Appeared in *Communication and Cryptography – Two Sides of One Tapestry*, (R. Blahut et al., Eds.), Kluwer Academic Publishers, pp. 271-285, 1994.

The Strong Secret Key Rate of Discrete Random Triples

Ueli M. Maurer

Institute for Theoretical Computer Science
ETH Zürich
CH-8092 Zürich, Switzerland

Dedicated to James L. Massey on the occasion of his 60th birthday.

Abstract. Three parties, Alice, Bob and Eve, know the sequences of random variables $X^N = [X_1, X_2, \ldots X_N]$, $Y^N = [Y_1, Y_2, \ldots Y^N]$ and $Z^N = [Z_1, Z_2, \ldots Z_N]$, respectively, where the triples $(X_i Y_i Z_i)$, for $1 \leq i \leq N$, are generated by a discrete memoryless source according to some probability distribution P_{XYZ}. Motivated by Wyner's and Csiszár and Körner's pioneering definition of, and work on, the secrecy capacity of a broadcast channel, the secret key rate of P_{XYZ} was defined by Maurer as the maximal rate M/N at which Alice and Bob can generate secret shared random key bits $S_1, \ldots, S_M$ by exchanging messages over an insecure public channel accessible to Eve, such that the rate at which Eve obtains information about the key is arbitrarily small, i.e., such that $\lim_{N\to\infty} I(S_1, \ldots, S_M; Z^N, C^t)/N = 0$, where C^t is the collection of messages exchanged between Alice and Bob over the public channel. However, this definition is not completely satisfactory because only the rate, but not the total amount of information about the key obtained by Eve is bounded. This paper introduces and investigates the *strong* secret key rate: it is required that the total amount of information about the key obtained by Eve be negligible, i.e. $\lim_{N\to\infty} I(S_1, \ldots, S_M; Z^N, C^t) = 0$, and that $[S_1, \ldots, S_M]$ be arbitrarily close to uniformly distributed, i.e. $\lim_{N\to\infty} M - H([S_1, \ldots, S_M]) = 0$. Using novel results on privacy amplification by Bennett, Brassard, Crépeau and Maurer we demonstrate that the known results for the secret key rate also hold for the stronger definition.

I. Introduction

Unlike a communications engineer who can prove the quality of a designed communication system for a given noisy channel simply by demonstrating the error-free transmission of information at a specified rate, a cryptographer is usually in a much less comfortable position. He (or she) can usually only affirm that the state-of-the-art in cryptography and in cryptanalysis has been taken into account in the design of a system, but is not able to *prove* the security of the system. It is conceivable that a cipher gets broken shortly after it was designed and, to make things even worse, it is possible that such a failure will not even become known to the designer or users of a system. No presently-used ciphers (except the one-time pad that is used in rare applications where security is paramount), including public-key cryptosystems, can be proven secure.

In his research, Jim Massey has always attacked the fundamental question behind a given problem and refrained from going for the more promising, but less exciting goal of making minor contributions along a path other researchers had previously taken. On our trip to Eurocrypt '86 held in Linköping (shortly after my entrance into Sweden had caused severe complications because my passport was expired, and it was only due to Lis Massey's diplomatic intervention that finally an exception was made), Jim explained to me the strong need for rigorous proofs in cryptography and asked me whether I would like to accept the challenge of working towards provable security in cryptography as the topic of my doctoral research. This challenge struck me immediately and has never since ceased to drive my research. As a doctoral student I had the invaluable opportunity to work with Jim on various aspects of provable security (cf. [12], [15], [16]). I am deeply grateful for his careful guidance and for demonstrating, as an outstanding example, how rewarding and enjoyable it can be to work in an academic environment.

This paper is concerned with *provable* security in cryptography. More precisely, we try to beat Shannon's bound [18] for perfect secrecy which states that a cipher can only be perfect, i.e., plaintext and ciphertext can only be statistically independent, if the entropy of the plaintext is at most equal to the entropy of the secret key. Shannon's bound applies only when an opponent can, except for the secret key, see precisely the same information as the legitimate receiver. This assumption is overly pessimistic in many situations and we therefore consider a scenario in which two parties, called Alice and Bob, exploit knowledge of some correlated random variables about which an opponent Eve also has partial information. Alice and Bob, who share no secret key initially, can generate a secret key by communicating only over an insecure channel, even when Eve has more information than Bob about Alice's random variable and also more information than Alice about Bob's random variable. Eve has essentially no information about the finally shared secret key which can thus be used as the key in a one-time pad system [19] to transmit messages in perfect secrecy.

II. Secret key agreement by public discussion

In this section we describe the general scenario investigated in this paper, which was first suggested in [13] and independently in [1]. The purpose of this paper is to derive more powerful results for the same scenario.

Consider the following general key agreement problem. Assume that Alice, Bob and Eve know random variables X, Y and Z, respectively[1], with joint probability distribution P_{XYZ}, and that Eve has no information about X and Y other than through her knowledge of Z. More precisely, $I(XY;U|Z) = 0$ where U summarizes Eve's complete information about the universe. X, Y and Z take on values in some finite alphabets $\mathcal{X}, \mathcal{Y}$ and $\mathcal{Z}$, respectively.

Alice and Bob share no secret key initially, but are assumed to know P_{XYZ} or at least an upper bound on the correlation between Z and X and Y. Eve is assumed to know everything about the protocol used by Alice and Bob. Every message communicated between Alice and Bob over an insecure channel can be intercepted by Eve, but it is assumed that Eve cannot insert fraudulent messages nor modify messages on this public channel without being detected. In a scenario where Eve is not restricted to passive eavesdropping, an unconditionally-secure authentication scheme with a short initially shared secret key [20] can be used to detect active tampering with messages with very high probability. In this case, our protocols can be viewed as a method for expanding a short secret key (retaining perfect secrecy) rather than generating a key from scratch. If only a computationally-secure authentication scheme were used, the unconditional security would only be retained against passive, but not against active wire-tapping.

A realistic scenario for the generation of random variables X, Y and Z is by using a satellite broadcasting random bits to the earth at a very low signal power. Alice, Bob and Eve can receive the bits over partially independent channels with certain bit error probabilities ϵ_A, ϵ_B and ϵ_E, respectively. Eve's channel must be assumed to be imperfect (i.e., $\epsilon_E > 0$) for Alice and Bob to be able to generate a secret key, but as demonstrated in Section V of [13], it is neither required that $\epsilon_E > \epsilon_A$ nor that $\epsilon_E > \epsilon_B$. In fact, the capacity of Eve's channel can be allowed to be significantly (e.g., a thousand times) larger than the capacities of Alice's and Bob's channel. In this paper we are not concerned with particular such situations, but we rather investigate the rate at which Alice and Bob can generate secret key in a scenario where either $I(X;Z) > I(X;Y)$ or $I(Y;Z) > I(Y;X)$. In fact, we demonstrate that any such difference can be fully exploited.

Alice and Bob use a protocol in which at each step either Alice sends a message to Bob depending on X and all the messages previously received from Bob, or vice versa (with X replaced by Y). Without loss of generality, we consider only protocols in which Alice sends messages at odd steps ($C_1, C_3, \ldots$) and Bob sends messages at even steps

[1] Sequences of random variable as described in the abstract will be considered later.

$(C_2, C_4, \ldots)$. Moreover, we can restrict the analysis to deterministic protocols since a possible randomizer which Alice's and/or Bob's strategy and messages might depend on can be considered as part of X and Y, respectively. In other words, Alice and Bob can without loss of generality extend their known random variables X and Y, respectively, by random bits that are statistically independent of X, Y and Z. At the end of the t-step protocol, Alice computes a key S as a function of X and $C^t \triangleq [C_1, \ldots, C_t]$ and Bob computes a key S' as a function of Y and C^t. Their goal is to maximize $H(S)$ under the conditions that S and S' agree with very high probability and that Eve has very little information about either S or S'. More formally we have

$$H(C_i|C^{i-1}X) = 0 \tag{1}$$

for odd i,

$$H(C_i|C^{i-1}Y) = 0 \tag{2}$$

for even i,

$$H(S|C^tX) = 0 \tag{3}$$

and

$$H(S'|C^tY) = 0, \tag{4}$$

and it is required that

$$P[S \neq S'] \leq \epsilon \tag{5}$$

and

$$I(S; C^tZ) \leq \delta \tag{6}$$

for some specified (small) δ and ϵ.

If one requires that $P[S \neq S'] = 0$ and $I(S; C^t) = 0$ (i.e., that $\epsilon = 0$ in (5) and $\delta = 0$ in (6)) it appears intuitive but not obvious that $I(X; Y)$ is an upper bound on $H(S)$. It appears to be similarly intuitive that $H(S) \leq I(X; Y|Z) = I(XZ; YZ) - H(Z)$ because even under the assumption that Alice and Bob could learn Z, the remaining information shared by Alice and Bob is an upper bound on the information they can share in secrecy. It was proved in [13] that for every key agreement protocol satisfying (1)-(4),

$$H(S) \ \leq \ \min[I(X; Y), I(X; Y|Z)] + H(S|S') + I(S; C^tZ), \tag{7}$$

and hence, by Fano's lemma (cf. [4], p. 156) and conditions (5) and (6), that

$$H(S) \ \leq \ \min[I(X; Y), I(X; Y|Z)] + \delta + h(\epsilon) + \epsilon \log_2(|\mathcal{S}| - 1),$$

where $|\mathcal{S}|$ denotes the number of distinct values that S takes on with nonzero probability. It is worth pointing out that $I(X; Y) < I(X; Y|Z)$ is possible.

It is more interesting to derive lower rather than upper bounds on $H(S)$. In order to be able to prove lower bounds on the achievable size of a key S shared by Alice and Bob in secrecy, we need to make more specific assumptions about the distribution

P_{XYZ}. One natural model is that of a discrete memoryless source generating triples $(X_iY_iZ_i)$ independently for $i = 1, 2, \ldots, N$ according to some distribution P_{XYZ}. In other words, Alice, Bob and Eve receive $X^N = [X_1, \ldots, X_N]$, $Y^N = [Y_1, \ldots, Y_N]$ and $Z^N = [Z_1, \ldots, Z_N]$, respectively, where

$$P_{X^N Y^N Z^N} \;=\; \prod_{i=1}^{N} P_{X_i Y_i Z_i}$$

and where $P_{X_iY_iZ_i} = P_{XYZ}$ for $1 \leq i \leq N$.

It is common practice in information theory to state results about a scenario of independent repetitions of a random experiment in terms of information *rates*. The definition of secrecy capacity of a broadcast channel introduced by Wyner [21][2], and later generalized by Csiszár and Körner [7], is natural in this sense. The secrecy capacity of a broadcast channel specified by the conditional distribution $P_{YZ|X}$ is defined as the maximal rate at which Alice (controlling the X-input of the channel) can send information to Bob (receiving the Y-output) such that the rate at which Eve (receiving the Z-output) obtains this secret information is arbitrarily small.

In cryptography it is usually assumed that the availability of secure channels such as a trusted courier is restricted but that insecure channels are freely available. Therefore the following generalized definition of secrecy capacity introduced in [13], which allows arbitrary communication between Alice and Bob over an insecure channel, appears to be natural.

Definition 1. The *secret key rate of X and Y with respect to Z*, denoted $S(X;Y||Z)$, is the maximum rate at which Alice and Bob can agree on a secret key S while keeping the rate at which Eve obtains information arbitrarily small, i.e., it is the maximal R such that for every $\epsilon > 0$ there exists a protocol for sufficiently large N satisfying (1)-(5) with X and Y replaced by X^N and Y^N, respectively, further satisfying

$$\frac{1}{N} I(S; C^t Z^N) \;\leq\; \epsilon, \tag{8}$$

and achieving

$$\frac{1}{N} H(S) \;\geq\; R - \epsilon.$$

Remark: If for some protocol the secret key generated by Alice and Bob were not uniformly distributed, an almost uniformly distributed key could be generated by applying the protocol a sufficient number of times and using an ideal data compression scheme. Hence the condition

$$\frac{1}{N} H(S) \;>\; \frac{1}{N} \log_2 |\mathcal{S}| - \epsilon \tag{9}$$

could be included in the above definition without loss of generality.

[2]We refer to [10] for a simplified treatment of the wire-tap channel.

Like Wyner's and Csiszár and Körner's definition, this definition is not completely satisfactory both from a theoretical and a practical viewpoint. Since the results are asymptotic, it is possible for Eve to obtain a non-negligible amount of information about the secret key S, even if the rate at which she receives information is arbitrarily small. In fact, according to the definition, her information is allowed to grow without bound as N goes to infinity, as long as the growth is less than linear in N. The confidentiality of a small part of a plaintext message could be of paramount importance and it is not guaranteed that this particular part is protected in a one-time pad that uses a generated secret key.

The purpose of this paper is to show that privacy amplification [3], [2] allows Alice and Bob to generate a secret key S, even when it is required that Eve's *total* information about S be negligibly small. Furthermore we require a uniformity condition on S that is much stricter than (9). We therefore introduce the following definition, where $|\mathcal{S}|$ denotes the cardinality of the set $\mathcal{S}$ of keys.

Definition 2. The *strong secret key rate of X and Y with respect to Z*, denoted $\overline{S}(X;Y||Z)$, is defined in the same way as the secret key rate in Definition 1, with the two modifications that condition (8) is replaced by

$$I(S;C^tZ^N) \ \leq \ \epsilon$$

and that

$$H(S) \geq \log_2|\mathcal{S}| - \epsilon.$$

We obviously have

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

where the second inequality is an immediate consequence of (7). One of the results of [13] states that if either Eve has less information about Y than Alice or, by symmetry, Eve has less information about X than Bob, then such a difference of information can be exploited:

$$\begin{aligned} S(X;Y||Z) &\geq \max[I(Y;X) - I(Z;X),\ I(X;Y) - I(Z;Y)] \\ &= I(X;Y) - \min[I(Z;X), I(Z;Y)]. \end{aligned} \tag{10}$$

The main result of this paper is a proof that the same lower bound holds also for the strong secret key rate.

III. Reconciliation and Privacy amplification

One particular protocol that allows Alice and Bob to generate a secret key consists of the following two phases. It should be pointed out that this type of protocol only allows to prove our main result, namely that the lower bound (10) also holds for the strong secret key rate, but that in situations where the right-hand side of (10) vanishes, more complicated protocols must be used to generate a secret key.

In a first phase, Alice sends $h(X^N)$ to Bob, where $h : \mathcal{X}^N \to \{0,1\}^L$ is a function designed to provide Bob (who knows Y^N) with a sufficient amount of redundant information about X^N to allow him to reconstruct X^N with high probability. The existence of such a function for L on the order of $N \cdot H(X|Y)$ is stated in the following theorem which implies that Bob can be informed by Alice about her string by sending bits (over a perfect channel) at a rate arbitrarily close to $H(X|Y)$. The proof of the theorem is omitted but will be given in a subsequent paper [14] which will provide a more general treatment of strong secret key rate and secrecy capacity.

Theorem 1: *Let the sequence $[(X_1, Y_1), \ldots, (X_N, Y_N)]$ be generated as described above. For every $\epsilon > 0$ and $\epsilon' > 0$, for sufficiently large N and for every L satisfying $L/N > (1+\epsilon)H(X|Y)$, there exists a function $h : \mathcal{X}^N \to \{0,1\}^L$ such that $[X_1, \ldots, X_N]$ can be decoded from $[Y_1, \ldots, Y_N]$ and $h(X^N)$ with error probability at most ϵ'.*

In a second phase, called privacy amplification, Alice and Bob compress the now shared string X^N, in a manner known to Eve, to result in a shorter binary string $S = [S_1, \ldots, S_M]$ with virtually uniform distribution about which Eve has essentially no information. Of course, this privacy amplification step must take into account Eve's total information about X^N consisting of Z^N and $h(X^N)$.

Privacy amplification was introduced in [3] and generalized in [2] and can be described as follows. Assume Alice and Bob share an N–bit string w about which an eavesdropper Eve has incomplete information characterized by a probability distribution P over the N–bit strings. For instance, Eve might have received some bits or parities of bits of w, she might have eavesdropped on some of the bits of w through a binary symmetric channel, or have some more complicated type of information about w. Alice and Bob have some knowledge of this distribution P, but they do not know exactly what is compromised about the secrecy of their string. Using a public channel, which is totally susceptible to eavesdropping, they wish to agree on a function $g : \{0,1\}^N \to \{0,1\}^M$ such that Eve, despite her partial knowledge about w and complete knowledge of g, almost certainly knows nearly nothing about $g(w)$. This process transforms a partially secret N–bit string w into a highly secret but shorter M–bit string $g(w)$.

Bennett, Brassard and Robert [3] solved the problem for the case where Eve is allowed to specify (secretly) an arbitrary eavesdropping function $e : \{0,1\}^N \to \{0,1\}^T$ from N bits to T bits such that only T, but not the function e is known to Alice and Bob, and where Eve obtains the result $e(w)$ of applying the eavesdropping function to w. Equivalently, Eve could be allowed to perform an arbitrary computation with w as input, as long as she keeps only T bits of the result and discards the input and all the intermediate results. The solution of [3] consists of Alice randomly selecting a function from a universal class of hash functions (see definition below) mapping N-bit strings to T-bit strings for an appropriate choice of T, and sending the description (or index) of the selected function to Bob (and hence also to Eve) over the insecure channel.

Definition 3 [6]: A class G of functions $\mathcal{A} \longrightarrow \mathcal{B}$ is *universal*$_2$ ("universal" for short) if, for any distinct x_1 and x_2 in $\mathcal{A}$, the probability that $g(x_1) = g(x_2)$ is at most $1/|\mathcal{B}|$ when g is chosen at random from G according to the uniform distribution.

Example: Let a be an elements of $GF(2^N)$ and also interpret x as an element of $GF(2^N)$. Consider the function $\{0,1\}^N \to \{0,1\}^M$ assigning to an argument x the first M bits of the element ax of $GF(2^N)$. The class of such functions for $a \in GF(2^N)$ with $a \neq 0$ is a universal class of functions for $1 \leq M \leq N$.

The results of [3] were generalized by Bennett, Brassard, Crépeau and Maurer [2] to include scenarios in which Eve's information about w is specified by some general probability distribution satisfying a certain constraint in terms of collision entropy defined below.

Definition 4 [2]: Let P_W be a probability distribution over some sample space $\mathcal{W}$. (Equivalently, we can consider the random variable W distributed according to P_W.) The *collision probability* of W, denoted $P_c(W)$, is the probability of drawing the same element if one samples twice in $\mathcal{W}$, with replacement, according to probability distribution P_W:

$$P_c(W) = \sum_{w \in \mathcal{W}} (P_W(w))^2.$$

The *collision entropy* [3] of W, denoted $H_c(W)$ is the negative logarithm of the collision probability, i.e.,

$$H_c(W) = -\log P_c(W). \quad ^4$$

It follows immediately from Jensen's inequality that

$$H(W) \geq H_c(W), \tag{11}$$

with equality if and only if P_W is the uniform distribution over $\mathcal{W}$ or a subset of $\mathcal{W}$, and where $H(W)$ is the (Shannon) entropy of a random variable W distributed according to P_W.

In analogy to Shannon entropy, one can also define conditional collision entropy. For an event $\mathcal{E}$, $H_c(W|\mathcal{E})$ is naturally defined as the collision entropy of the conditional distribution $P_{W|\mathcal{E}}$, for instance

$$H_c(W|V = v) = -\log \sum_{w \in \mathcal{W}} (P_{W|V}(w, v))^2,$$

and the collision entropy conditioned on a random variable can be defined as the expected value of the conditional collision entropy:

$$H_c(W|V) = \sum_v P_V(v) H_c(W|V = v).$$

[3]also known as Renyi entropy of order 2

[4]All logarithms in this paper are to the base 2.

One can also define collision information in analogy to Shannon information.

Like Shannon entropy, collision entropy conditioned on a random variable is "well-behaved": it is proved in [5] that

$$H_c(W) - H_c(W|V) \le H(V).$$

It should be pointed out, however, that the more intuitive inequality $H_c(W) - H_c(W|V) \le H_c(V)$ is false in general [5]. However, an important problem we will have to deal with is that, like for Shannon entropy, the condition $V = v$ can induce an arbitrarily large decrease of collision entropy: If V can take on 2^L values, $H_c(W) - H_c(W|V = v) \gg L$ is possible for certain values v.

We will make use in a crucial manner of an interesting and counter-intuitive property of collision entropy pointed out and used in [2]. As opposed to Shannon entropy, collision entropy can *increase* when extra information is revealed, i.e., $H_c(W|V) > H_c(W)$ is possible. (Of course, this property rules out collision entropy as a measure of information that could be useful in investigating source and channel coding.)

We now return to the discussion of privacy amplification. One of the main results of [2] can be restated as follows.

Theorem 2 [2]: *Let P_W be a probability distribution over $\mathcal{W}$ with collision entropy $H_c(W)$, and let G be the random variable corresponding to a universal class of hash functions from $\mathcal{W}$ to $\{0,1\}^M$ with uniform distribution over the class. Then*

$$H(G(W)|G) \ge H_c(G(W)|G) \ge M - \frac{2^{M - H_c(W)}}{\ln 2}.$$

Remark: While this theorem applies of course also to conditional probability distributions, i.e.,

$$H_c(G(W)|G, V = v) \ge M - 2^{M - H_c(W|V=v)}/\ln 2,$$

it should be pointed out that it cannot be generalized to collision entropy conditioned on a random variable: $H_c(G(W)|GV) \ge M - 2^{M - H_c(W|V)}/\ln 2$ is false in general.

Theorem 2 states that if Alice and Bob share a particular string w and Eve's information about w can be modeled by the distribution $P_{W|V=v}$ (where v denotes the particular value of her information vector) about which Alice and Bob know nothing except a lower bound T on the collision entropy, i.e. $H_c(W|V = v) \ge T$, then Alice and Bob can generate a secret key of roughly T bits. More precisely, if Alice and Bob compress w slightly more to an M-bit key with $M < T$, then Eve's total information about this key decreases exponentially in the excess compression $T - M$.

IV. A lower bound on strong secret key rate

Our goal is to apply privacy amplification to the string X^N shared by Alice and Bob after the error-correction phase in which $h(X^N)$ is sent from Alice to Bob, taking into account Eve's knowledge consisting of Z^N and $h(X^N)$. However, several major problems arise:

- First, Eve's initial collision entropy $H_c(X^N|Z^N = z^N)$ depends on the particular string z^N that she has received. Unfortunately, as pointed out above, privacy amplification does not apply when only a bound on the *average* collision entropy $H_c(X^N|Z^N)$ is known.

- Second, the reduction of Eve's collision entropy about $X_1, \ldots, X_N$ due to receiving a particular value of the error-correction information, $h(X^N) = a$, sent from Alice to Bob over the public channel, must be analyzed. Knowing that $H_c(X^N|Z^N) - H_c(X^N|Z^N, h(X^N)) \leq H(h(X^N)) \leq L$ is not sufficient for the same reason as mentioned above. Because $H_c(X^N|Z^N = z^N, h(X^N) = a)$ could potentially be much smaller than $H_c(X^N|Z^N = z^N)$ one has to consider all possible values of $h(X^N)$.

- Third, Theorem 2 suggests that $H_c(X^N|Z^N=z^N, h(X^N) = a)$ is an upper bound on the size of the secret key that can be generated by privacy amplification. Unfortunately, the collision entropy is generally smaller than the Shannon entropy. In particular, $H_c(X^N|Z^N)$ will generally be substantially smaller than $H(X^N|Z^N)$; hence it appears impossible to exploit Eve's full Shannon entropy $H(X^N|Z^N)$, reduced by the amount of extra information (on the order of $H(X|Y)$) provided by $h(X^N)$, as would be necessary in order to prove that the lower bound (10) also holds for the strong secret key rate.

- Fourth, one needs to guarantee that the finally shared string $S = [S_1, \ldots, S_M]$ has virtually maximal entropy.

We solve these problems by exploiting the fact described earlier that collision entropy can increase when extra information is revealed. It is therefore conceivable to consider an oracle who gives Eve some side information (called spoiling knowledge in [2]) about X^N for free. Revealing extra information can certainly not hurt Eve since she could always discard it. However, if chosen carefully, this extra information may increase Eve's collision entropy. This demonstrates that a longer key than suggested by considering Eve's collision entropy about X^N (without the oracle's "help") can safely be distilled by application of Theorem 2. Clearly, Eve's Shannon entropy will be reduced by receiving the oracle's side information, but in our case this reduction will be negligible in terms of rate, i.e., when divided by N.

In the following we will make use of typical sequence arguments. There exist several definitions of typical sequences, and we use that of [4] for strongly typical sequences. Consider a probability distribution P_U over some finite set $\mathcal{U}$, which we assume without loss of generality to be $\mathcal{U} = \{1, \ldots, t\}$ for some t. We further assume that $P_U(i) > 0$ for $1 \leq i \leq t$. Consider a sequence u^N of N digits of $\mathcal{U}$ and define $n_i(u^N)$, for $i = 1, \ldots, t$, to be the number of occurrences of the digit i in u^N. A sequence u^N is called a δ-typical sequence if and only if

$$(1-\delta)P_U(i) \leq \frac{n_i(u^N)}{N} \leq (1+\delta)P_U(i)$$

for $1 \leq i \leq t$. Consider now a sequence $U^N = [U_1, \ldots, U_N]$ of N i.i.d. random variables, each distributed according to P_U. Using the Chernoff bound (cf. [4]) one can prove that the total probability of all δ-typical sequences approaches 1 as N goes to infinity. More precisely, the total probability of the non-δ-typical sequences goes to 0 faster than $1/N$: For every $\delta > 0$ and $\epsilon > 0$, we have

$$N \cdot P[U^N \text{ is not } \delta - \text{typical}] < \epsilon \tag{12}$$

for sufficiently large N.

We now return to the discussion of our secret key agreement scenario with independent random triples $(X_iY_iZ_i)$, for $i = 1, \ldots, N$, being generated according to P_{XYZ}. We first focus on the sequence of pairs (X_iZ_i). Without loss of generality we let the alphabets for X and Z be $\mathcal{X} = \{1, \ldots, t_1\}$ and $\mathcal{Z} = \{1, \ldots, t_2\}$, respectively. Let m_j for $1 \leq j \leq t_2$ denote the number of occurrences of digit j in the sequence $Z_1, \ldots, Z_N$, and let n_{ij} for $1 \leq i \leq t_1$ and $1 \leq j \leq t_2$ denote the number of occurrences of the pair (i, j) in the sequence $[(X_1, Z_1), \ldots, (X_N, Z_N)]$. We have

$$\sum_{i=1}^{t_1} n_{ij} = m_j \tag{13}$$

for $1 \leq i \leq t_1$, and

$$\sum_{j=1}^{t_2} m_j = N. \tag{14}$$

Let $\mathcal{E}$ be the event that the sequence $[(X_1, Z_1), \ldots, (X_N, Z_N)]$ is δ-typical for the alphabet $\mathcal{X} \times \mathcal{Z}$ and the distribution P_{XZ}. According to (12), $P[\bar{\mathcal{E}}]$ can be made arbitrarily small for any fixed $\delta > 0$ by choosing a sufficiently large block length N. In the following we will consider probability distributions and entropies conditioned on the event $\mathcal{E}$. By definition, this condition implies that

$$(1-\delta)P_{XZ}(i,j) \leq \frac{n_{ij}}{N} \leq (1+\delta)P_{XZ}(i,j) \tag{15}$$

and, as a consequence, that

$$(1-\delta)P_Z(j) \leq \frac{m_j}{N} \leq (1+\delta)P_Z(j). \tag{16}$$

Eve knows a particular sequence z^N with corresponding values $m_1, \ldots, m_{t_2}$. Assume now that the oracle mentioned above tells Eve the numbers n_{ij} for free. This extra information, denoted as O, decreases Eve's Shannon entropy somewhat, i.e.,

$$H(X^N|Z^N = z^N, O) < H(X^N|Z^N = z^N)$$

and

$$H(X^N|Z^N = z^N, O, \mathcal{E}) < H(X^N|Z^N = z^N, \mathcal{E})$$

but increases her collision entropy significantly, i.e.,

$$H_c(X^N|Z^N = z^N, O, \mathcal{E}) > H_c(X^N|Z^N = z^N, \mathcal{E}).$$

In fact, for a particular value $O = o$ provided by the oracle, Eve's distribution of X^N, i.e. $P_{X^N|Z^N=z^N,O=o,\mathcal{E}}$, is such that all sequences $[x_1, \ldots, x_N]$ that are consistent with her information are equally probable. It is easy to see that the number of such sequences is

$$Q = \prod_{j=1}^{t_2} \frac{m_j!}{\prod_{i=1}^{t_1} n_{ij}!}.$$

Therefore both the Shannon and the collision entropy of this distribution are equal to $\log Q$.

As pointed out before, privacy amplification applies to conditional distributions only when a bound on the collision entropy of a random variable, given the particular value of the conditioning random variable, is known. In order to be able to apply privacy amplification to the distribution $P_{X^N|Z^N=z^N,O=o,\mathcal{E}}$, we state the following result for specific values $Z^N = z^N$ and $O = o$. Of course, it also hold when averaged over all values of Z^N and O.

Lemma 3. *For $0 < \delta \leq 1/2$ and for all values z^N and o,*

$$H_c(X^N|Z^N = z^N, O = o, \mathcal{E}) > N\left[H(X|Z) - \delta(H(X) + H(XZ) + 4)\right] - t_1 t_2 \log N.$$

This lemma implies that for sufficiently small δ and for sufficiently large N, Eve's per-digit Shannon and collision entropy are both arbitrarily close to $H(X|Z)$. Note again that $H_c(X^N|Z^N O, \mathcal{E})$ is significantly larger than $H_c(X^N|Z^N, \mathcal{E})$.

Proof: Stirling's formula for $n!$ (cf. [9], p. 467) implies that

$$n(\log n - \alpha) \ < \ \log n! \ < \ n(\log n - \alpha) + \log n$$

for all sufficiently large n, where $\alpha = 1/\ln 2$ and $\ln 2$ denotes the logarithm of 2 to the base e. Using (13), (14), (15) and (16) we get

$$\begin{aligned}
\log Q &= \sum_{j=1}^{t_2} \left(\log(m_j!) - \sum_{i=1}^{t_1} \log(n_{ij}!) \right) \\
&> \sum_{j=1}^{t_2} m_j(\log m_j - \alpha) - \sum_{i=1}^{t_1} \sum_{j=1}^{t_2} \left[n_{ij}(\log n_{ij} - \alpha) + \log n_{ij}\right] \\
&> \sum_{j=1}^{t_2} m_j(\log m_j - \log N) - \sum_{i=1}^{t_1} \sum_{j=1}^{t_2} \left[n_{ij}(\log n_{ij} - \log N) + \log N\right]
\end{aligned}$$

$$
\begin{aligned}
&= N\left[\sum_{j=1}^{t_2}\frac{m_j}{N}\log\frac{m_j}{N} - \sum_{i=1}^{t_1}\sum_{j=1}^{t_2}\frac{n_{ij}}{N}\log\frac{n_{ij}}{N}\right] - t_1t_2\log N \\
&\geq N\left[\sum_{j=1}^{t_2}(1+\delta)P_Z(j)\log((1-\delta)P_Z(j))\right. \\
&\qquad \left. - \sum_{i=1}^{t_1}\sum_{j=1}^{t_2}(1-\delta)P_{XZ}(i,j)\log((1+\delta)P_{XZ}(i,j))\right] - t_1t_2\log N \\
&\geq N\left[(1+\delta)\sum_{j=1}^{t_2}P_Z(j)\log P_Z(j) - (1-\delta)\sum_{i=1}^{t_1}\sum_{j=1}^{t_2}P_{XZ}(i,j)\log P_{XZ}(i,j)\right] \\
&\qquad + N\Big[(1+\delta)\log(1-\delta) - (1-\delta)\log(1+\delta)\Big] - t_1t_2\log N
\end{aligned}
$$

Note that because $\sum_{j=1}^{t_2} m_j = \sum_{i=1}^{t_1}\sum_{j=1}^{t_2} n_{ij}$, the two occurrences of α on the second line can both be deleted or replaced by any other expression (like $\log N$ in our case). We have also made use of the trivial fact that $\log n_{ij} \leq \log N$. It is easy to check that $(1+\delta)\log(1-\delta) - (1-\delta)\log(1+\delta) > -4\delta$ for all $\delta \leq 1/2$. Hence

$$
\begin{aligned}
H_c(X^N|Z^N = z^N, O = o, \mathcal{E}) &= \log Q \\
&> N\left[-(1+\delta)H(Z) + (1-\delta)H(XZ) - 4\delta\right] - t_1t_2\log N \\
&= N\left[H(X|Z) - \delta(H(X) + H(XZ) + 4)\right] - t_1t_2\log N
\end{aligned}
$$

for all $\delta \leq 1/2$, as was to be shown. $\Box$

In order to apply privacy amplification according to Theorem 2 to compress the string X^N now shared by Alice and Bob to a shorter string S about which Eve has essentially no information, it remains to investigate the reduction of Eve's collision entropy about X^N due to seeing $h(X^N)$ sent from Alice to Bob over the public channel. As pointed out before, for any particular value a taken on by $h(X^N)$, the reduction of collision entropy induced by obtaining side-information $h(X^N) = a$, i.e. $H_c(X^N) - H_c(X^N|h(X^N) = a)$, could generally be arbitrarily large. (The fact that $H_c(X^N) - H_c(X^N|h(X^N)) \leq H(h(X^N))$ is of little use here.) However, it is easy to prove (cf. [2],[5]) that for a uniform distribution, i.e., one for which all non-zero probabilities are identical, revealing L bits of information can reduce the collision entropy by at most L. Thus

$$
H_c(X^N|Z^N = z^N, O = o, h(X^N) = a, \mathcal{E}) \;\geq\; H_c(X^N|Z^N = z^N, O = o, \mathcal{E}) - L. \quad (17)
$$

The main result of this paper can be summarized in the following theorem. Only a sketch of the proof is given and we refer to [14] for a complete proof.

Theorem 4: $\overline{S}(X;Y||Z) \;\geq\; \max[I(Y;X) - I(Z;X),\; I(X;Y) - I(Z;Y)]$.

Proof sketch: We only prove that $I(Y;X) - I(Z;X)$ is an achievable rate; the proof for $I(X;Y) - I(Z;Y)$ follows by symmetry. Alice and Bob choose a suitable error-correction function $h : \mathcal{X}^N \to \{0,1\}^L$ and, after having sent and received $h(X^N)$, choose a compression function G at random from a universal class of hash functions $\mathcal{X}^N \to \{0,1\}^M$, for appropriate parameters L and M. Then they compute $S = G(X^N)$. The quantity to be bounded is $I(S;GZ^N h(X^N))$. It can be shown to be arbitrarily small by proving that $H(S|GZ^N h(X^N))$ is arbitrarily close to M.

$$
\begin{aligned}
H(S|GZ^N h(X^N)) &\geq (1 - P[\mathcal{E}]) H(S|GZ^N h(X^N), \mathcal{E}) \\
&\geq H(S|GZ^N h(X^N), \mathcal{E}) - P[\mathcal{E}] \cdot H(S) \\
&\geq H(S|GZ^N h(X^N) O, \mathcal{E}) - P[\mathcal{E}] \cdot N \cdot H(X).
\end{aligned}
$$

Theorem 2 implies that if $\mathcal{E}$ occurs, then for every value $[z^N, o, a]$ which the random triple $[Z^N, O, h(X^N)]$ can take on,

$$
H(S|G, Z^N = z^N, O = o, h(X^N) = a, \mathcal{E}) \geq M - 2^{M - H_c(X^N|Z^N = z^N, O = o, h(X^N) = a, \mathcal{E})} / \ln 2.
$$

Let $\epsilon > 0$ be an arbitrary but fixed parameter. For an appropriate choice of L/N ϵ-close to $H(X|Y)$ and of M/N ϵ-close to $H(X|Z) - H(X|Y)$, and by using (17) and Lemma 3, one can show that the above exponent goes to minus infinity as N goes to infinity. Hence $H(S|G, Z^N = z^N, O = o, h(X^N) = a, \mathcal{E})$ and thus also $H(S|GZ^N h(X^N) O, \mathcal{E})$ can be made arbitrarily close to M for sufficiently large N. Furthermore, (12) implies that $N \cdot P[\mathcal{E}] \cdot H(X)$ vanishes when N goes to infinity, and Theorem 1 implies that Bob can decode X^N from Y^N and $h(X^N)$ with probability arbitrarily close to 1. $\square$

V. Conclusions

We have pointed out that previous definitions of secrecy capacity of broadcast channels and secret key rate of random triples are not satisfactory because the total amount of information an opponent can obtain is not bounded, let alone arbitrarily small. For a correspondingly stronger definition of secret key rate it was proved that the results previously obtained for a weak definition of secret key rate also hold for the new stronger definition. The techniques of [2] used in the proof appear to be novel and we believe that they will have other applications in information theory. Results for a strengthened definition (in analogy to Definition 2) of secrecy capacity in the broadcast channel models of Wyner [21] and Csiszár and Körner [7] will be described in [14].

The strong secret key rate is an asymptotic definition. However, concrete protocols based on techniques described in Section V of [13] and in [8] and on efficiently decodable error-correcting codes can be constructed and analyzed using the techniques of [5]. Perfectly-secure secret-key agreement is possible even when Eve initially has more information about Alice's string than Bob and also more information about Bob's string than Alice [13]. In this case, however, several rounds of interaction between Alice and Bob are required.

Acknowledgment

It is a pleasure to thank Jim Massey for providing the initial motivation for this research, and Charles H. Bennett, Gilles Brassard, Christian Cachin, Claude Crépeau and Martin Gander for interesting discussions.

References

[1] R. Ahlswede and I. Csiszàr, Common randomness in information theory and cryptography – Part I: secret sharing, *IEEE Transactions on Information Theory*, Vol. 39, No. 4, pp. 1121-1132, July 1993.

[2] C.H. Bennett, G. Brassard, C. Crépeau and U.M. Maurer, Privacy amplification against probabilistic information, preprint, 1993.

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

[4] R.E. Blahut, *Principles and Practice of Information Theory*, Reading, MA: Addison-Wesley, 1987.

[5] C. Cachin and U.M. Maurer, Linking information reconciliation and privacy amplification, preprint, 1994.

[6] J.L. Carter and M. N. Wegman, Universal classes of hash functions, *Journal of Computer and System Sciences*, Vol. 18, 1979, pp. 143–154.

[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] M. Gander and U.M. Maurer, On the secret-Key rate of binary random variables (extended abstract), to be presented at the 1994 Int. Symp. on Information Theory.

[9] R. L. Graham, D. E. Knuth and O. Patashnik, *Concrete mathematics*, Reading, MA: Addison-Wesley, 1990.

[10] 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.

[11] J.L. Massey, Contemporary cryptology – an Introduction, in *Contemporary cryptology – the science of information integrity*, G.J. Simmons (Ed.), IEEE Press, 1992.

[12] U.M. Maurer, *Provable security in cryptography*, Ph. D. dissertation, No. 9260, Swiss Federal Institute of Technology (ETH), Zurich, 1990.

[13] U.M. Maurer, Secret key agreement by public discussion from common information, *IEEE Transactions on Information Theory*, Vol. 39, No. 3, pp. 733-742, May 1993.

[14] U.M. Maurer, Strengthening the definition of secret key rate and secrecy capacity, in preparation.

[15] U.M. Maurer and J.L. Massey, Local randomness in pseudo-random sequences, *Journal of Cryptology*, Vol. 4, No. 2, 1991, pp. 135-149.

[16] U.M. Maurer and J.L. Massey, Cascade ciphers: the importance of being first, *Journal of Cryptology*, Vol. 6, No. 1, pp. 55-61, 1993.

[17] 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, pp. 120-126, 1978.

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

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

[20] 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, pp. 265-279, 1981.

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