Appeared in *Journal of Cryptology*, vol. 5, no. 1, pp. 53-66, 1992.

Conditionally-Perfect Secrecy and a Provably-Secure Randomized Cipher [1]

Ueli M. Maurer

Institute for Theoretical Computer Science
Swiss Federal Institute of Technology (ETH)
CH-8092 Zurich, Switzerland

Abstract. Shannon's pessimistic theorem, which states that a cipher can be perfect only when the entropy of the secret key is at least as great as that of the plaintext, is relativized by the demonstration of a randomized cipher in which the secret key is short but the plaintext can be very long. This cipher is shown to be "perfect with high probability". More precisely, the eavesdropper is unable to obtain any information about the plaintext when a certain security event occurs, and the probability of this event is shown to be arbitrarily close to one unless the eavesdropper performs an infeasible computation. This cipher exploits the assumed existence of a publicly-accessible string of random bits whose length is much greater than that of all the plaintext to be encrypted; this is a feature that our cipher has in common with the previously considered "book ciphers". Two modifications of this cipher are discussed that may lead to practical provably-secure ciphers based on either of two assumptions that appear to be novel in cryptography, viz., the (sole) assumption that the enemy's memory capacity (but not his computing power) is restricted and the assumption that an explicit function is, in a specified sense, controllably-difficult to compute, but not necessarily one-way.

Keywords: provable security, perfect secrecy, information theory, randomized encryption, public randomness, book cipher.

[1] A preliminary version of this paper was presented at EUROCRYPT'90, May 21-24, 1990, Århus, Denmark, and will appear in the proceedings.

1. Introduction

One of the most important practical and theoretical open problems in cryptography is to devise a cipher that is both provably-secure and practical. The significance of a result on provable security crucially depends on the definition of security used, on the assumptions about the enemy's knowledge and resources, and on the practicality of the cipher. Excluding approaches that are based on an unproven hypothesis such as the intractability of a certain problem (e.g., factoring), one observes that every approach to provable security that has previously been proposed (except maybe the recent approach by Maurer [5] briefly discussed below) is either impractical or is based on a generally unrealistic assumption about the enemy's *a priori* and/or obtainable knowledge. To list a few examples: the one-time pad [10] is, because of its large key size, impractical in most applications; perfect local randomizers [6] are based on the generally unrealistic assumption that an eavesdropper can only obtain a small number of ciphertext bits; Wyner's wire-tap channel [12] is based on the generally unrealistic assumption that the eavesdropper's channel is noisier than the main channel; and the Rip van Winkle cipher proposed by Massey and Ingemarsson [3, 4] is completely impractical since the legitimate receiver's deciphering delay is on the order of the square of the time the enemy must spend in order to break the cipher. It is conceivable that quantum cryptography introduced by Bennett, Brassard and others in a series of papers (see [1] and its references) can eventually become practical, but their proofs of security rely on the uncertainty principle of quantum physics. Finally, the result that a cascade of additive stream ciphers is at least as secure as any of its component ciphers [7] yields provably-secure ciphers only when a set of additive stream ciphers can be constructed that provably contains at least one computationally secure cipher (which may be impossible to identify).

The only practical and provably-secure cryptosystem that is not based on an unrealistic or unproven assumption is Maurer's recently proposed information-theoretically secure key exchange protocol [5]. It allows two parties A and B receiving the output of a random source (e.g., a satellite broadcasting random bits) over different noisy channels to agree (by public discussion) on a secret key Z such that an enemy is left with arbitrarily little information in Shannon's sense about Z. This key can subsequently be used to set up a perfect one-time pad. The protocol is secure even when the enemy is receiving the output of the same source over a much more reliable (less noisy) channel than A and B, and when the enemy can listen to the public discussion between A and B. The proof is based on the sole realistic assumption that the noise on the enemy's channel is at least to some degree independent of the noise on A's and B's channels.

In this paper, we present a new approach to provable security that was motivated by [4] and is based on the assumed availability of a very large publicly-accessible string of random bits. The need for this public randomizer is the only but at present serious detriment to the practicality of the proposed cipher. The randomizer could, for instance, be stored on a high-density storage medium, copies of which are publicly available, or it could be broadcasted by a satellite. Alternatively, a natural publicly-accessible source of randomness could be used.

The enemy's computational effort needed to break the cipher is measured in terms of the number of randomizer bits that he must examine or, equivalently, as the depth of the decision tree corresponding to his computation (see [11]). The basic idea of our approach is to prove that, even if he uses an optimal strategy for examining randomizer bits, an eavesdropper obtains no information in Shannon's sense about the plaintext with probability very close to one unless he accesses a substantial fraction of all the randomizer bits. More precisely, we prove that if a certain event occurs, then the eavesdropper's entire observation, consisting of the cryptogram and the examined randomizer bits, is statistically independent of the plaintext. The probability of this event is lower bounded by a quantity that depends only on the number of bits examined by the eavesdropper, and it is very close to one unless he examines a substantial fraction (e.g., 2/3) of the entire randomizer. It is obviously impossible to prove that the number of bits that the eavesdropper must examine is greater than the total number of randomizer bits, and thus our result is close to optimum within our framework of provable security.

Since the effort to examine a random bit is in current technology roughly equal to that required to generate one, our lower bound on the enemy's computational effort appears to be on the same order as the effort needed to generate the randomizer. Therefore, our strongly-randomized cipher is truly practical only when either a source of randomness is available from Nature (for example, a deep-space radio source or the surface of the moon) or should a much simpler way of generating large amounts of random data be discovered. It is not the purpose of this paper to discuss further the technical problem of generating a huge amount of publicly available random data. Rather, our interest is in exploring the question whether provable security is possible in such a model. However, as will be explained in Section 4, the broadcast version of our cipher requires a randomizer length that is only somewhat larger than the enemy's memory capacity and can therefore be practical. Unfortunately, the security of this cipher could so far be proved only under the unrealistic assumption that the enemy stores actual randomizer bits in his memory rather than cleverly chosen boolean functions of the randomizer bits.

The results of this paper appear to be somewhat surprising for two reasons. First, they demonstrate that, although perfect secrecy can be achieved only when the entropy of the secret key is at least equal to that of the plaintext (see [9]), relaxing the notion of perfectness only slightly allows one to build a provably-secure cipher whose secret key is very short compared to the length of the plaintext. Second, although information-theoretic security usually implies that the enemy has infinite computing power, our proposed cipher is secure for an information-theoretic notion of security only when the enemy is computationally restricted.

Previously proposed book ciphers share an essential feature with our cipher: the use of a long publicly-accessible random string (a book) and a short secret key that selects portions of the random string which are used for enciphering a message. The approach of this paper however differs from previously discussed book ciphers in that a novel way of accessing the public randomness is proposed that, under the assumption that the "book" is truly random, allows to prove the security of the cipher.

In Section 2, our model of a cipher with public randomizer is introduced, and a particular randomized cipher is presented. After describing a general model of attacks against randomized ciphers, a proof of the information-theoretic security of our cipher against all feasible attacks is given in Section 3. In Section 4, techniques are suggested for basing the (provable) security of ciphers on either one of two assumptions, viz., that the enemy's memory capacity is restricted or that a certain function is difficult to compute in a specified sense, but not necessarily one-way.

2. Description of the Randomized Cipher

Throughout this paper, random variables are denoted by capital letters, whereas the corresponding small letters denote specific values that can be taken on by these random variables. Underlined capital letters or superscripted capital letters denote random vectors; a superscript indicates the number of components. Our model of a strongly-randomized cipher is as follows. As in a conventional symmetric cryptosystem, the communicating parties share a short randomly-selected secret key. The randomizer $\underline{R}$ is a binary random string of length L, whose bits can be read in a random-access manner by the legitimate parties as well as by an eavesdropper, i.e., $\underline{R}$ is assumed to be publicly accessible. The cryptogram is a function of the plaintext, the secret key and the randomizer such that given the cryptogram, the key and the randomizer, the plaintext is uniquely determined. The goal of the design of a randomized cipher is to devise an encryption transformation such that the cryptogram depends on only a few randomizer bits whose positions in turn depend on the secret key in such a manner that without the secret key it is impossible to obtain any information about the plaintext without examining a very large number of randomizer bits.

We now describe our specific strongly-randomized cipher. It is a binary additive stream cipher in which the plaintext $\underline{X} = [X_1, \ldots, X_N]$, the cryptogram $\underline{Y} = [Y_1, \ldots, Y_N]$ and the keystream $\underline{W} = [W_1, \ldots, W_N]$ are binary sequences of length N. The cryptogram $\underline{Y}$ is obtained by adding $\underline{X}$ and $\underline{W}$ bitwise modulo 2:

$$Y_n = X_n \oplus W_n \quad \text{for } 1 \leq n \leq N.$$

The publicly-accessible binary random string $\underline{R}$ consists of K blocks of length T and thus has total length $L = KT$ bits. These blocks are denoted by $R[k, 0], \ldots, R[k, T-1]$ for $1 \leq k \leq K$, i.e., the randomizer can be viewed as a two-dimensional array of binary random variables (see Figure 1). The secret key $\underline{Z} = [Z_1, \ldots, Z_K]$, where $Z_k \in \{0, \ldots, T-1\}$ for $1 \leq k \leq K$, specifies a position within each block of $\underline{R}$, and is chosen to be uniformly distributed over the key space $S_{\underline{Z}} = \{0, \ldots, T-1\}^K$. Thus the number of bits needed to represent the key is $K \log_2 T$.

$$\begin{array}{cccc} R[1,0] & R[1,1] & \ldots & R[1,T-1] \\ R[2,0] & R[2,1] & \ldots & R[2,T-1] \\ \cdot & \cdot & & \cdot \\ \cdot & \cdot & & \cdot \\ R[K,0] & R[K,1] & \ldots & R[K,T-1] \end{array}$$

Figure 1. The randomizer $\underline{R}$, viewed as a two-dimensional array.

The keystream $\underline{W}$, which is a function of the secret key $\underline{Z}$ and the randomizer $\underline{R}$, is the bitwise modulo 2 sum of the K subsequences of length N within the randomizer starting at the positions specified by the key, where each block (row) of $\underline{R}$ is considered to be extended cyclically, i.e., the second index is reduced modulo T:

$$W_n = \sum_{k=1}^{K} R[k,\ (n-1+Z_k) \bmod T] \tag{1}$$

for $1 \leq n \leq N$, where the summation is modulo 2. The sub-array of the randomizer that determines $\underline{W}$ is denoted by $R^{\underline{Z}}$ and is depicted in Figure 2. A diagram of the sending site of the cipher system is shown in Figure 3. Note that the legitimate receiver who knows the secret key needs to examine only KN of the L random bits, i.e., a fraction N/T of the entire randomizer which is very small when $T \gg N$ as we shall assume.

$$\begin{array}{cccc} R[1,Z_1] & R[1,Z_1+1] & \ldots & R[1,Z_1+N-1] \\ R[2,Z_2] & R[2,Z_2+1] & \ldots & R[2,Z_2+N-1] \\ \cdot & \cdot & & \cdot \\ \cdot & \cdot & & \cdot \\ R[K,Z_K] & R[K,Z_K+1] & \ldots & R[K,Z_K+N-1] \end{array}$$

Figure 2. The sub-array $R^{\underline{Z}}$ of the randomizer $\underline{R}$ is selected by the secret key $\underline{Z}$. All second indices are to be reduced modulo T. The keystream $\underline{W} = [W_1, \ldots, W_N]$ is formed by adding the K rows of $R^{\underline{Z}}$ bitwise modulo 2.

3. Model of Attacks and Main Results

An eavesdropper trying to break the cipher may have (possibly partial) knowledge of the plaintext statistics and may also have some other *a priori* information about the plaintext. Let $P_{\underline{X}}$ be the probability distribution of the plaintext and let V be a random variable, jointly distributed with $\underline{X}$ according to $P_{\underline{X}V}$, that summarizes the eavesdropper's other *a priori* information about $\underline{X}$. Since precise knowledge of $P_{\underline{X}V}$ and thus also of $P_{\underline{X}}$ can only help the eavesdropper and because we assume that he precisely knows these distributions, our proof of security remains valid when the eavesdropper actually has only partial knowledge about $P_{\underline{X}V}$.

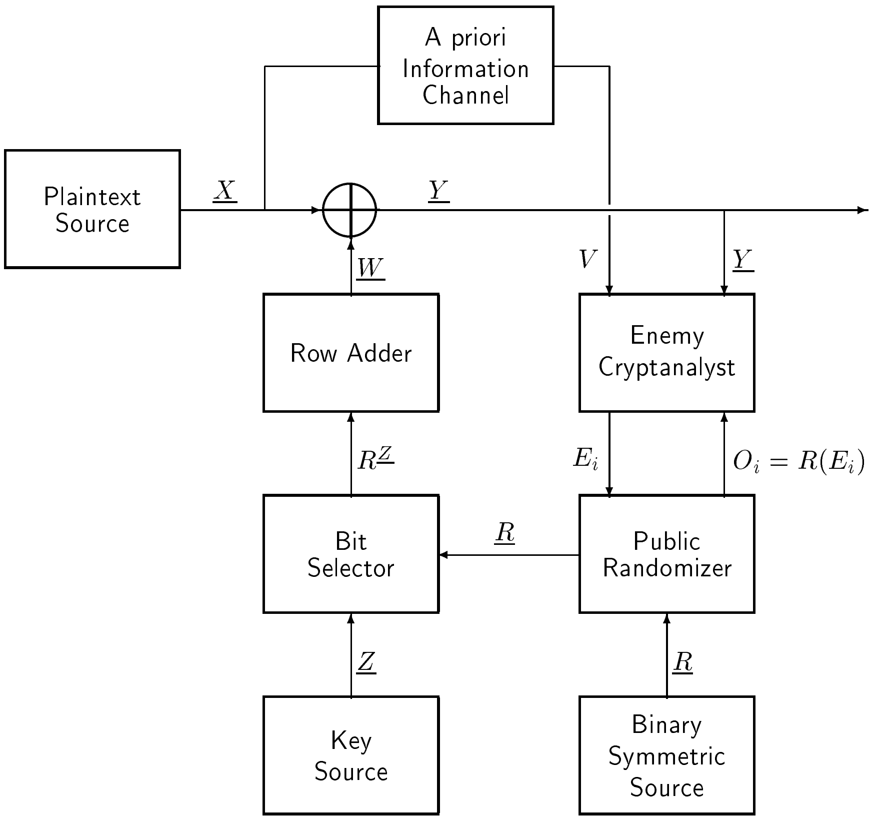

Figure 3. A block diagram of the specific strongly-randomized cipher investigated in Section 3. The public randomizer $\underline{R}$ is an array of independent and completely random binary random variables. The keystream $\underline{W}$ is formed by letting the key $\underline{Z}$ select the sub-array $R^{\underline{Z}}$ of bits of $\underline{R}$ consisting of K rows of length N, and adding these rows bitwise modulo 2. The enemy cryptanalyst or eavesdropper uses an arbitrary, possibly probabilistic, sequential strategy to determine the positions $E_1, E_2, \ldots$ of the randomizer bits $O_1, O_2, \ldots$ that he examines.

Our model of the eavesdropper's *attack* is described in the sequel. We allow the eavesdropper to use an arbitrary, possibly probabilistic, sequential strategy for selecting the positions of the randomizer bits that he examines. At each step of the attack, the eavesdropper can make use of the entire available information, i.e., the cryptogram $\underline{Y}$, the side-information V, and the positions and values of the bits observed so far. Let $E_i = [A_i, B_i]$ denote the address of the i-th randomizer bit examined by the eavesdropper, where A_i and B_i satisfy $1 \leq A_i \leq K$ and $0 \leq B_i \leq T-1$ for $i = 1, 2, \ldots$. Let further $O_i = R(E_i) = R[A_i, B_i]$ denote the observed value of the randomizer bit at position E_i that is examined by the eavesdropper at the i-th step of his attack. Note that the randomizer bit O_i is a binary random variable whose address E_i is a random variable rather than a constant. However, we will make use several times of the fact that, given $E_i = e_i$, O_i corresponds to the randomizer bit $R(e_i)$ at the specific address e_i. We use the notation $E^m = [E_1, \ldots, E_m]$ and $O^m = [O_1, \ldots, O_m]$ for all $m \geq 1$. For a particular sequence $e^m = [e_1, \ldots, e_m]$ of m bit positions, where $e_i = [a_i, b_i]$ with $1 \leq a_i \leq K$ and $0 \leq b_i \leq T-1$ for $1 \leq i \leq m$, $R(e^m) = [R(e_1), \ldots, R(e_m)]$ denotes the corresponding sequence of randomizer bits. Correspondingly, we have $O^m = R(E^m)$ for $m \geq 1$.

The bit position E_m is for $m \geq 1$ determined by the eavesdropper as a (possibly randomized) function of the entire information he possesses at this time, i.e., the cryptogram $\underline{Y}$, the values O^{m-1} of all previously examined bits together with their addresses E^{m-1}, and the *a priori* information V. The eavesdropper's strategy is hence completely specified by the sequence of conditional probability distributions $P_{E_1|\underline{Y}V}$, $P_{E_2|\underline{Y}VE_1O_1}$, $P_{E_3|\underline{Y}VE_1E_2O_1O_2}$ and so on. The following theorem is the main result of this paper.

Theorem: *There exists an event $\mathcal{E}$ such that, for all joint probability distributions $P_{\underline{X}V}$ and for all (possibly probabilistic) strategies for examining bits $O_1, \ldots, O_M$ of $\underline{R}$ at addresses $E_1, \ldots, E_M$,*

$$I\left(\underline{X}; \underline{Y}E^M O^M \,\middle|\, V, \mathcal{E}\right) \;=\; 0 \qquad \text{and} \qquad P(\mathcal{E}) \;\geq\; 1 - N\delta^K,$$

where $\delta = M/KT$ is the fraction of randomizer bits examined by the eavesdropper.

Here $I(\underline{X}; \underline{Y}E^M O^M \,|\, V, \mathcal{E})$ denotes the (mutual) information that $\underline{Y}$, E^M and O^M together give about $\underline{X}$, given that V is known and given that the event $\mathcal{E}$ occurs. The theorem states that if the event $\mathcal{E}$ occurs, then the eavesdropper's total observation $[\underline{Y}, E^M, O^M]$ gives no information about the plaintext $\underline{X}$ beyond the information already provided by V. Clearly, if the eavesdropper knew the value of a random variable V that uniquely determines $\underline{X}$, i.e., such that $H(\underline{X}|V) = 0$, it would make little sense to use a cipher at all. But the point is that, no matter what *a priori* information about the plaintext the eavesdropper has, this does not help him to obtain any *additional* information. For instance, even if the eavesdropper knew all but one bit of the plaintext, he would still get no information about this remaining bit if $\mathcal{E}$ occurs, and the probability of $\mathcal{E}$ could not be reduced by exploiting his virtually complete knowledge about the plaintext. This cipher hence provides perfect secrecy in Shannon's sense if the event $\mathcal{E}$ occurs. Note that the theorem asserts the existence of a high-probability event $\mathcal{E}$, but does not specify it since the specification of $\mathcal{E}$ has no impact on the significance of the result. However, in the proof we will specify such an event. The result that will be proved is in fact stronger than the theorem. It will be proved that if $\mathcal{E}$ occurs, then the eavesdropper would have no information about the plaintext even if he were given the secret key. (Of course, he could avoid the occurrence of the event $\mathcal{E}$ if he knew the secret key before requesting his observations of the randomizer.)

Example: Assume that $K = 50$, $T = 10^{20}$ and let the plaintext be one gigabit, i.e., $N = 2^{30} \approx 10^9$. The key size of this cipher is $50 \cdot \log_2 10^{20} \approx 3320$ bits. The legitimate users need to examine only 50 randomizer bits per plaintext bit. An eavesdropper, however, even if he used an optimal strategy for examining a fraction $\delta = 1/4$ of all bits, i.e., $M = KT/4 = 1.25 \cdot 10^{21}$ bits in total, would have a chance of obtaining any new information about the plaintext not greater than $2^{30} \cdot (1/4)^{50} < 10^{-21}$.

The proof of the above theorem is divided into four steps: the proofs of Lemmas 1 to 4.

Definition: The sequence $e^M = [e_1, \ldots, e_M]$ of $M \geq 1$ bit positions yields a *consistency check* for the key $\underline{z} = [z_1, \ldots, z_K]$ if and only if there exists an integer $n \in [1, N]$ and a subset $\{[1, t_1], [2, t_2], \ldots, [K, t_K]\}$ of $\{e_1, \ldots, e_M\}$ such that

$$t_k - z_k \equiv n - 1 \pmod{T} \quad \textit{for } 1 \leq k \leq K.$$

In other words, e^M yields a consistency check for $\underline{z}$ if and only if when $\underline{Z} = \underline{z}$, $R(e^M)$ determines at least one (the n-th) bit of the keystream $\underline{W} = [W_1, \ldots, W_N]$ or, equivalently, if and only if $R(e^M)$ completely determines at least one column of $R^{\underline{Z}}$ (cf. Figure 2). Furthermore, let $\mathcal{Z}(e^M) \subseteq S_{\underline{Z}}$ denote the set of keys for which e^M yields at least one consistency check, i.e.,

$$\mathcal{Z}(e^M) = \left\{ \underline{z} \in S_{\underline{Z}} : e^M \text{ yields at least one consistency check for } \underline{z} \right\}.$$

The idea behind this definition is that if the eavesdropper knew the plaintext (and hence also the keystream because he knows the ciphertext) and the set $R(e^M)$ of randomizer bits, then, for every key $\underline{z} \in \mathcal{Z}(e^M)$, he could perform one consistency check per keystream bit that he could compute from $R(e^M)$, by comparing the computed keystream bit for the key $\underline{z}$ with the actual keystream bit. If all computed (for key $\underline{z}$) keystream bits agree with the actual keystream bits, the key $\underline{z}$ is still a possible candidate, but if any of the computed keystream bits differs from the corresponding actual keystream bit, then $\underline{z}$ cannot be the actual key. Note that when e^M consists of one bit in each block of $\underline{R}$, then e^M yields exactly one consistency check for N different keys. In general, if e^M consists of m_k bits in the k-th block for $1 \leq k \leq K$, then e^M yields a total number $N \prod_{k=1}^{K} m_k$ of consistency checks, but several of these checks could be for the same key. The event $\mathcal{E}$ introduced in the main theorem will later be defined as the event that the actual key does not belong to the set of keys for which the eavesdropper's set E^M of observed bits yields a consistency check. The following lemma is proved in the Appendix.

Lemma 1: *For all joint probability distributions $P_{\underline{X}V}$, for every sequence $e^M = [e_1, \ldots, e_M]$ of $M \geq 1$ bit positions, and for all $\underline{x}, v, \underline{y}$, $r^M \in \{0,1\}^M$ and $\underline{z} \notin \mathcal{Z}(e^M)$ (such that the* *conditioning event has non-zero probability),*

$$P\left[\underline{X}=\underline{x}\,\middle|\,V=v,\underline{Y}=\underline{y},E^M=e^M,O^M=r^M,\underline{Z}=\underline{z}\right] \;=\; P\left[\underline{X}=\underline{x}|V=v\right].$$

Lemma 2: *For all probability distributions $P_{\underline{X}V}$ and for all (possibly probabilistic) strategies for examining $M \geq 1$ bits $O_1,\ldots,O_M$ of $\underline{R}$ at positions $E_1,\ldots,E_M$, we have*

$$I\left(\underline{X};\underline{Y}E^MO^M\underline{Z}\,\middle|\,V,\underline{Z}\notin\mathcal{Z}(E^M)\right) \;=\; 0.$$

This lemma establishes the first part of the main theorem when $\mathcal{E}$ is defined as the event that $\underline{Z}\notin\mathcal{Z}(E^M)$. It states that if the eavesdropper does not succeed in choosing the bit positions E^M such that $\underline{Z}\in\mathcal{Z}(E^M)$, then he does not obtain any information whatsoever about the plaintext beyond the information already conveyed by V, even if an oracle would give him the key $\underline{Z}$ for free after he has finished his observation. Note that this lemma is true even if the eavesdropper knew the key beforehand, but that clearly a clever eavesdropper would exploit the knowledge of the secret key in order to make the event $\mathcal{E}$ have zero probability.

Proof: The conditional mutual information of Lemma 2 can be written as a difference of conditional uncertainties:

$$I\left(\underline{X};\underline{Y}E^MO^M\underline{Z}\,\middle|\,V,\underline{Z}\notin\mathcal{Z}(E^M)\right)$$

$$= \; H\left(\underline{X}\,\middle|\,V,\underline{Z}\notin\mathcal{Z}(E^M)\right) - H\left(\underline{X}\,\middle|\,V\underline{Y}E^MO^M\underline{Z},\underline{Z}\notin\mathcal{Z}(E^M)\right).$$

It is an immediate consequence of Lemma 1 that both uncertainties are equal. □

Remark: Lemma 2 seems to be an immediate consequence of the fact that when $\underline{Z}\notin\mathcal{Z}(E^M)$, then at least one randomizer bit that has not been observed by the eavesdropper contributes to every keystream bit, and that therefore $\underline{X}$ is concealed by a "perfect one-time pad". However, the proof of Lemma 1 is quite involved due to the fact that the random variable E^M depends on $\underline{Z}$ and thus the event $\mathcal{E}$ is non-trivial. on an alternative and possibly simplified proof of Lemma 2.

It remains to prove the second part of the theorem, viz., to establish a lower bound on the probability of the event $\mathcal{E}$ that $\underline{Z}\notin\mathcal{Z}(E^M)$. For a given set S, let $|S|$ denote its cardinality.

Lemma 3: *For all probability distributions $P_{\underline{X}V}$ and for all (possibly probabilistic) strategies for examining $M \geq 1$ bits $O_1,\ldots,O_M$ of $\underline{R}$ at positions $E_1,\ldots,E_M$,*

$$P\left[\underline{Z}\notin\mathcal{Z}(E^M)\right] \;\geq\; 1-\frac{\max_{e^M}|\mathcal{Z}(e^M)|}{T^K}.$$

Proof: We consider the conditional probability distribution of the key, given all the eavesdropper's information after having observed M bits of $\underline{R}$:

$$\begin{aligned} P\left[\underline{Z}=\underline{z}\,\middle|\,\underline{Y}=\underline{y}, V=v, E^M=e^M, O^M=r^M\right] &= P\left[\underline{Z}=\underline{z}\,\middle|\,\underline{Y}=\underline{y}, V=v, E^M=e^M, R(e^M)=r^M\right] \\ &= P\left[\underline{Z}=\underline{z}\,\middle|\,\underline{Y}=\underline{y}, V=v, R(e^M)=r^M\right] \end{aligned}$$

for all $\underline{z}, \underline{y}, v, e^M$ and r^M. The last step follows from expanding $P[E^M=e^M, \underline{Z}=\underline{z}|\underline{Y}=\underline{y}, V=v, \underline{R}(e^M)=r^M]$ in two different ways and applying equation (8) of the Appendix. Similarly expanding $P[\underline{Z}=\underline{z}, \underline{Y}=\underline{y}\,|\,V=v, R(e^M)=r^M]$ in two different ways and applying equation (5) of the Appendix yields

$$\begin{aligned} P\left[\underline{Z}=\underline{z}\,\middle|\,\underline{Y}=\underline{y}, V=v, R(e^M)=r^M\right] &= P\left[\underline{Z}=\underline{z}\,\middle|\,V=v, R(e^M)=r^M\right] \\ &= P[\underline{Z}=\underline{z}] \\ &= T^{-K} \end{aligned} \tag{2}$$

for all $e^M, \underline{y}, v, r^M$ and $\underline{z} \notin \mathcal{Z}(e^M)$. The second step follows from the fact that $\underline{Z}$ is statistically independent of V and $\underline{R}$. The fact that

$$P\left[\underline{Z}=\underline{z}\,\middle|\,\underline{Y}=\underline{y}, V=v, E^M=e^M, O^M=r^M\right] = T^{-K} \tag{3}$$

for all $e^M, \underline{y}, v, r^M$ and $\underline{z} \notin \mathcal{Z}(e^M)$ appears to be somewhat counter-intuitive since it states that the *a posteriori* probabilities of the keys $\underline{z} \notin \mathcal{Z}(e^M)$ are equal to the *a priori* probabilities even when there exists a key $\underline{z} \in \mathcal{Z}(e^M)$ that satisfies many consistency checks and therefore appears to be the correct key. Equation (3) implies that

$$P[\underline{Z}=\underline{z}|E^M=e^M] = T^{-K}$$

for all e^M and $\underline{z} \notin \mathcal{Z}(e^M)$. Summing these probabilities over all keys $\underline{z} \notin \mathcal{Z}(e^M)$, i.e., over $T^K - |\mathcal{Z}(e^M)|$ terms, we obtain

$$\begin{aligned} P\left[\underline{Z} \notin \mathcal{Z}(E^M)\,\middle|\,E^M=e^M\right] &= \sum_{\underline{z} \notin \mathcal{Z}(e^M)} P\left[\underline{Z}=\underline{z}\,\middle|\,E^M=e^M\right] \\ &= 1 - \frac{|\mathcal{Z}(e^M)|}{T^K}. \end{aligned} \tag{4}$$

Since $P[\underline{Z} \notin \mathcal{Z}(E^M)]$ is equal to the average of $P[\underline{Z} \notin \mathcal{Z}(E^M)|E^M=e^M]$ over all values of e^M, we immediately have

$$P\left[\underline{Z} \notin \mathcal{Z}(E^M)\right] \geq 1 - \frac{\max_{e^M} |\mathcal{Z}(e^M)|}{T^K}. \quad \square$$

Equation (4) demonstrates that the eavesdropper's optimal strategy for making the event $\underline{Z} \notin \mathcal{Z}(E^M)$ as unlikely to occur as possible is simply to make the set $\mathcal{Z}(E^M)$

as large as possible. Notice that, surprisingly, this strategy is independent of $\underline{Y}, O^M$ and V. In other words, letting the selected bit positions $E_1, \ldots, E_M$ depend on the observed bits $O_1, \ldots, O_M$, the cryptogram $\underline{Y}$ and on the *a priori* information V cannot help the eavesdropper in reducing the probability of the event $\mathcal{E}$. However, to base the strategy on $\underline{Y}, O^M$ and V can increase the amount of information that he gets about the plaintext in case that $\mathcal{E}$ does not occur, i.e., when $\underline{Z} \in \mathcal{Z}(E^M)$. Note that although $P[\underline{Z} \notin \mathcal{Z}(E^M)|E^M = e^M]$ equals the number of keys that are not in $\mathcal{Z}(e^M)$ divided by the total number of keys, equation (4) is non-trivial since E^M is a random variable that, because it depends on $\underline{Y}$, also depends on $\underline{Z}$.

Lemma 4: *For every sequence $e^M = [e_1, \ldots, e_M]$ of $M \geq 1$ bit positions,*

$$|\mathcal{Z}(e^M)| \;\leq\; N\left(\frac{M}{K}\right)^K.$$

Proof: Let m_k, for $1 \leq k \leq K$, be the number of randomizer bits specified by e^M that belong to the k-th block of $\underline{R}$, i.e., whose first address component is equal to k. Every subset of elements of e^M of the form $\{[1,t_1],[2,t_2],\ldots,[K,t_K]\}$ yields a consistency check for exactly N keys, namely for the keys $\underline{z} = [(t_1 - x) \bmod T,\ (t_2 - x) \bmod T, \ldots, (t_K - x) \bmod T]$ for $0 \leq x \leq N-1$. There are exactly $\prod_{k=1}^K m_k$ different subsets of the described form and hence there are at most $N\prod_{k=1}^K m_k$ keys for which e^M yields a consistency check. $\prod_{k=1}^K m_k$ is maximized for real m_k under the restriction $\sum_{k=1}^K m_k = M$ by the choice $m_1 = \cdots = m_K = M/K$ for which $\prod_{k=1}^K m_k = (M/K)^K$. Clearly, this maximum is also an upper bound on $\prod_{k=1}^K m_k$ under the restriction that $m_1, \ldots, m_K$ must be integers satisfying $\sum_{k=1}^K m_k = M$. $\square$

Proof of the Theorem: Lemma 2 shows that if we define $\mathcal{E}$ as the event that $\underline{Z} \notin \mathcal{Z}(E^M)$, then $I(\underline{X}; \underline{Y}E^M O^M \underline{Z}|V, \mathcal{E}) = 0$ and therefore also $I(\underline{X}; \underline{Y}E^M O^M|V, \mathcal{E}) = 0$. This last step follows from the two basic facts that mutual information is always non-negative and that joining additional random variables (here $\underline{Z}$) to the information-giving set cannot reduce the information about $\underline{X}$. Lemmas 3 and 4 finally give

$$P[\mathcal{E}] \;\geq\; 1 - \frac{\max_{e^M}|\mathcal{Z}(e^M)|}{T^K} \;\geq\; 1 - N\left(\frac{M}{KT}\right)^K = 1 - N\delta^K. \quad \square$$

4. Two Modifications and Conclusions

In this section, we suggest two modifications of the randomized cipher presented in Section 2 that are more practical in that the size of the public randomizer required to achieve a sufficient level of security can be much smaller. A rigorous proof of security for the first suggested modification would lead to the first cipher that is provably-secure under the sole assumption that the enemy's memory capacity, but not necessarily his computing power, is restricted. The second suggested modification has the potential of leading to an existence proof for secure cryptosystems without necessarily leading to a specific realization.

We first discuss a version of our strongly-randomized cipher in which, instead of having the randomizer *stored* in a publicly-accessible way, it is *broadcasted* by a sender (e.g., a satellite), i.e., the randomizer evolves in time rather than in space. There may exist natural sources of randomness, such as a deep-space radio source, that could be used. Alternatively, the randomizer could be transmitted as a burst of random data over the (insecure) communication channel prior to the transmission of the actual cryptogram. It should be pointed out that the special case where $K = 1$ and the randomizer is interleaved with the cryptogram, corresponds to a finite-plaintext version of the Rip van Winkle cipher [4], which was formulated for an infinite plaintext sequence. Because in this version of our cipher, the randomizer is not available at the time that the ciphertext is transmitted, an enemy must not only examine but also *store* a substantial fraction of the randomizer in order to be able to obtain any information about the plaintext from the subsequently transmitted ciphertext. Thus, if the enemy's memory capacity is at most δ times the number of randomizer bits, then there exists no strategy for storing randomizer bits such that these will later be of any use to the enemy with probability more than $N\delta^K$, where N is the length of the plaintext.

In a situation as considered in [2], [6] and [8] where an enemy is for some reason restricted in the number of bits he may obtain of the message (in our case the randomizer) that is transmitted over the insecure channel but where he is free as to choose the positions of the bits, this "broadcast" cipher is provably-secure even for a moderate size of the randomizer.

However, in order to prove the security of the cipher under the *sole* assumption that the enemy's memory capacity is at most M bits but without restriction on his computing power, one would have to prove the stronger result that, for every function $\{0,1\}^L \rightarrow \{0,1\}^M$, the evaluation of the function with the randomizer as its argument would with overwhelming probability (over choices of the secret key) give only a negligible amount of information about the plaintext. We strongly believe that a result of this type holds and suggest its precise formulation and proof as an open research problem.

A second modification of our cipher is based on the observation that the size of the randomizer can be strongly reduced if the bit access operation can be made more difficult. Assume that a function $f : \{0,1\}^m \rightarrow \{0,1\}$ is available with the property that one can prove a lower bound that is linear in n for $n \ll 2^m$ on the complexity of computing the values $f(x_1), f(x_2), \ldots, f(x_n)$ for n randomly selected arguments $x_1, x_2, \ldots, x_n$ with substantial probability of success. Assume further that an algorithm is known that computes the function f and whose complexity is not much greater than that guaranteed by the lower bound. In other words, to compute the function f (possibly only approximately) for n randomly selected arguments is about n times as difficult as to compute f for one argument. One can then replace every randomizer bit in our cipher by a string of m random bits and the bit-access operation by the evaluation of the function f for the corresponding

m-bit argument. The computational security of such a cipher relies on the fact that the number of times the legitimate users need to evaluate the function is much smaller than the number of randomizer bits that the enemy must compute in order to obtain sufficient information about the plaintext. A possible candidate for f could be the output function of some nonlinear finite automaton after Q clock cycles where the function argument is the initial state and where Q is a fixed large number.

It is well-known (see for example [11]) that, even for moderate m, almost all boolean functions $\{0,1\}^m \to \{0,1\}$ are very difficult to compute, yet no explicit function is known to be difficult to compute. It may similarly be possible to prove the existence of a function f of the type described above without exhibiting a specific example so that this above approach to provable security, as opposed to approaches based on the existence of one-way functions, may lead to the first true existence proof for computationally-secure ciphers, although it would not provide a specific example.

Another interesting property of the randomized cipher of Section 2 is that its security is not compromised when it is used repeatedly with only one initial secret key and with the second and subsequent secret keys being transmitted as part of preceding plaintexts.

Finally, we would like to point out that randomization techniques similar to those presented in this paper may be useful for the construction of practical ciphers, even when the randomizer is not sufficiently long to guarantee a reasonable lower bound on the enemy's computational effort required to break the cipher or when the randomizer is replaced by a pseudo-random sequence.

Appendix

Proof of Lemma 1: Every bit of the keystream $\underline{W}$ is the sum of K randomizer bits (see equation (1)). The crucial observation is that when $\underline{Z} \notin \mathcal{Z}(e^M)$, then, for $1 \leq n \leq N$, at least one of the K randomizer bits contributing to W_n is not contained in the sequence $R(e^M)$ of randomizer bits. Therefore, for every sequence e^M of bit positions, all keystream sequences $w^N \in \{0,1\}^N$, and therefore also all cryptograms $y^N \in \{0,1\}^N$, are equally likely candidates when $\underline{Z} \notin \mathcal{Z}(e^M)$, i.e.,

$$P\left[\underline{Y}=\underline{y} \,\middle|\, \underline{X}=\underline{x}, V=v, R(e^M)=r^M, \underline{Z}=\underline{z}\right] \;=\; 2^{-N} \tag{5}$$

for all $\underline{y}, \underline{x}, v, r^M$ and $\underline{z} \notin \mathcal{Z}(e^M)$.

The fact that the position E_m of the m-th examined bit is determined by a (possibly probabilistic) strategy based on $\underline{Y}, V, E^{m-1}$ and O^{m-1} can be expressed as

$$\begin{aligned} &P\left[E_m = e_m \,\middle|\, E^{m-1}=e^{m-1}, \underline{X}=\underline{x}, \underline{Y}=\underline{y}, V=v, R(e^{m-1})=r^{m-1}, \tilde{R}=\tilde{r}, \underline{Z}=\underline{z}\right] \\ &\qquad = \; P\left[E_m = e_m \,\middle|\, E^{m-1}=e^{m-1}, \underline{Y}=\underline{y}, V=v, R(e^{m-1})=r^{m-1}, \tilde{R}=\tilde{r}\right] \end{aligned} \tag{6}$$

for $m \geq 1$ and for all $e_m, e^{m-1}, \underline{x}, \underline{y}, v, r^{m-1}, \tilde{r}$ and $\underline{z}$, where $\tilde{R}$ is an arbitrary fixed subset of bits of $\underline{R}$. The condition $\tilde{R}=\tilde{r}$ on the right-hand side of (6) can be omitted, but we will make use of (6) in the given form. Note that (6) does not imply that knowing the secret key would not help the enemy in selecting the bit positions, but rather that the bit positions of the real attack do not depend on the secret key $\underline{Z}$ and the plaintext $\underline{X}$ other than by the fact that $\underline{Y}$ depends on $\underline{X}$ and $\underline{Z}$, and V depends on $\underline{X}$. Equation (6) can be used to obtain

$$\begin{aligned}
&P\left[E^M{=}e^M \,\middle|\, \underline{X}{=}\underline{x}, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \\
&\quad = \prod_{m=1}^{M} P\left[E_m = e_m \,\middle|\, E^{m-1}{=}e^{m-1}, \underline{X}{=}\underline{x}, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \\
&\quad = \prod_{m=1}^{M} P\left[E_m = e_m \,\middle|\, E^{m-1}{=}e^{m-1}, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \\
&\quad = P\left[E^M{=}e^M \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \qquad (7)
\end{aligned}$$

for all $e^M, \underline{x}, \underline{y}, v, r^M$ and $\underline{z}$ such that the conditioning event has non-zero probability. In the second step, we have made use of (6) where $\tilde{R}$ is the set $\{R(e_m), R(e_{m+1}) \ldots, R(e_M)\}$. The condition $\underline{Z}{=}\underline{z}$ can be omitted in the last two lines, and this shows that

$$P\left[E^M{=}e^M \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] = P\left[E^M{=}e^M \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M\right] \qquad (8)$$

for all $e^M, \underline{y}, v, r^M$ and $\underline{z}$ such that the conditioning event has non-zero probability. Equation (8) is used in the proof of Lemma 3. We continue the proof of Lemma 1 by noting that, given the event $E^M{=}e^M$, O^M is a fixed set of randomizer bits, namely $O^M = R(e^M)$, and by application of (7), which leads to

$$\begin{aligned}
&P\left[\underline{X}{=}\underline{x} \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, E^M{=}e^M, O^M{=}r^M, \underline{Z}{=}\underline{z}\right] \\
&\qquad = P\left[\underline{X}{=}\underline{x} \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, E^M{=}e^M, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \\
&\qquad = P\left[\underline{X}{=}\underline{x} \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \qquad (9)
\end{aligned}$$

for all $e^M, \underline{x}, \underline{y}, v, r^M$ and $\underline{z}$ such that the first conditioning event has non-zero probability. Using (5) gives

$$\begin{aligned}
P\left[\underline{X}{=}\underline{x} \,\middle|\, \underline{Y}{=}\underline{y}, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] &= P\left[\underline{X}{=}\underline{x} \,\middle|\, V{=}v, R(e^M){=}r^M, \underline{Z}{=}\underline{z}\right] \\
&= P\left[\underline{X}{=}\underline{x} \,\middle|\, V{=}v\right] \qquad (10)
\end{aligned}$$

for all $e^M, \underline{x}, \underline{y}, v, r^M$ and $\underline{z} \notin \mathcal{Z}(e^M)$ such that the first conditioning event has non-zero probability. The last step follows from the facts that $\underline{X}$, $\underline{R}$, and $\underline{Z}$, and thus also $\underline{X}$, $R(e^M)$ and $\underline{Z}$, are statistically independent and that V is statistically independent of $\underline{Z}$ and $\underline{R}$. The lemma follows from (9) and (10). $\square$

Acknowledgements

I would like to thank Jim Massey for motivating this research and for many helpful discussions, and Gilles Brassard for many suggestions improving the presentation of this paper.

References

[1] C.H. Bennett, F. Bessette, G. Brassard, L. Savail and J. Smolin, Experimental quantum cryptography, *Advances in Cryptology - EUROCRYPT'90*, Lecture Notes in Computer Science, Springer Verlag, Berlin, to appear.

[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] J.L. Massey, An introduction to contemporary cryptology, *Proceedings of the IEEE*, Vol. 76, No. 5, 1988, pp. 533-549.

[4] J.L. Massey and I. Ingemarsson, The Rip van Winkle cipher - a simple and provably computationally secure cipher with a finite key, in IEEE Int. Symp. Info. Th., Brighton, England, (Abstracts), June 24-28, 1985, p.146.

[5] U.M. Maurer, Perfect cryptographic security from partially independent channels, preprint, Princeton University, Nov. 1990.

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

[7] U.M. Maurer and J.L. Massey, Cascade ciphers: the importance of being first, presented at the 1990 IEEE Int. Symp. Inform. Theory, San Diego, CA, Jan. 14-19, 1990.

[8] L.H. Ozarow and A.D. Wyner, Wire-tap channel II, *AT&T Bell Laboratories Technical Journal*, Vol. 63, No. 10, 1984, pp. 2135-2157.

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

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

[11] I. Wegener, *The complexity of boolean function*, John Wiley & Sons, Inc., New York, 1987.

[12] A. Wyner, The wire-tap channel, *Bell Systems Technical Journal*, Vol. 54, No. 8, Oct. 1975, pp. 1355-1387.