Perfect Cryptographic Security from Partially Independent Channels *

Ueli M. Maurer †

Department of Computer Science
Princeton University
Princeton, NJ 08544

Abstract

Several protocols are presented that allow two parties Alice and Bob not sharing any secret information initially (except possibly a short key to be used for authentication) to generate a long shared secret key such that even an enemy Eve with unlimited computing power is unable to obtain a non-negligible amount of information (in Shannon's sense) about this key.

Two different models are considered. In a first model we assume that Alice can send information to Bob over a noisy main channel but that Eve is able to receive the same information over a parallel independent noisy channel from Alice to Eve. In a second, more general model we assume that Alice, Bob and Eve receive the output of a random source (e.g., a satellite broadcasting random bits) over three independent individual channels. The condition that the channels be independent can be replaced by the condition that they be independent only to a known, arbitrarily small degree.

We demonstrate that *even when Eve's channel is superior* (i.e., less noisy) to Alice's and Bob's channel(s), they can generate an information-theoretically secure secret key by communicating over a public (error-free) channel to which Eve is assumed to have unrestricted access. The results of this paper suggest to base the security of cryptographic systems on realistic statistical assumptions about the partial independence of two (three) channels and about a reasonable lower bound on the noise power on the enemy's channel, as an alternative to commonly used approaches based on an intractability hypothesis.

The paper suggests two general conclusions:

(1) for cryptographic purposes, a given noisy communication channel should not be converted into an error-free channel (by means of error-correcting codes) on which a conventional cryptographic protocol is executed, but rather *cryptographic coding and error-control coding should be combined*, and

(2) a mere *difference* in the signals received by the enemy and the legitimate receiver, but not necessarily with an advantage to the receiver (such as his sharing of a secret key with a sender or knowledge of a trapdoor), may be sufficient for achieving cryptographic security. This observation seems to have broader applications in cryptography.

1. Introduction

One of the classical problems in cryptography is to transmit a message M securely to a legitimate receiver such that an enemy is unable to obtain useful information about M. In the classical model of a cryptosystem introduced by Shannon [17], the enemy has access to the insecure transmission channel; thus

*Preprint of a paper to appear in the Proceedings of the 23rd Annual ACM Symposium on Theory of Computing (STOC '91), New Orleans, May 6-8, 1991.

†Supported by a postdoctoral fellowship from the Swiss National Science Foundation.

he is assumed to receive the entire ciphertext C, i.e., an identical copy of the message received by the legitimate receiver. Shannon defined a cipher system to be *perfect* when the ciphertext gives no information about the plaintext, i.e., when

$$I(M;C) \;\stackrel{\triangle}{=}\; H(M) - H(M|C) \;=\; 0$$

or, equivalently, when M and C are statistically independent. $H(X)$ denotes the entropy of the random variable X and is defined [16] by

$$H(X) \;=\; -\sum_{x} P(X=x)\log_2 P(X=x),$$

and $H(X|Y)$ denotes the conditional entropy of X when given the random variable Y and is defined by

$$H(X|Y) = -\sum_{(x,y)} P(X=x,Y=y)\log_2 P(X=x|Y=y).$$

We refer to [8] for an introduction to information theory.

Shannon proved that a necessary condition for a cipher to be perfect is that the secret key K, which is shared initially by the legitimate communicants and about which the enemy is assumed to have no *a priori* information, satisfies

$$H(K) \;\geq\; H(M), \tag{1}$$

i.e., the number of binary digits in the secret key must be at least as great as the number of bits of information in the message.

A perfect cipher is *unconditionally secure* because even an enemy with unlimited computational resources is unable to break it. In virtually all applications it is completely impractical to use a secret key that satisfies (1). Because of Shannon's pessimistic inequality (1), perfect secrecy is often prejudged as being impractical. It is one of the main goals of this paper to relativize this pessimism by pointing out that Shannon's analysis assumes that, except for the secret key, the enemy has access to precisely the same information as the legitimate receiver, and that this apparently innocent assumption is *much more restrictive than has generally been realized.*

Virtually all presently-used ciphers use a short secret key and can therefore theoretically be broken, for instance by an exhaustive key search. The goal of designing such a practical cipher is therefore to prove that there exists *no efficient algorithm* for breaking it, for a reasonable definition of breaking. However, for no existing cipher can the security be proved without invoking an unproven intractability hypothesis. For instance the security of the famous Diffie-Hellman public-key distribution system [7] is based on the (unproven) difficulty of the discrete logarithm problem in the multiplicative group modulo a large prime p for which $p-1$ also has a large prime factor. Surprisingly, even the *existence* of *computationally secure* ciphers is an open problem, but it has been proved [13] that, for virtually all definitions of breaking a cipher and of computational difficulty, the cascade of several additive binary stream ciphers is at least as computationally secure as the strongest component cipher.

Information-theoretic or unconditional security is more desirable in cryptography than computational security for two reasons. First, for the former no assumption about the enemy's computing power need be made, and second, no definition of security need be given because the utmost one can wish to prove is that the enemy has no information (in Shannon's sense) about the plaintext. Note that there are many possible definitions for the computational security of a cipher, and it is not obvious which one to choose.

As mentioned above, one way of achieving unconditional cryptographic security is the use of an impractically large amount of secret key. Some alternative approaches try to *ensure* that an enemy cannot obtain the same information as the legitimate receiver, but these systems are at the moment still not truly practical. Quantum cryptography introduced by Wiesner and put forward by Bennett, Brassard *et al.* [1, 5], which is not practical (although a prototype exists) because it requires the transmission of single photons, is based on the uncertainty principle of quantum physics. Maurer's strongly randomized cipher makes use of a public random string that is too long to be read entirely in feasible time and is impractical because no source of such an immense amount of randomness has yet been discovered. Both these systems allow two parties initially sharing a short secret key to generate a much longer and unconditionally secure shared secret key. In the first of these systems, the secret key is required for authentication (as in a realistic implementation of our protocols) and in the second system it is used to select a feasible number of random bits to generate the keystream.

Some other approaches to unconditional cryptographic security are based on the *assumption* that the enemy is restricted in the amount of information about the ciphertext that he can obtain. The drawback of these approaches is that such an assumption is usually completely *unrealistic*. Perfect local randomizers considered by Maurer and Massey [14] are based on the unrealistic assumption that an enemy can obtain only a small number of ciphertext bits. Wyner [20] and subsequently Csiszár and Körner [6] considered a situation where the enemy receives the message transmitted by a sender over a channel that is *noisier* than the legitimate receiver's channel. The assumption that the enemy's channel is worse than the main channel may be reasonable in an application where the quality of the enemy's channel can be controlled (for instance by monitoring the received signal power at the output of an optical fiber and thus limiting the signal power that an enemy can extract from the fiber [9] without being detected), but is unrealistic in general.

In this paper, new approaches to provable information-theoretic security are presented that are based on novel and generally much more realistic assumptions about the enemy's obtainable information. Our assumptions are (1) that the noise power on the enemy's channel is at least a certain known *fraction* of the noise power on the legitimate user's channel(s) and (2) that the noise on the enemy's channel is at least to some known, but arbitrarily small degree *independent* of the noise on the other channel(s). A first model is introduced in Section 2 in which Alice can send information to Bob through a noisy main channel where Eve (the enemy) receives the same information through an independent (possibly superior) channel. In Section 3 it is demonstrated that the previously needed condition that Eve's channel be worse than the main channel can be removed. An analysis is given for the case of binary symmetric channels (BSC). In the generalized model of Section 4, Alice, Bob and Eve are assumed to receive the output of a random source (for example a satellite broadcasting random bits) over independent individual channels, where again Eve's channel may be much more reliable than the other two channels. A further generalization and some open problems are suggested in Section 5.

An essential feature of our protocol is the use of an (error-free) *public communication channel* between Alice and Bob. It is assumed that all messages sent over the public channel are received by Eve without error, but that she cannot corrupt the messages or introduce faked messages. Authentication and data integrity can be ensured by using an unconditionally secure authentication scheme [19] based on universal hashing, which requires that Alice and Bob share a short secret key beforehand. In this case, the purpose of our protocol is to stretch (rather than to generate) a secret key unconditionally securely. Part of the new key can be used for authentication in a subsequent instance of the protocol. The use of a public channel by two parties for extracting a secret key from an initially shared partially secret string was previously considered by Leung-Yan-Cheong [10] and independently by Bennett, Brassard and Robert [2]. Another somewhat related paper is [15].

A novel technique used in our protocols is the conversion of the error-free public channel into a *conceptual noisy broadcast channel* by combining noisy information from the actual noisy channels with information to be sent over the conceptual channel, and sending the result over the public channel. Another novel technique is that of *reliability estimation* which allows the receiver at the output of a noisy channel to *select* only those messages that have been received sufficiently reliably and to discard the other messages.

This paper is concerned with key distribution rather than encryption. Note however that an unconditionally secure secret key can be used as the one-time pad in the well-known perfect Vernam cipher [18].

2. Secret Communication over Broadcast Channels

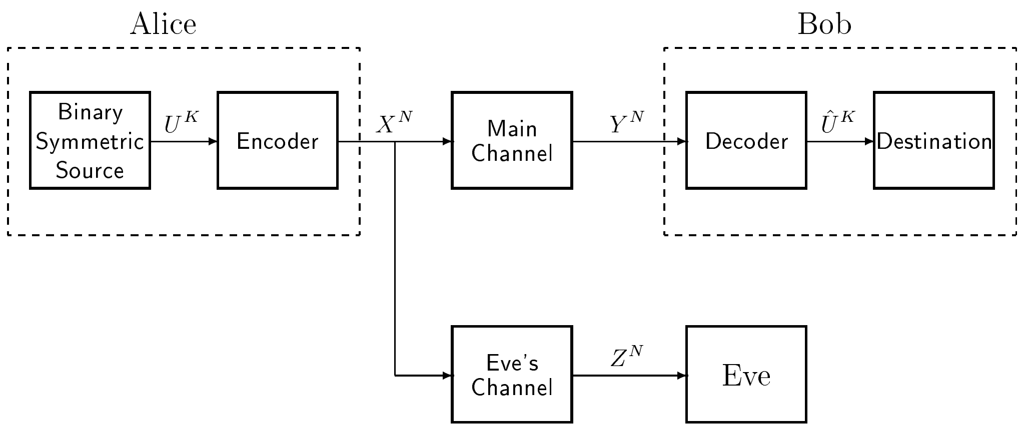

Figure 1. Model of the discrete memoryless broadcast channel introduced by Csiszár and Körner [6]. The common input to the main channel and to Eve's channel is the random variable X sent by Alice. The superscripts K and N indicate that the information word U^K is encoded into a codewords X^N which is transmitted by N (independent) uses of the channel. The channel behavior is completely specified by the conditional probability distribution $P_{YZ|X}$.

In his celebrated paper [20], Wyner considered a situation in which Alice can send information to Bob over a discrete memoryless channel (DMC) [8] such that a wire-tapper Eve can receive Bob's output only through an additional cascaded independent DMC. A channel is memoryless if successive uses of the channel are independent of each other. Wyner's setting is in information theory called a *degraded* discrete memoryless broadcast channel. Wyner proved that in such a (generally unrealistic) setting Alice can send information to Bob in virtually perfect secrecy, even though Alice and Bob share no secret key initially, thus seemingly contradicting Shannon's result (1).

Wyner's model and results were generalized by Csiszár and Körner [6] who considered a *discrete memoryless broadcast channel* (cf. Figure 1) in which the wire-tapper Eve's received message is not necessarily a degraded version of the legitimate receiver's message. The common input to the main channel and Eve's channel is the random variable X chosen by Alice according to some probability distribution P_X, and the random variables received by the legitimate receiver Bob and by the enemy Eve are Y and Z, respectively. The superscript N in Figure 1 indicates that the channel is used (independently) N times to transmit an N-digit codeword X^N resulting from encoding a K-bit information word U^K. Without loss of generality the information word is assumed to be a binary random sequence. The channel behavior is completely specified by the conditional probability distribution $P_{YZ|X}$. Note that in Wyner's setting, X, Y and Z form a Markov chain, i.e., $P_{Z|XY} = P_{Z|Y}$. (Equivalently, $I(X; Z|Y) = 0$.)

The *secrecy capacity* C_s of the broadcast channel of Figure 1 is defined as the maximum rate K/N at which Alice can reliably send information to Bob such that the rate at which Eve obtains information is arbitrarily small. In other words, the secrecy capacity is the maximal number of bits that Alice can secretly send to Bob *per use of the channel*. More formally, the secrecy capacity can be defined as the maximal rate $R = K/N$ such that for every $\gamma > 0$ and for sufficiently large N there exists a code of length N with 2^K codewords together with a decoding function $d: \mathcal{Y}^N \to \{0,1\}^K : Y^N \mapsto \hat{U}^K = d(Y^N)$ such that $P[U^K \neq \hat{U}^K] < \gamma$ and $H(U^K|Z^N)/K > 1-\gamma$ where $\mathcal{Y}$ denotes the main channel output alphabet.

For very general probability distributions $P_{YZ|X}$ (and more generally for every pair of conditional probability distributions $(P_{Y|X}, P_{Z|X})$ such that

$P_{YZ|X}$ is not necessarily defined), the secrecy capacity C_s of the broadcast channel of Figure 1 is given by

$$\begin{aligned} C_s &= \max_{P_X} \,[I(X;Y) - I(X;Z)] \\ &= \max_{P_X} \,[H(X|Z) - H(X|Y)] \end{aligned} \qquad (2)$$

(see [6]), where $I(U;V) \stackrel{\triangle}{=} H(U) - H(U|V)$ denotes the (mutual) information [8] that the random variable V gives about the random variable U (and vice versa). Moreover, the intuitive result holds that the secrecy capacity is zero unless the enemy's channel is noisier than the main channel. In other words, a message can be transmitted secretly from Alice to Bob only under the generally unrealistic assumption that the enemy's channel is worse than the main channel. It is one of the achievements of this paper that this unrealistic assumption is no longer needed.

As a motivating example in which the use of feedback from Bob to Alice allows Alice and Bob to generate a secret key even when Eve's channel in Figure 1 is superior to the main channel from Alice to Bob (and thus provably no secret key can be generated without feedback), let both channels be additive white Gaussian noise (AWGN) channels with statistically independent noise. For example, Alice's sender could be on a satellite. Assume further that Alice uses binary antipodal signaling to transmit an uncoded sequence of independent and completely random bits ($K = N$ in Figure 1). In order to convert the enemy's advantage into a disadvantage Bob picks only those bits out of the data stream that he receives very reliably, but discards less reliable bits. Note that since the receiver output is analog (Gaussian distribution with mean +1 or -1 according to the bit sent, and with variance proportional to the noise power) rather than two-level quantized, the reliability of a decision about the bit sent by Alice can be determined as a function of the absolute value of the receiver's output. Bob now uses the public feedback channel to inform Alice of which bits he picked. Although Bob's bit error probability is on the average worse than Eve's bit error probability, it is nevertheless better when averaged only over the selected bits. Note that, by the independence of the two channels, knowledge of the positions of the bits received reliably by Bob gives no information to Eve about the values of these bits. By adding modulo 2 several of the selected bits or, more generally, by applying an appropriate linear hashing function, Alice and Bob can now reduce Eve's information about the extracted bits to an arbitrarily small fraction of a bit while keeping the probability that Alice's and Bob'b extracted strings disagree sufficiently small.

Clearly, the protocol just described is completely impractical when the main channel is much worse than the enemy's channel because the rate at which Alice and Bob can agree on secret bits is extremely small since the event that Bob receives a bit sent by Alice more reliably than Eve has very small probability. A more efficient protocol is described in the following section for the more interesting case of binary symmetric channels.

3. Key Distribution Protocols for Binary Symmetric Broadcast Channels

In this section, the important special case of the broadcast channel of Figure 1 is considered where both the main channel and Eve's channel are binary symmetric channels (BSC) with bit error probabilities ϵ and δ, respectively. Without loss of generality we assume that $\epsilon \leq 1/2$ and $\delta \leq 1/2$. We first assume that the two channels are independent, but the more general case of dependent channels is discussed later.

When no feedback is allowed, the secrecy capacity can be determined according to (2). Let $h(.)$ denote the binary entropy function defined by

$$h(x) = -x \log_2 x - (1-x)\log_2(1-x)$$

for $0 < x < 1$ and by $h(0) = h(1) = 0$.

Theorem 1. *The secrecy capacity (without feedback) of the broadcast channel of Figure 1, where the main channel is a BSC with error probability $\epsilon \leq 1/2$ and Eve's channel is an independent BSC with error probability $\delta \leq 1/2$, is given by*

$$C_s = \begin{cases} h(\delta) - h(\epsilon) & \text{if } \delta > \epsilon, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. P_X is characterized by the single parameter $P(X=0)=p$. (It follows that $P(X=1)=1-p$.) The entropy of a binary random variable taking on the two values with probabilities p and $1-p$, respectively, is given by $h(p)$. $H(XY)$ can be expressed in two different ways, viz. $H(XY)=H(X)+H(Y|X)=H(Y)+H(X|Y)$. Hence $H(X|Y)=H(X)-H(Y)+H(Y|X)$. We have $H(X)=h(p)$, $H(Y|X)=h(\epsilon)$ because $H(Y|X=x)=h(\epsilon)$ independent of x, and $H(Y)=h(p+\epsilon-2p\epsilon)$ since $P(Y=0)=p(1-\epsilon)+(1-p)\epsilon=p+\epsilon-2p\epsilon$. Thus $H(X|Y)=h(p)+h(\epsilon)-h(p+\epsilon-2p\epsilon)$ and similarly one obtains that $H(X|Z)=h(p)+h(\delta)-h(p+\delta-2p\delta)$. The theorem follows from the fact for every $p<1/2$, $h(p+\xi-2p\xi)$ is a monotonically increasing function for $0\le\xi<1/2$ such that $H(X|Z)-H(X|Y)=h(\delta)-h(\epsilon)+h(p+\epsilon-2p\epsilon)-h(p+\delta-2p\delta)$ is maximized for $p=1/2$, in which case the last two terms vanish. □

It should be pointed out that the proofs for (2) given in [6] and [20] are non-constructive existence proofs based on a random-coding argument. We also refer to [11] for a simplified treatment of Wyner's wire-tap channel. The problem of finding actual efficiently encodable and decodable codes that perform well in a particular situation is not solved in general. However, one can often find practical codes that achieve a constant fraction of the secrecy capacity. The situation is thus similar as for the well-known problem of designing practical error-correcting codes [4], which is actually a special case of our problem: it is known that for a given channel codes with rate arbitrarily close to the channel capacity [8] exist whose block error probability (for the optimal decoder) is arbitrarily small, but it is not known whether efficiently encodable and decodable such codes exist.

We now return to our problem of using a public channel to generate a secret key shared by Alice and Bob, even when Eve's channel is superior to the main channel. Unfortunately, the method discussed in the previous section for AWGN channels based on a reliability estimation does not work for BSCs because every bit received through a BSC is equally reliable. More precisely, the conditional probability distribution over the inputs, given the output of the channel, is the same (up to a permutation of the inputs) for every output digit. However, when we consider $N>1$ consecutive uses of the BSC as one use of a conceptual super-channel with 2^N inputs and 2^N outputs, a reliability decision can be made when only a subset of the inputs are chosen with non-zero probability. This approach, which will be discussed in more detail in the next section, is equivalent to the use of an error-correcting code. We now present a more efficient protocol that works well when the channels are sufficiently independent

For a given broadcast channel specified by $P_{YZ|X}$ we define the *secrecy capacity with public discussion*, denoted as $\overline{C}_s$, as the maximal rate (in bits per channel use) at which Alice and Bob can agree on a secret key using arbitrary public discussion such that Eve receives only an arbitrarily small amount of information. The following theorem gives a general upper bound on $\overline{C}_s$.

Theorem 2. *The secrecy capacity with public discussion of a discrete memoryless broadcast channel specified by $P_{YZ|X}$ is upper bounded by*

$$\overline{C}_s \le \max_{P_X} I(X;Y|Z) \stackrel{\triangle}{=} \max_{P_X}[H(X|Z)-H(X|YZ)],$$

i.e., by the mutual information between X and Y, given that Z is known.

Proof. Assume that Eve provides Z to Alice and Bob for free. Knowing Z cannot reduce the secrecy capacity because Alice and Bob could always discard Z. The remaining amount of information shared by Alice and Bob, $I(X;Y|Z)$, is thus an upper bound on the amount of information shared by Alice and Bob in secrecy. □

It should be pointed out that the upper bound $\max_{P_X} I(X;Y|Z)$ on $\overline{C}_s$ is not achieved in general. Consider as an example a binary broadcast channel that selects Y randomly and independently of X and forms $Z=X+Y$. Clearly, $I(X;Y)=0$ and therefore also $\overline{C}_s=0$, but note that $I(X;Y|Z)=1$ since, when Z is given, additionally giving Y uniquely determines X. It is therefore somewhat surprising that the upper bound of Theorem 2 is achieved in the case of independent binary symmetric channels:

Theorem 3. *The secrecy capacity with public discussion of the broadcast channel of Figure 1, where the main channel is a BSC with error probability $\epsilon \leq 1/2$ and Eve's channel is an independent BSC with error probability $\delta \leq 1/2$, is given by*

$$\overline{C}_s \;=\; \max_{P_X} \; I(X;Y|Z) \;=\; h(\epsilon + \delta - 2\epsilon\delta) - h(\epsilon).$$

Moreover, $\overline{C}_s$ is strictly positive unless $\epsilon = 0.5$ or $\delta = 0$, i.e., unless X and Y are statistically independent or the enemy receives $Z = X$, respectively.

Proof. Note first of all that $I(X;Y|Z) \stackrel{\triangle}{=} H(Y|Z) - H(Y|XZ) = H(Y|Z) - H(Y|X)$ where $H(Y|XZ) = H(Y|X)$ follows from the independence of the channels. By an argument similar to that used in the proof of Theorem 1 one can show that $H(Y|Z) - H(Y|X)$ is maximized for the choice $P_X(0) = P_X(1) = 1/2$ for which it takes on the value $h(\epsilon+\delta-2\epsilon\delta) - h(\epsilon)$. The term $\epsilon+\delta-2\epsilon\delta$ corresponds to the error probability of a BSC that is the cascade of two independent BSC's with error probabilities ϵ and δ.

The key observation for proving that the upper bound can be achieved is that, by appropriately choosing the feedback message, Bob can send bits to Alice (and to Eve since the feedback channel is assumed to be public) such that Alice receives these bits with error probability ϵ, whereas Eve receives the bits with the larger error probability $\epsilon + \delta - 2\epsilon\delta$. In other words, Bob can create a *conceptual broadcast channel* in which the conceptual channel to Alice is equivalent to the actual main channel from Alice to Bob but Eve's conceptual channel is equivalent to a cascade of the actual main channel and Eve's actual channel. Bob can use this conceptual backward channel to send secret information to Alice at a rate equal to its secrecy capacity, which is according to Theorem 1 equal to $h(\epsilon + \delta - 2\epsilon\delta) - h(\epsilon)$.

To illustrate how Bob can send a codeword V^N over this conceptual backward channel, assume that Alice transmits a sequence X^N of independent and completely random bits, i.e., the encoder in Figure 1 is replaced by a straight wire and $K = N$. The noise on the channels is for each use characterized by two binary, statistically independent random variables D and E with $P(D = 1) = \delta$ and $P(E = 1) = \epsilon$, where $Y = X + E$ and $Z = X + D$. Here and in the sequel, addition (except addition of entropies) is modulo 2 and vector addition is componentwise. Thus the blocks received by Bob and Eve are $X^N + E^N$ and $X^N + D^N$, respectively. Bob adds the received word $Y^N = X^N + E^N$ to V^N and sends $W^N = V^N + Y^N = X^N + V^N + E^N$ over the public channel. Alice, who knows X^N, can add X^N to W^N and thereby obtains $V^N + E^N$, which is equivalent to receiving V^N with error probability ϵ. Eve on the other hand knows $Z^N = X^N + D^N$ and $W^N = X^N + V^N + E^N$ and can do no better than to compute $Z^N + W^N = V^N + E^N + D^N$ and to discard Z^N and W^N, which is equivalent to receiving V^N over a cascade of the actual main channel and Eve's actual channel.

In order to prove that without loss of optimality Eve can form $Z^N + W^N$ and discard Z^N and W^N, we show that $H(V^N|Z^N W^N) = H(V^N|Z^N + W^N)$ as follows.

$$\begin{aligned} H(V^N|Z^N W^N) \;&=\; H(V^N|Z^N + W^N, W^N) \\ &= H(V^N W^N|Z^N + W^N) - H(W^N|Z^N + W^N) \\ &= H(V^N|Z^N + W^N) + H(W^N|V^N, Z^N + W^N) \\ &\quad - H(W^N|Z^N + W^N). \end{aligned}$$

The first step follows from the fact that the pair (Z^N, W^N) uniquely determines the pair $(Z^N + W^N, W^N)$ and vice versa and the other steps are applications of standard information-theoretic identities [8]. The result now follows upon noting that $H(W^N|V^N, Z^N{+}W^N) = H(X^N{+}V^N{+}D^N|V^N, V^N{+}E^N{+}D^N) = N$ and thus also $H(W^N|Z^N{+}W^N) = N$ since X^N is completely random and statistically independent of V^N, E^N and D^N.

To prove the last claim, note that $h(x)$ is a monotonically increasing function for $0 \leq x < 1/2$, and that $\epsilon + \delta - 2\epsilon\delta \geq \epsilon$ with equality if and only if either $\delta = 0$ or $\epsilon = 1/2$. □

We now consider the more general case of dependent channels. When the main BSC and Eve's BSC are dependent, this situation is equivalent to the following model with three independent BSCs. The main channel and Eve's channel each consist of a cascade of two BSCs, where the first BSC with error probability ϵ_A is common for both Bob and Eve, and where the two cascaded channels have error probabilities ϵ_B and ϵ_E, respectively. The main and Eve's channels have error probabilities ϵ and δ, respectively, and hence we have $\epsilon = \epsilon_A + \epsilon_B - 2\epsilon_A\epsilon_B$ and $\delta = \epsilon_A + \epsilon_E - 2\epsilon_A\epsilon_E$. Note that by creating conceptual channels as described above, Bob can create a situation that is symmetric to the one considered here in which Bob is the sender. Thus a protocol similar to the one described above for independent channels can be used when either $\epsilon_A < \epsilon_E$ or $\epsilon_B < \epsilon_E$. The proof of the following theorem is similar to that of Theorem 3. The lower bound is positive if and only if either $\epsilon_A < \epsilon_E$ or $\epsilon_B < \epsilon_E$.

Theorem 4. *The secrecy capacity with public discussion of the broadcast channel discussed above satisfies*

$$\overline{C}_s \geq \max\{h(\epsilon_A+\epsilon_E-2\epsilon_A\epsilon_E), h(\epsilon_B+\epsilon_E-2\epsilon_B\epsilon_E)\} \\ -h(\epsilon_A + \epsilon_B - 2\epsilon_A\epsilon_B).$$

We have demonstrated a method for converting a situation in which Eve's channel is superior to the main channel into a situation where her channel is inferior, which can be exploited according to [6]. However, as mentioned earlier, the proofs given in [6] are non-constructive, and it is an open problem whether efficiently encodable and decodable codes exist that exploit the full secrecy capacity. A method that is different from that proposed in [6] and [20] for exploiting the availability of a superior channel is to use a reconciliation and universal hashing information reduction protocol as described in [2]. One open problem that will be addressed in [3] is to generalize the proof for the information reduction protocol given in [2] to cases where Eve knows K bits of information rather than the output of a 2^K-valued function.

4. Generating a Mutual Secret Key from Randomness Received over Independent Channels

Consider a random source (e.g., a satellite broadcasting random bits) that is received by Alice, Bob and Eve over three at least to some degree independent individual channels. In this section we present a protocol by which Alice and Bob can generate a mutual secret key by public discussion after individually receiving a noisy version of the same sequence of random bits from the source. The generated secret key is information-theoretically secure. Thus even an enemy with unlimited computing power cannot obtain more than a negligible amount of information in Shannon's sense about the key. Somewhat surprisingly, the protocol works even for cases in which Eve's channel is better than Alice's and Bob's channel, i.e., even when Eve's version of the same random string contains less errors.

Because thermal noise in different receivers is statistically independent and also the atmospheric noise is to some degree independent for Alice, Bob and Eve, the above assumption appears to be realistic in many cases. An additional assumption about Eve's minimal error probability (noise power) can be realistically mild because Eve's channel need not be assumed to be worse than the Alice's and Bob's channels. The significance of the protocol presented below is therefore that it allows to base cryptographic security on a realistic statistical assumption as an alternative to the common approach of founding the security on an unproven intractability hypothesis.

Consider a binary symmetric source emitting a sequence $X_1, X_2, \ldots$ of random bits. Alice, Bob and Eve receive noisy versions of these bits, i.e., Alice receives $X_i + F_i$, Bob receives $X_i + G_i$ and Eve receives $X_i + H_i$ for $i = 1, 2, \ldots$, where F_i, G_i and H_i are in the following assumed to be statistically independent random variables taking on the value 1 with probabilities ϵ_A, ϵ_B and ϵ_E, respectively, and where addition is modulo 2. When F_i, G_i and H_i are not completely independent, our protocol still works but the analysis may become more involved. Without loss of generality we assume that $\epsilon_A, \epsilon_B, \epsilon_E \leq 1/2$. However, it is not assumed that $\epsilon_E \geq \epsilon_A$ or $\epsilon_E \geq \epsilon_B$, i.e., that either Alice or Bob has an advantage compared to Eve. Let $\delta_A = 1 - \epsilon_A, \delta_B = 1 - \epsilon_B$ and $\delta_E = 1 - \epsilon_E$.

In order to create a situation conceptually similar to that considered in the previous section where Alice can send bits V_i to Bob over a main channel such that Eve receives these bits over another (possibly better) channel, Alice can send the sum $V_i + X_i + F_i$ over the public channel. Both Bob and Eve can recover a noisy version of V_i by adding their respective noisy versions of X_i, i.e., Bob can compute $V_i + F_i + G_i$ and Eve can compute $V_i + F_i + H_i$. By an argument similar to that used in the proof of Theorem 3 one can show that without loss of optimality Bob and Eve can discard all other received information. Note that the noise seen by Bob and Eve on their respective conceptual channels is $F_i + G_i$ and $F_i + H_i$, respectively, and is not independent. The situation is thus equivalent to the broadcast channel considered in Section 3 (cf. Figure 1) but where now the main channel and Eve's channel are not independent. When both $\epsilon_E < \epsilon_A$ and $\epsilon_E < \epsilon_B$ then the lower bound of Theorem 4 is 0 and the protocol of the previous section is hence useless.

Nevertheless the following protocol allows Alice and Bob to generate a secret key. Alice randomly selects a codeword V^N from the set of codewords of an appropriate error-correcting code $\mathcal{C}$ with codewords of length N and sends it to Bob over the above mentioned conceptual channel. Bob receives $V^N + F^N + G^N$ and Eve receives $V^N + F^N + H^N$. In order to create an advantage over the enemy even when $\epsilon_B > \epsilon_E$, Bob only accepts a block when it is received very reliably, i.e., when it is very close to an actual codeword of $\mathcal{C}$. Bob then publicly announces which blocks he accepted. Although Eve receives codewords V^N more reliably than Bob on the average, her conceptual channel is nevertheless worse than Bob's channel if averaged only over those codewords accepted by Bob. Because consecutive uses of the channel are independent, the blocks discarded by Bob are also useless for Eve. One problem with this approach is that for a general code it seems very difficult to compute the amount of information obtained by Eve, but in some cases bounds can be given.

Note that Alice cannot use this restricted (to certain block) broadcast channel because she does not know in advance which blocks will be received reliably by Bob. However, the advantage created by the reliability estimation and block selection can be exploited by Alice and Bob by using a reconciliation and information reduction protocol as described in [2] or by creating yet another conceptual channel from Bob back to Alice (and to Eve).

Consider now the following special case. Alice uses a repeat code of length N, i.e., for $j = 1, 2, \ldots$ she randomly selects an information bit R_j and sends $V_j^N = 00\ldots0$ if $R_j = 0$ and $V_j^N = 11\ldots1$ otherwise over the conceptual channel to Bob. Bob accepts a received block if and only if it is exactly equal to one of the codewords, i.e., if and only if the block consists of either only 0's or only 1's. Although a sufficiently reliable decision could also be made in the case where the received block contains only few 0's or few 1's, such blocks are discarded in order to deprive Eve of blocks that might be more useful for her than for Bob. The probability that a codeword is received by Bob without error is given by

$$p_{\text{correct}} = (\delta_A\delta_B + \epsilon_A\epsilon_B)^N$$

and similarly the probability that a codeword is received as its complement equals

$$p_{\text{error}} = (1 - \delta_A\delta_B - \epsilon_A\epsilon_B)^N.$$

The probability that Bob accepts a codeword is thus given by $p_{\text{accept}} = p_{\text{correct}} + p_{\text{error}}$. The channel from Alice to Bob thus corresponds to a binary symmetric channel with error probability $\beta = p_{\text{error}}/p_{\text{accept}}$.

Let α_{rs} for $r, s \in \{0, 1\}$ be the probability that a single bit 0 sent by Alice is received by Bob as r and by Eve as s. Thus

$$\begin{aligned}
\alpha_{00} &= \delta_A\delta_B\delta_E + \epsilon_A\epsilon_B\epsilon_E, \\
\alpha_{01} &= \delta_A\delta_B\epsilon_E + \epsilon_A\epsilon_B\delta_E, \\
\alpha_{10} &= \delta_A\epsilon_B\delta_E + \epsilon_A\delta_B\epsilon_E \\
\text{and} \quad \alpha_{11} &= \delta_A\epsilon_B\epsilon_E + \epsilon_A\delta_B\delta_E.
\end{aligned}$$

Let further p_w for $0 \le w \le N$ be the probability that a $00\ldots0$ codeword sent by Alice is accepted by Bob (correct or not) and is received by Eve as a particular given block of Hamming weight w. We have

$$p_w = \alpha_{00}^{N-w}\alpha_{01}^{w} + \alpha_{10}^{N-w}\alpha_{11}^{w}.$$

Eve's error probability when she guesses the bit sent by Alice is

$$\gamma = \sum_{w=\lceil N/2\rceil}^{N} \binom{N}{w} p_w.$$

In order to allow for a successful protocol, N should be chosen such that $\gamma > \alpha$. However, it is important to note that this condition is not sufficient because Eve is not forced to make a hard decision about the bit but can rather wait until the end of the protocol to make a *soft* decision about the final string shared by Alice and Bob based on the entire information that she collected. The relevant quantities are therefore the capacities C_{AB} and C_{AE} of the restricted channels from Alice to Bob and to Eve, averaged only over those blocks accepted by Bob. The capacity of a BSC with error probability ϵ is equal to $1-h(\epsilon)$ [8], hence

$$C_{AB} \;=\; 1-h(\beta).$$

C_{AE} can be computed as the average capacity of $\sum_{w=0}^{\lfloor N/2\rfloor}\binom{N}{w}$ BSCs where each channel is weighted with the probability $(p_w+p_{N-w})/p_{\text{accept}}$ that it is used:

$$C_{AE} = \sum_{w=0}^{\lfloor N/2\rfloor}\binom{N}{w}\frac{p_w+p_{N-w}}{p_{\text{accept}}}\left(1-h(\frac{p_w}{p_w+p_{N-w}})\right).$$

When $C_{AB} > C_{AE}$, Alice and Bob can use *a posteriori* error-correction as proposed in [2] with subsequent universal hashing to extract a secret key about which Eve's information is arbitrarily small.

Example: Consider a case in which Eve receives the output of the random source more reliably than Alice and Bob; let $\epsilon_A=\epsilon_B=0.2$ and $\epsilon_E=0.15$. Note that Alice and Bob each receive $1-h(\epsilon_A)=0.278$ bits of information about each random bit whereas Eve receives $1-h(\epsilon_E)=0.39$, i.e., 40 % more. Let $N=5$. Then $p_{\text{correct}}=0.14539$, $p_{\text{error}}=0.003355$, $p_{\text{accept}}=0.14875$, $\alpha_{00}=0.55$, $\alpha_{01}=0.13$, $\alpha_{10}=\alpha_{11}=0.16$, $p_0=0.05043$, $p_1=0.01200$, $p_2=0.002917$, $p_3=0.0007695$, $p_4=0.00026619$ and $p_5=0.00014198$. Hence $\beta=2.25\%$ compared to $\gamma=6.63\%$ and thus Bob receives the selected bits much more reliably than Eve. One further obtains $C_{AB}=0.845$ and $C_{AE}=0.745$, i.e., Eve's capacity is 12% less than Bob's capacity.

When this protocol is used by Alice and Bob and when an ideal code were used to eliminate the remaining errors, then the number of random bits from the source that are required to generate one bit of shared secret key is ideally given by $N/[(C_{AB}-C_{AE})p_{\text{accept}}]$ which for the above example evaluates to 336.

Clearly there exist more efficient protocols for the above choice of error probabilities when more reasonable codes are used, but an analysis becomes very tedious. It seems also possible to gain secrecy at the cost of analyzability by using protocols with several rounds of interaction, making repeated use of conceptual channels and reliability estimation.

5. Generalizations and Open Problems

The settings of the previous two sections can be generalized in the following way. Alice, Bob and Eve receive (repeatedly at discrete time instances) random variables X, Y and Z, respectively. Rather than being the output of some channels, these random variables are assumed to be generated according to some given joint probability distribution P_{XYZ}. We define the secrecy capacity of X and Y with respect to Z, denoted $S(X,Y|Z)$, as the maximal rate (per use of a set of such random variables) at which Alice and Bob can by public discussion generate key bits about which Eve has an arbitrarily small amount of information. The following result is stated without proof.

Theorem 5. *The secrecy capacity of X and Y with respect to Z is lower bounded by*

$$S(X,Y|Z) \;\geq\; \max\left\{\max_{P_{V|X}}[I(V;Y)-I(V;Z)]\ ,\ \max_{P_{W|Y}}[I(W|X)-I(W;Z)]\right\}.$$

where the inner maximizations are over choices of the indicated conditional probability distributions.

We believe that the observations made in this paper are of rather fundamental nature in cryptography, theoretical computer science as well as in information theory. Many theoretical and practical problems remain open: characterizing the secrecy capacity for general settings, finding codes that perform well theoretically or also practically in a given setting, and determining the importance of the number of rounds of interaction.

Acknowledgements

Major parts of this work were performed while the author was with the Institute for Signal and Information Processing, Swiss Federal Institute of Technology, Zürich. The support provided by the Swiss National Science Foundation is gratefully acknowledged. I would like to thank Charles Bennett, Gilles Brassard and James Massey for helpful discussions and comments.

References

[1] C.H. Bennett, F. Bessette, G. Brassard, L. Salvail and J. Smolin, Experimental quantum cryptography, to appear in *Journal of Cryptology.*

[2] C.H. Bennett, G. Brassard and J.-M. Robert, Privacy amplification by public discussion, *SIAM Journal on Computing*, Vol. 17, No. 2, 1988, pp. 210-229.

[3] C.H. Bennett, G. Brassard, U.M. Maurer and L. Salvail, Generalized privacy amplification, in preparation.

[4] R.E. Blahut, *Theory and Practice of Error Control Codes*, Reading, MA: Addison-Wesley, 1984.

[5] G. Brassard, *Modern Cryptology: A Tutorial*, Lecture Notes in Computer Science, Vol. 325, Berlin: Springer-Verlag, 1988.

[6] I. Csiszár and J. Körner, Broadcast channels with confidential messages, *IEEE Transactions on Information Theory*, Vol. 24, No. 3, 1978, pp. 339-348.

[7] W. Diffie and M.E. Hellman, New directions in cryptography, *IEEE Transactions on Information Theory*, Vol. 22, No. 6, 1976, pp. 644-654.

[8] R.G. Gallager, *Information Theory and Reliable Communication*, New York: John Wiley & Sons, 1968.

[9] P.L. Heinzmann, Fiber optics and secure communications, IBM Research Report RZ 1759, Nov. 1988.

[10] S.K. Leung-Yan-Cheong, Multi-user and wire-tap channels including feedback, Tech. Rep. No. 6603-2, Stanford University, Information Systems Lab., July 1976.

[11] J.L. Massey, A simplified treatment of Wyner's wire-tap channel, *Proc. 21st Annual Allerton Conf. on Comm., Control, and Computing*, Monticello, IL, Oct. 5-7, 1983, pp. 268-276.

[12] U.M. Maurer, Conditionally-perfect secrecy and a provably-secure randomized cipher, to appear in *Journal of Cryptology.*

[13] U.M. Maurer and J.L. Massey, Cascade ciphers: the importance of being first, submitted to Journal of Cryptology.

[14] U.M. Maurer and J.L. Massey, Local randomness in pseudo-random sequences, to appear in *Journal of Cryptology.*

[15] A. Orlitsky, Worst-case interactive communication I: two messages are almost optimal, *IEEE Transactions on Information Theory*, Vol. 36, No. 5, 1990, pp. 1111-1126.

[16] C.E. Shannon, A mathematical theory of communication, *Bell System Technical Journal*, Vol. 27, No. 3, July 1948, pp. 379-423 and 623-656.

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

[18] G.S. Vernam, Cipher printing telegraph systems for secret wire and radio telegraphic communications, *J. Amer. Inst. Elec. Eng.*, Vol. 55, 1926, pp. 109-115.

[19] 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.

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