenticator of Alice with a bit error fraction close to ϵ_{AB}.

When Eve has intercepted a message together with its authenticator, her best strategy for creating an authenticator for a different message (hoping that it will be accepted by Bob) is to copy those bits from the received authenticator that are also contained in the new authenticator and to take as guesses for the remaining bits her copies of the bits (in $[Z_q, \ldots, Z_{q+m-1}]$), introducing bit errors in those bits with probability ϵ_{BE}. The maximal probability of successful deception is hence determined by the number d of bits that Eve must guess and the total number l of bits in the forged authenticator.

The expected value and the standard deviation of the number of bits in the correct autenticator that agree with Bob's corresponding bits are

$$\mu = l(1 - \epsilon_{AB})$$

and

$$\sigma = \sqrt{l\epsilon_{AB}(1 - \epsilon_{AB})},$$

respectively. When Eve tries to deceive Bob, the expected value and the standard deviation of the fraction of bits in the forged autenticator that agree with Bob's corresponding bits are

$$\mu' = (l - d)\epsilon_{AB} + d\epsilon_{BE}$$

and

$$\sigma' = \sqrt{(l - d)\epsilon_{AB}(1 - \epsilon_{AB}) + d\epsilon_{BE}(1 - \epsilon_{BE})},$$

respectively. Bob accepts an authenticator if and only if the number of his bits that agree with the corresponding authenticator bits is within q standard deviations of μ, where q is a security parameter that grows with l. The difference between the two expected values is $d\epsilon_{BE}$ and the standard deviation is $\sigma = \Omega(\sqrt{l})$. When d grows substantially faster than $\sqrt{l}$ one can let $q = \Omega(d/\sqrt{l})$. The law of large numbers implies that Eve's cheating probability decreases exponentially in q.

We now investigate how this can be achieved. An appropriate set of such subsets of bit positions (i.e., subsets of $\{1, \ldots, m\}$) can be interpreted as a code: each subset corresponds to a codeword of length m, where a 1 (or a 0) indicates that the bit at the corresponding position is (is not) contained in the subset. The weight of a codeword is equal to the length of the corresponding authenticator.

The desired distance property of the code differs from the Hamming distance considered in the theory of error-correcting codes. Instead, we define the *0-1 distance* from a codeword c_1 to a codeword c_2, denoted $d(c_1 \to c_2)$, as the number of bits that Eve must guess when trying to convert the authenticator corresponding to c_1 into the authenticator corresponding to c_2. The distance $d(c_1 \to c_2)$ is hence defined as the number of transitions from 0 to 1 when going from c_1 to c_2, hence not counting the transitions from 1 to 0. Note that this distance is not symmetric, i.e. $d(c_1 \to c_2) \neq d(c_2 \to c_1)$ in general. It is required that the $0-1$ distance from any codeword to any other codeword be large, say at least d. A conventional linear code cannot be used because the $0-1$ distance from any codeword to the zero-codeword is zero.

We now give a simple construction of codes that are good with respect to this distance measure. One can convert any code of length l and minimum distance d into a (non-linear) code of length $m = 2l$ and minimum $0-1$ distance d, where each codeword has weight l. This is achieved by replacing every bit in the original code by pair of bits, namely by replacing 0 by 01 and 1 by 10.

In the context of this proof, a possible code to be used for the construction is an extended Reed-Solomon code over a finite field $GF(2^r)$ [5]. For any K there exists such a code encoding K information digits into codewords of length $N = 2^r$ and with minimum distance $N - K + 1$. By interpreting elements of $GF(2^r)$ as binary substrings of length r, we obtain a binary code with 2^{rK} codewords of length $2rN$ and with minimum $0-1$ distance at least d.

By taking r as a security parameter and letting $N = 2^r$, $K = N/2$ and $k = rK$ we obtain $l = 2k = rN$ and $m = 2l = 2rN$. This is sufficient to complete the proof. □

By symmetry, the same technique can be used to authenticate messages sent from Bob to Alice, provided that $\epsilon_E > \epsilon_B$. This theorem shows that the rate at which random bits are needed for authentication is a constant factor times the bit rate at which Alice sends messages to Bob. Therefore, the secret key rate of P_{XYZ} for active adversaries is a constant (≤ 1) times the secret key rate of P_{XYZ} for passive adversaries. In the proof of the following theorem we need to show that the number of bits needed for authentication is asymptotically negligible compared to the number of bits needed for secret-key agreement (in the passive case).

Theorem 8. *When both $\epsilon_E > \epsilon_B$ and $\epsilon_E > \epsilon_A$ in the described scenario, then $S^*(P_{XYZ}) = S(P_{XYZ})$, i.e., an active adversary is not more powerful than a passive adversary. Otherwise, if either $\epsilon_E > \epsilon_B$ or $\epsilon_E > \epsilon_A$, then $S^*(P_{XYZ}) = 0$.*

Proof. The fact that $S^*(P_{XYZ}) = 0$ when either $\epsilon_E < \epsilon_B$ or $\epsilon_E < \epsilon_A$ follows from Theorem 5 because P_{XYZ} is either X-simulatable or Y-simulatable by Eve. The fact that $S^*(P_{XYZ}) = S(P_{XYZ})$ when $\epsilon_E > \epsilon_B$ and $\epsilon_E > \epsilon_A$ can be proved as follows. A suboptimal protocol based on the authentication method of Theorem 7 can be used to generate a relatively small t-bit secret key K, using $O(t)$ bits of the random string. This key can then be used, similar to a bootstrapping process, for instance based on the protocols of [10], to authenticate the messages exchanged in an optimal passive-adversary protocol $\mathcal{P}$ achieving $S(P_{XYZ})$. The size of K must only be logarithmic in the maximal size of a message exchanged in $\mathcal{P}$ [10] and linear in the number of rounds of $\mathcal{P}$. No matter what amount of secret key must be generated by $\mathcal{P}$, this can be achieved by using messages of size proportional to the key size in a constant number of rounds. Therefore, the ratio of size of K and the size of the generated key vanishes asymptotically. □

It is known from [14] that

$$\min[h(\epsilon_{AE}), h(\epsilon_{BE})] - h(\epsilon_{AB}) \;\leq\; S(P_{XYZ}) \;\leq\; 1 - h(\epsilon_{AB}).$$

It was recently proved that $S(P_{XYZ}) > 0$ unless $\epsilon_E = 0$ [17], even when both $\epsilon_E < \epsilon_B$ and $\epsilon_E < \epsilon_A$, i.e., even when the above lower bound vanishes (or is negative).

4.4 A completeness result for the general case

Let P_{XYZ} be an arbitrary probability distribution of a random experiment that is repeated many times. In general, only lower and upper bounds on $S(P_{XYZ})$ are known and $S(P_{XYZ})$ is known exactly only for special cases. The following theorem characterizes $S^*(P_{XYZ})$ completely in terms of P_{XYZ} and $S(P_{XYZ})$ and characterizes the power of active adversaries in comparison to passive ones for the described noisy-channel initialization scenario. Determining the exact power of a passive adversary remains an open problem.

Theorem 9. *When P_{XYZ} is either X-simulatable or Y-simulatable by Eve, then $S^*(P_{XYZ}) = 0$. Otherwise, $S^*(P_{XYZ}) = S(P_{XYZ})$.*

Proof sketch. The proof of this theorem relies on the theory of typical sequences[7] and is similar to the proof of Theorem 8, which is a special case of this theorem, but the technical details are omitted from this extended abstract. In order to authenticate a k-bit message by an $l = 2k$-bit authenticator using $m = 4k$ bits of X^n (or of Y^n when Bob is the sender), the described approach based on error correcting codes can be used to select the positions of a subsequence $[X_{i_1}, \ldots, X_{i_l}]$ of X^n. The receiver accepts the message if and only if the sequence of pairs $[(X_{i_1}, Y_{i_1}), \ldots, (X_{i_l}, Y_{i_l})]$ is γ-typical for the distribution P_{XY} for some suitable small γ. One can prove that for every distribution P_{XYZ} that is neither X-simulatable nor Y-simulatable by Eve, there exists a positive γ such that for sufficiently large k Eve's cheating probability is arbitrarily small. The same argument as in the proof of Theorem 8 can be used to prove that the ratio of bits needed for authentication and of bits used for secret-key agreement vanishes asymptotically. □

Acknowledgement

I would like to thank Christian Cachin and Stefan Wolf for interesting discussions and helpful comments.

[7] Loosely speaking, a sequence $U_1, \ldots, U_r$ of digits of an alphabet $\mathcal{U}$ is γ-typical for a given distribution P_U over $\mathcal{U}$ if for every $u \in \mathcal{U}$ the fraction of occurrences of u in $U_1, \ldots, U_r$ deviates by at most γ from $P_U(u)$ (see for instance [6]).

References

1. R. Ahlswede and I. Csiszár, Common Randomness in information theory and cryptography – part I: secret sharing, *IEEE Transactions on Information Theory*, Vol. IT–39, 1993, pp. 1121–1132.
2. C.H. Bennett, F. Bessette, G. Brassard, L. Salvail and J. Smolin, "Experimental quantum cryptography", *Journal of Cryptology*, Vol. 5, no. 1, 1992, pp. 3–28.
3. C.H. Bennett, G. Brassard, C. Crépeau, and U.M. Maurer, "Generalized privacy amplification", to appear in *IEEE Transactions on Information Theory*, Nov. 1995.
4. C.H. Bennett, G. Brassard and J.–M. Robert, "Privacy amplification by public discussion", *SIAM Journal on Computing*, Vol. 17, no. 2, April 1988, pp. 210–229.
5. R.E. Blahut, *Theory and Practice of Error Control Codes*, Reading, MA: Addison-Wesley, 1983.
6. R.E. Blahut, *Principles and Practice of Information Theory*, Reading, MA: Addison-Wesley, 1987.
7. J.L. Carter and M.N. Wegman, "Universal classes of hash functions", *Journal of Computer and System Sciences*, Vol. 18, 1979, pp. 143–154.
8. I. Csiszár and J. Körner, "Broadcast channels with confidential messages", *IEEE Transactions on Information Theory*, Vol. IT–24, no. 3, 1978, pp. 339–348.
9. W. Diffie and M.E. Hellman, "New directions in cryptography", *IEEE Transactions on Information Theory*, Vol. IT–22, 1976, pp. 644–654.
10. P. Gemmell and M. Naor, Codes for interactive authentication *Advances in Cryptology — Proceedings of Crypto '93*, Lecture Notes in Computer Science, Vol. 773, Springer–Verlag, Berlin, 1994, pp. 355–367.
11. E.N. Gilbert, F.J. MacWilliams, and N.J.A. Sloane, Codes which detect deception, *Bell Syst. Tech. J.*, Vol. 53, No. 3, 1974, pp. 405–424.
12. R.L. Graham, D.E. Knuth and O. Patashnik, *Concrete mathematics*, Reading, MA: Addison-Wesley, 1990.
13. U.M. Maurer, Protocols for secret key agreement by public discussion based on common information, *Advances in Cryptology - CRYPTO '92*, Lecture Notes in Computer Science, Berlin: Springer-Verlag, vol. 740, pp. 461–470, 1993.
14. U.M. Maurer, Secret key agreement by public discussion from common information, *IEEE Transactions on Information Theory*, vol. IT–39, 1993, pp. 733–742.
15. U.M. Maurer, The strong secret key rate of discrete random triples, *Communications and Cryptography, Two Sides of one Tapestry*, R.E. Blahut et al. (editors), Kluwer Academic Publishers, 1994, pp. 271–285.
16. U.M. Maurer and P.E. Schmid, A calculus for security bootstrapping in distributed systems, *Journal of Computer Security*, vol. 4, no. 1, pp. 55–80, 1996.
17. U.M. Maurer and S. Wolf, Towards characterizing when information-theoretic secret key agreement is possible, *Advances in Cryptology - ASIACRYPT '96*, K. Kim and T. Matsumoto (Eds.), Lecture Notes in Computer Science, Berlin: Springer-Verlag, vol. 1163, pp. 145–158, 1996.
18. U.M. Maurer and S. Wolf, The intrinsic conditional mutual information and perfect secrecy, to appear in *Proc. 1997 IEEE Symposium on Information Theory*, (Abstracts), Ulm, Germany, June 29–July 4, 1997.
19. U.M. Maurer and S. Wolf, Privacy amplification secure against active adversaries, preprint, 1997.
20. R.L. Rivest, A. Shamir, and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, *Communications of the ACM*, Vol. 21, No. 2, 1978, pp. 120–126.

21. C. E. Shannon, Communication theory of secrecy systems, *Bell System Technical Journal*, Vol. 28, October 1949, pp. 656–715.
22. G. J. Simmons, Authentication theory/coding theory, in *Advances in Cryptology – CRYPTO 84*, G.R. Blakley and D. Chaum (Eds.), Lecture Notes in Computer Science, No. 196, Berlin: Springer Verlag, 1985, pp. 411–431.
23. D. R. Stinson, Universal hashing and authentication codes, *Advances in Cryptology — Proceedings of Crypto '91*, Lecture Notes in Computer Science, Vol. 576, Springer–Verlag, Berlin, 1994, pp. 74 – 85.
24. M. N. Wegman and J. L. Carter, New hash functions and their use in authentication and set equality, *Journal of Computer and System Sciences*, Vol. 22, 1981, pp. 265–279.
25. A. D. Wyner, The wire-tap channel, *Bell System Technical Journal*, Vol. 54, no. 8, 1975, pp. 1355–1387.