Information Theoretically Secure Communication in the Limited Storage Space Model

Yonatan Aumann[1] and Michael O. Rabin[2,3]*

[1] Department of Computer Science, Bar-Ilan University, Ramat Gan 52900, Israel
aumann@cs.biu.ac.il
[2] Department of Computer Science, The Hebrew University of Jerusalem
Jerusalem 91904, Israel
rabin@cs.huji.ac.il
[3] DEAS, Harvard University, Cambridge, MA 02138, USA
rabin@deas.harvard.edu

Abstract. We provide a simple secret-key two-party secure communication scheme, which is provably *information-theoretically* secure in the *limited-storage-space* model. The limited-storage-space model postulates an eavesdropper who can execute arbitrarily complex computations, and is only limited in the total amount of storage space (not computation space) available to him. The bound on the storage space can be arbitrarily large (e.g. terabytes), as long as it is fixed. Given this bound, the protocol guarantees that the probability of the eavesdropper of gaining any information on the message is exponentially small. The proof of our main results utilizes a novel combination of linear algebra and Kolmogorov complexity considerations.

1 Introduction

The most basic problem in cryptography is that of communication over an insecure channel, where a Sender $\mathcal{S}$ wishes to communicate with a Receiver, $\mathcal{R}$, while an Eavesdropper, $\mathcal{E}$, is tapping the line. To achieve privacy, the Sender and Receiver may share a common key. In a seminal work, Shannon [10] proved that if the eavesdropper has complete access to the communication line, and is not bounded in any way, then *perfect, information theoretically* secure communication is only possible if the entropy of the key space is at least as large as that of the Plaintext space. In essence, this means that if the eavesdropper is unbounded then the *one-time-pad* scheme, where the size of the secretly shared pad equals the size of the message, is the best possible scheme. This, of course, is impractical for most applications. Thus, to obtain practical solutions one must place some bounds on the eavesdropper's power. Most of modern cryptography has proceeded along the line of assuming that the eavesdropper is computationally bounded and devising schemes that are computationally hard to break.

* Research supported, in part, by NSF Grant NSF-CCR-97-00365, at Harvard University.

Michael Wiener (Ed.): CRYPTO'99, LNCS 1666, pp. 65–79, 1999.
© Springer-Verlag Berlin Heidelberg 1999

These results, though unquestionably impressive, suffer from two drawbacks. First, they are based on unproven complexity assumptions. Second, many of the methods tend to require considerable computations, even from the Sender and Receiver. In theory, any polynomial computation is considered feasible for the participants. In practice, some of the polynomials are so large as to render the solution impractical. In other cases, even when the solutions are practical (e.g. RSA), the amount of computation limits their use to high-security applications or key-exchange sessions.

The Limited Storage Space Model. Recently, there has been development in constructing secure communication protocols *not* based on computational complexity. In this paper we consider the *limited storage space* model, where the security guarantees are based on the limited *storage space* available to the eavesdropper. This model, first introduced by Maurer [7], assumes that there is a known bound, possibly very large - but fixed, on the amount of storage-space available to both the Sender and Receiver, and a (possibly much larger but fixed) bound on the storage space available to the Eavesdropper. It is important to differentiate between this model and the well-known so called *bounded space* model (e.g. log-space). The bounded space model considers cases where the space is usually very limited, and serves as *computing* space. In this case, the space-bound is in effect a limitation on computational power. The *limited storage space* model, by contrast, allows very large amounts of storage space, placing the limit only to practically feasible capacities (e.g. a *terabyte* or several *terabytes*) of storage. At the same time we do not stipulate any limitations on the computational power of the eavesdropper who is trying to subvert the secrecy of the protocol. Given the bound on the eavesdropper's storage space, the model enables to obtain results which are *information theoretic*, not depending on any unproven assumptions. Furthermore, the model enables us to construct a very simple and efficient protocol, requiring very little computation and storage space for the Sender and Receiver. Informally, the limited storage space model assumes a publicly accessible source of random bits, such as a high rate broadcast of a string α of random bits, equally accessible to the Sender, Receiver and Eavesdropper. Let the length of α be $|\alpha| = nm$, where m is the length of the message to be securely sent. It is assumed that the Eavesdropper is limited to storing $E < n$ bits, say $E = n/5$. The Eavesdropper can listen to the whole of α and compute and store any function $f(\alpha)$, provided that $|f(\alpha)| \leq E$. In one version, the model postulates that the sender and receiver share a secret key s where $|s| = O(logn)$. Using s and listening to α, the sender and receiver both read and store ℓ chosen locations of α and compute from those bits a one-time pad $X, |X| = m$. The pad X is used by the Sender to encrypt a message $M, |M| = m$. We show that the pad X can be used as a secure one-time-pad for secret communication between the Sender and the Receiver.

The Limited Storage Space Model - Previous Work. In a ground-breaking work, Maurer [7] presented the first protocol for private-key secure communication in the Limited-Storage-Space model described above. However, in [7], the proof of the security of the protocol is provided only for the case where the bound on the eavesdropper is not only on the *space* available to her, but also on the total number of random bits she can *access*. In particular, it is assumed that the eavesdropper can only access a constant fraction of the random bits of α. The analysis of [7] does not provide a proof for the general limited-storage-space case, where the eavesdropper can access all the bits of α, can perform any computation on these bits, and store the results of the computation in the bounded space available to him. It was left as an open question in [7] if any security result can be proved using the limited-storage-space assumption alone. Recently, Cachin and Maurer [4] provided a protocol for which they prove security based solely on the limited-storage-space assumption. However, this protocol is considerably more complex than the original protocol, employing advanced privacy amplification techniques. Their proof uses sophisticated Renyi entropy considerations. To assure that the probability of revelation to the eavesdropper of the secret message M be smaller than ϵ, the protocol of [4] requires the Sender and Receiver to store $\ell \log n$ and transmit ℓ bits, where $\ell = 3/\epsilon^2$. Thus if we prudently require $\epsilon = 10^{-6}$, we get that $\ell = 3 \cdot 10^{12}$; the length n of the random string α can be 2^{40}, so that $\log n = 40$. Thus for this choice of ϵ, the Sender and Receiver have to store and transmit very large numbers of bits. Furthermore, the protocol calls for a multiplication operation in a large field of ℓ-bit numbers.

Our Results. In this paper we return to the simple original protocol of [7], and show that security can be proved, for a slightly modified protocol, based only on the limited storage space assumption. Altogether, we obtain a secure secret key communication scheme secure against any eavesdropper with limited storage space. The protocol is very simple for the Sender and Receiver, necessitating only the elementary XOR operations, and minimal storage space. Specifically, the secret shared key s has $k \log n$ bits, where k is a security parameter. For a secret transmission of a message M of length m, the Sender and Receiver have each to read from α and store just km bits and the one-time pad X is computed from those bits, like in [7], by just XOR operations. Finally, the probability of the adversary Eavesdropper to gain even one-bit information on the message M, is smaller than $2^{-k/5}$, i.e. exponentially small in k. The results are obtained using novel techniques, employing linear algebra and Kolmogorov complexity arguments. An exact formulation of the results is provided in Section 2. We note, however, that our protocol requires a longer random string α, than does that of [7]. It remains a open problem to reduce this number. It will also be interesting to see whether the new methods can considerably improve the constants in the protocol of [4], or some variation of this protocol.

Related Work. In a series of papers, Maurer and colleagues [7,8,2,4] consider secure communication solutions not based on computation complexity. In [8,2] the authors consider the setting where all parties (honest and eavesdropper) communicate over noisy channels. The model of limited-storage-space was first introduced in [7] and further developed in [4], as discussed above. The limited

```
Communication Protocol.
Message $M = (M_1, \ldots, M_m) \in \{0,1\}^m$. Secret key $s = (\sigma_1, \ldots, \sigma_k) \in \{1, \ldots, n\}^k$
for $i = 1$ to $m$ do
    for $j = 1$ to $n$ do
        Broadcast random $\alpha_j^{(i)}$ (either produced by $\mathcal{S}$, $\mathcal{R}$ or an outside source).
        If $j \in s$ then
            $\mathcal{R}$ and $\mathcal{S}$ store $\alpha_j^{(i)}$ in memory
    end for loop
    $\mathcal{S}$ and $\mathcal{R}$ set $X_i := \bigoplus_{j=1}^{k} \alpha_{\sigma_j}^{(i)}$
end for loop
$\mathcal{S}$ and $\mathcal{R}$ set $X = (X_1, \ldots, X_m)$
$\mathcal{S}$ computes $Y = X \oplus M$. Sends $Y$ to $\mathcal{R}$.
$\mathcal{R}$ decryptes $M = X \oplus Y$
```

Fig. 1. Communication Protocol

storage space model in the context of Zero-Knowledge proofs was first studied by De Santis, Persiano and Yung [5], and then by Aumann and Feige [1].

2 The Protocol

We consider a secret key setting, where the sender and receiver share a small secret key. Sender $\mathcal{S}$ wants to send a message $M \in \{0,1\}^m$ to the receiver $\mathcal{R}$ over an insecure channel. An eavesdropper $\mathcal{E}$ may be tapping the communication line between $\mathcal{S}$ and $\mathcal{R}$. We assume that there is a known bound, E, on the total storage space available to $\mathcal{E}$. Let k be a security parameter, and let $n = 5E$. Given a private key of size $k \log n$ chosen uniformly at random, our scheme guarantees information-theoretic secrecy with probability $\geq 1 - 2^{k/5}$.

The scheme is essentially that of [7] (with one exception, which is discussed in the end of this section). First, $\mathcal{S}$ and $\mathcal{R}$ produce a shared "one-time-pad", $X = (X_1, \ldots, X_m)$. Then $\mathcal{S}$ computes $Y = X \oplus M$. He then sends Y to $\mathcal{R}$, who then computes $M = X \oplus Y$ to obtain the original message.

To produce each shared bit, X_i, $i = 1, \ldots, m$, the protocol employs a long random "noise" string $\alpha^{(i)}$, of size $n = 5E$, broadcasted from $\mathcal{S}$ to $\mathcal{R}$. Alternatively, the string $\alpha^{(i)}$ may be available to both and $\mathcal{R}$ from an outside source, e.g. random channel noise, or a special public "noise" broadcast. We assume that $\mathcal{E}$ has full access to $\alpha^{(i)}$ while it is being broadcast. In particular, $\mathcal{E}$ can access all the bits of $\alpha^{(i)}$ and perform on them *any* possible computation, polynomial or non-polynomial. The sole restriction is that the total space available to $\mathcal{E}$, for storing the output of her computation on $\alpha^{(i)}$ is bounded by E. Let s be the secret key. We interpret s as a sequence $\sigma_1, \ldots, \sigma_k$ of integers in $[1, n]$. As the random string $\alpha^{(i)}$ is broadcasted, both the $\mathcal{S}$ and the $\mathcal{R}$ retain the k bits $\alpha_{\sigma_1}^{(i)}, \ldots, \alpha_{\sigma_k}^{(i)}$. Both players then compute $X_i = \bigoplus_{j=1}^{k} \alpha_{\sigma_j}^{(i)}$, to produce the i-th random bit of the one-time-pad X. A detailed description of the protocol is provided in Figure 1.

We prove that the bit X_i can be used as a secure one-time-pad bit for communication between $\mathcal{S}$ and $\mathcal{R}$. In particular, we show that the probability of $\mathcal{E}$

computing X_i correctly is $\leq 1/2 + 2^{-k/5}$. This holds true even if s is revealed to $\mathcal{E}$ after his computation on the string $\alpha^{(i)}$ is completed. Note that for the honest players the scheme is very simple, requiring only $k \log n$ storage space and XOR computations.

To understand the protocol, consider the case where $\mathcal{E}$ can only store original bits of α, without performing any computations. The storage space available to $\mathcal{E}$ allows her to store at most one fifth of the bits of α. Thus, for any i, the probability that all bits $\alpha^{(i)}_{\sigma_i}$ are available to $\mathcal{E}$ is exactly 5^{-k}. If $\mathcal{E}$ happens to have stored exactly the right bits, then she knows the value of X_i. Otherwise, it can be shown that $\mathcal{E}$ has no information on X_i, and can only guess its value with probability 1/2. Thus, if $\mathcal{E}$ can only store original bits of α then her probability of guessing X_i is $1/2 + 5^{-k}$. A result in this spirit was proven by Maurer [7]. In our model, however, we allow $\mathcal{E}$ to perform any computation on the bits of α, and retain any *function* of these bits, provided that the total size of the output does not exceed E bits. We prove that the extra computations can give $\mathcal{E}$ at most a very marginal advantage. We prove that in any case, the probability of $\mathcal{E}$ of correctly computing even *one bit* of knowledge about X is at most $1/2 + 2^{-k/5}$. The bound is information theoretic and does not depend on any unproven complexity assumptions.

The Main Theorem. We prove a strong statement concerning the security of the scheme. We show that even if the Eavesdropper is provided with the full key s after α is broadcasted, she still has negligible probability of gaining any information on the message M. Furthermore, we apply the storage restriction only at one point in time - namely, the time immediately after α is transmitted. Thus, the dynamics can be described as follows:

Phase I:
(a) The stream α is generated and transmitted ($\alpha \in \{0,1\}^{nm}$).
(b) The Eavesdropper can perform any computation on α, with no restrictions on space or time.

Following this phase, the eavesdropper can store $n/5$ bits. We denote by η, $|\eta| = n/5$, the information stored after this phase.

Phase II: The Eavesdropper is provided with $Y = X \oplus M$ and the key s. Based on η, Y and s, the Eavesdropper tries to gain information on M.

Thus, any algorithm, A of the eavesdropper, is actually a pair of algorithms $A = (A_1, A_2)$, where A_1 is the algorithm for the first phase and A_2 is the algorithm for the second phase (after $\mathcal{E}$ receives s). The first algorithm gets α as an input, and outputs $\eta = A_1(\alpha)$, with $|\eta| = n/5$. The second algorithm A_2 gets s, η and Y as inputs, and outputs a single bit, $\delta = A_2(\eta, s, Y)$. Accordingly, we denote the entire eavesdropper's algorithm as $A(\alpha, s, Y)$.

We prove that, for any message M, sent from $\mathcal{S}$ to $\mathcal{R}$, the probability of $\mathcal{E}$ to gain any additional information on M from the protocol is exponentially small (in k). Specifically, consider any two possible distributions on messages, $D^{(0)}$ and $D^{(1)}$. The Eavesdropper wishes to know from which of the two distribution the message was drawn. For a message M and an eavesdropper's algorithm $A = (A_1, A_2)$, we denote by $A(M)$ the entire output of the algorithm for a message M, i.e. $A(M) = A_2(A_1(\alpha), s, X(s, \alpha) \oplus M)$, with α and s chosen uniformly at random.

Theorem 1. *For n large enough, for any distributions $D^{(0)}, D^{(1)}$, and any Eavesdropper's algorithm A, which uses at most $E = n/5$ storage space,*

$$\left|\Pr\left[A(M) = 1 | M \in D^{(1)}\right] - \Pr\left[A(M) = 1 | M \in D^{(0)}\right]\right| \leq 2^{-k/5},$$

where the probability is taken over the random choices of α, the random secret key s, and the random choices of M from the distributions $D^{(0)}$ and $D^{(1)}$.

In particular the theorem says that if there are only two possible messages, $M^{(0)}$ and $M^{(1)}$, and the Eavesdropper wishes to know which message was sent, then she has only an exponentially small probability of success. Note that there is no limit on the time complexity of the eavesdropper's algorithm A. The only limit is that the storage is bounded by $E = n/5$. Also note that the result is *non-uniform*, in the sense that the algorithm A may be tailored to the specific distributions $D^{(0)}$ and $D^{(1)}$.

We note that, in addition to providing provable secrecy, our scheme provides two important security features, not usually provided in complexity-based schemes. First, as noted, the secrecy is guaranteed even if following the transmission the secret-key is fully revealed. Thus, the system is secure against future leakage of the key. Secondly, the system is also secure against the event in which the eavesdropper $\mathcal{E}$ subsequently obtains more storage space. The bound we have on $\mathcal{E}$'s storage space need only be true for her *current* available storage. Any future additional storage will not give her any advantage. Thus, future advances in storage technology do not threaten the secrecy of current communications. This is in contrast to most-all complexity-based schemes, were messages can be stored now and deciphered later, if and when computing technology allows (e.g. using quantum computers to factor numbers).

How many random bits are necessary? The protocol as described above uses n random bits for each message bit, for a total of nm random bits. The [7] protocol, in contrast, uses only n bits in total, for the entire message. This is achieved in the following way. A single string α of length n is broadcast. The bit X_1 is defined as in our protocol. For the subsequent bits, X_i is defined to be $X_i := \bigoplus_{j=1}^{k} \alpha_{s_{j+i-1}}$. Thus, all the X_i's are obtained from a single α of length n. For the model of [7], where the eavesdropper can only *access* $E = n/5$ bits, Maurer proves that the reduced bit protocol suffices to guarantee security. The proof, however, does not carry over to our setting, where the eavesdropper can access all the bits and the sole bound is on the space available to the eavesdropper. Our proof, as described in the next section, necessitates n random bits for each message bit. It remains an important open problem to extend the proof to allow for a similar number of bits as in the original [7] protocol.

3 The Proof

We provide the proof in several stages. First, we prove that it is sufficient to prove the theorem for the case where there are only two possible messages. Next, we prove the theorem for the case of one-bit messages. We do this by proving that if for a given α, the knowledge of η ($|\eta| = n/5$) helps the eavesdropper in reconstructing the one-bit message M for many different secret-keys s, then the Kolmogorov Complexity of α must be small. Since most α's have high Kolmogorov complexity, this shows that the eavesdropper's probability of being correct is small. Next, having proven the theorem for single bit messages, we consider the case of long messages that differ in a single bit. Finally, we prove the full theorem.

Notations: For a string $\alpha^{(i)} = (\alpha_1^{(i)}, \ldots, \alpha_n^{(i)})$, ($\alpha^{(i)} \in \{0,1\}^n$), and $s = (\sigma_1, \ldots, \sigma_k)$ we denote $s(\alpha^{(i)}) = \bigoplus_{j=1}^{k} \alpha_{\sigma_j}^{(i)}$. For $\alpha = (\alpha^{(1)}, \ldots, \alpha^{(m)})$, we denote $s(\alpha) = (s(\alpha^{(1)}), \ldots, s(\alpha^{(m)}))$. We also denote $X(s, \alpha) = s(\alpha)$.

3.1 From Distributions to Messages.

Lemma 1. *Theorem 1 holds iff for n large enough, for any two messages $M^{(0)}$ and $M^{(1)}$ and any Eavesdropper's algorithm A, which uses at most $n/5$ storage space,*

$$\left|\Pr\left[A(M^{(1)}) = 1\right] - \Pr\left[A(M^{(0)}) = 1\right]\right| \leq 2^{-k/5},$$

where the probability is taken over the random choices of α and the secret key s.

Proof. Clearly, if Theorem 1 holds, then in particular it holds when the distribution $D^{(1)}$ is concentrated solely on $M^{(1)}$ and $D^{(0)}$ solely on $M^{(0)}$.

Conversely, suppose that

$$\left|\Pr\left[A(M^{(1)}) = 1\right] - \Pr\left[A(M^{(0)}) = 1\right]\right| \leq 2^{-k/5},$$

for any two messages. Let $D^{(0)}$ and $D^{(1)}$ be two distributions. W.l.o.g. assume that

$$\Pr\left[A(M) = 1 | M \in D^{(1)}\right] - \Pr\left[A(M) = 1 | M \in D^{(0)}\right] \geq 0.$$

Let $M^{(1)}$ be the message such that $\Pr\left[A(M^{(1)}) = 1\right]$ is the largest, and let $M^{(0)}$ be the message such that $\Pr\left[A(M^{(0)}) = 1\right]$ is the smallest. Then,

$$\Pr\left[A(M) = 1 | M \in D^{(1)}\right] - \Pr\left[A(M) = 1 | M \in D^{(0)}\right] =$$
$$\sum_M \Pr\left[A(M) = 1\right] \Pr_{D^{(1)}}\left[M\right] - \sum_M \Pr\left[A(M) = 1\right] \Pr_{D^{(0)}}\left[M\right] \leq$$
$$\left|\Pr\left[A(M^{(1)}) = 1\right] - \Pr\left[A(M^{(0)}) = 1\right]\right| \leq 2^{-k/5}.$$

□

Thus, it is sufficient to focus on the case of just two possible messages.

3.2 Single Bit Secrecy

We now prove the theorem for the case of a single bit message, i.e. $m = 1$. We use the following notations. Since $m = 1$, we have $\alpha = \alpha^{(1)}$. Thus, we omit the superscript from α and write $\alpha = (\alpha_1, \ldots, \alpha_n)$ ($\alpha_j \in \{0,1\}$). Similarly, we denote $X = X_1$. Let $K = n^k$ and $N = 2^n$. Let $S = (s_1, s_2, \ldots, s_K)$ be an enumeration of all possible secret keys, and $\mathcal{A} = (\alpha_1, \ldots, \alpha_N)$ be an enumeration of all strings of length n. For a bit b we denote $\bar{b} = (-1)^b$ (i.e. we replace 1 by -1 and 0 by 1). For a sequence $B = (b_1, \ldots, b_K)$ we denote $\bar{B} = (\bar{b}_1, \ldots, \bar{b}_K)$. For a sequence α we denote $v(\alpha) = (\overline{s_1(\alpha)}, \ldots, \overline{s_K(\alpha)})$.

Preliminaries. For $v \in \{1,-1\}^K$ define the *discrepancy* of v as $d(v) = |\sum_{i=1}^K v_i|$.

Lemma 2. *Let $\alpha \in \{0,1\}^n$ be such that the fraction of 1's and the fraction of 0's in α is no less then $1/8$, then*

$$d(v(\alpha)) < \frac{K}{2^{0.4k}}.$$

Proof. Assume k is odd. Let p be the fraction of 1's in α. Set $q = 1-p$. Consider a random choice of $s = (s_1, \ldots, s_k) \in S$. Since k is odd,

$$s(\alpha) = 1 \Leftrightarrow |\{i : \alpha_{s_i} = 1\}| \text{ is odd}$$

For any $0 \le t \le k$,

$$\Pr\left[|\{i : \alpha_{s_i} = 1\}| = t\right] = \binom{k}{t} p^t q^{k-t}.$$

Thus,

$$\Pr\left[s(\alpha) = 1\right] = \sum_{t \text{ odd}} \binom{k}{t} p^t q^{k-t}. \tag{1}$$

Now,

$$1 = (p+q)^k = \sum_t \binom{k}{t} p^t q^{k-t} \tag{2}$$

$$(p-q)^k = \sum_t \binom{k}{t} (-1)^{k-t} p^t q^{k-t} \tag{3}$$

Since k is odd, when t is even, $(-1)^{k-t} = -1$. Thus, adding (2) and (3),

$$\frac{1+(p-q)^k}{2} = \sum_{t \text{ odd}} \binom{k}{t} p^t q^{k-t}.$$

Together with (1), we get,

$$\frac{1}{2} - \frac{1}{2^{0.4k+1}} < \Pr\left[s(\alpha) = 1\right] = \frac{1+(p-q)^k}{2} < \frac{1}{2} + \frac{1}{2^{0.4k+1}}.$$

For k even, an analogous argument works by considering $\Pr[s(\alpha) = 0]$ and the number of zeros in α.

Thus,

$$D(v(\alpha)) = K|\Pr[s(\alpha) = 1] - \Pr[s(\alpha) =]| < \frac{K}{2^{0.4k}}.$$

□

Let

$$\mathcal{D} = \left\{\alpha \in \mathcal{A} : d(v(\alpha)) > \frac{K}{2^{0.4k}}\right\} \tag{4}$$

be the set of vectors α with a large discrepancy.

Corollary 3 *For $c \geq 0.798$, $|\mathcal{D}| \leq 2^{cn}$.*

Proof. For a string α denote by $z(\alpha)$ the number of zeros in α. By Lemma 2, $\alpha \in \mathcal{D}$ only if $z(\alpha) < n/8$ or $z(\alpha) > 7n/8$. For a random α, $E(z(\alpha)) = n/2$. Thus, by the Chernoff bound ([9] p. 70, Theorem 4.2),

$$\Pr[z(\alpha) < n/8 \text{ or } z(\alpha) > 7n/8] = 2\Pr\left[z(\alpha) < \frac{n}{8}\right] \leq$$

$$2\Pr\left[z(\alpha) < \left(1 - \frac{3}{4}\right)E(z(\alpha))\right] \leq 2\exp\left(-\frac{n}{2}\left(\frac{3}{4}\right)^2\frac{1}{2}\right) < 2^{-0.202n}.$$

for n sufficiently large. □

Lemma 4. *Let $\alpha, \beta \in \mathcal{A}$,*

$$v(\alpha) \otimes v(\beta) = v(\alpha \oplus \beta),$$

where $\otimes$ is the coordinate-wise multiplication.

Proof. For each $s \in S$

$$\overline{s(\alpha)} \cdot \overline{s(\beta)} = (-1)^{s(\alpha)} \cdot (-1)^{s(\beta)} = (-1)^{(s(\alpha))\oplus(s(\beta))} = (-1)^{s(\alpha\oplus\beta)} = \overline{s(\alpha \oplus \beta)}$$

□

Single Bit Case. For the case of single bit messages the only two possible messages are $M = 0$ and $M = 1$. For a given Y, $M = 1$ iff $X = 1 - Y$. Thus, in order for the Eavesdropper to distinguish between $M = 1$ and $M = 0$, she must be able to distinguish between $X = 0$ and $X = 1$. Consider an algorithm $A = (A_1, A_2)$ of the Eavesdropper for guessing $X = s(\alpha)$. Given α, let $\eta = A_1(\alpha)$ be the information stored by the Eavesdropper following the first phase. By definition $|\eta| = n/5$. Next, when provided with s, and η, algorithm $A_2(\eta, s)$ outputs a bit X, in hope that $X = s(\alpha)$ (note that in this case A_2 does not get Y, as it only needs to guess X). For a message M, we denote $A(M) = A_2(A_1(\alpha), s)$, where α and s are chosen at random.

Definition 1 *We say that A_2 is* good *for α if there exists an $\eta \in \{0,1\}^{n/5}$ such that $\Pr[A_2(\eta, s) = s(\alpha)] \geq \frac{1}{2} + \frac{1}{2^{0.4k/2}}$, where the probability is taken over the random choices of s.*

We prove that for any A_2, for almost all α's, A_2 is not good.

Let us concentrate on a given η. For the given η, let

$$B = B_\eta = (A_2(\eta, s_1), \ldots, A_2(\eta, s_K))$$

(where $s_1, \ldots, s_K$ is the enumeration of all possible keys). The vector B is an enumeration of the answers of A_2 on input η, given the various keys s_i. For a given answer vector B, let

$$L_B = \left\{ \alpha : |\bar{B} \cdot v(\alpha)| \geq \frac{2K}{2^{0.4k/2}} \right\}.$$

L_B is the set of α's for which the answers in B are good. By definition, if A_2 is good for α then $\alpha \in L_{B_\eta}$, for some η. We now bound the size of L_B, for any B.

Let V be the $K \times N$ matrix, whose columns are $v(\alpha_1), \ldots, v(\alpha_N)$, i.e.

$$V = \begin{pmatrix} \overline{s_1(\alpha_1)} & \overline{s_1(\alpha_2)} & \cdots\cdots & \overline{s_1(\alpha_N)} \\ \overline{s_2(\alpha_1)} & \vdots & & \vdots \\ \vdots & \vdots & & \vdots \\ \vdots & \vdots & & \vdots \\ \overline{s_K(\alpha_1)} & \cdots & \cdots\cdots & \overline{s_K(\alpha_N)} \end{pmatrix}$$

Consider a specific $B \in \{0,1\}^K$. Let $L^+ = \left\{\alpha : \bar{B} \cdot v(\alpha) \geq \frac{2K}{2^{0.4k/2}}\right\}$ and $L^- = L_B - L^+$. We bound the size of L^+. The proof for L^- is analogous.

Let 1_{L^+} be the characteristic vector of L^+ ($1_{L^+} \in \{0,1\}^N$). For any i, $\bar{B} \cdot V \cdot e_i = \bar{B} \cdot v(\alpha_i)$ (where e_i is the unit vector with 1 in the i-th coordinate). Thus, by definition of L^+,

$$\bar{B} \cdot V \cdot 1_{L^+} \geq |L^+| \cdot \frac{2K}{2^{0.4k/2}}. \tag{5}$$

On the other hand, by the Cauchy-Schwartz inequality,

$$\bar{B} \cdot V \cdot 1_{L^+} \leq \left\|\bar{B}\right\| \cdot \left\|V \cdot 1_{L^+}\right\|. \tag{6}$$

Since $\bar{B} \in \{1, -1\}^K$, we have

$$\left\|\bar{B}\right\| = \sqrt{K}. \tag{7}$$

Next, by definition,

$$\left\|V \cdot 1_{L^+}\right\|^2 = 1_{L^+}^T \cdot V^T V \cdot 1_{L^+}.$$

Consider the matrix $H = V^T V$. Set $H = (h_{i,j})$. By definition and Lemma 4, $|h_{i,j}| = d(v(\alpha_i) \otimes v(\alpha_j)) = d(v(\alpha_i \oplus \alpha_j))$. Thus,

$$|h_{i,j}| \leq \begin{cases} K & \alpha_i \oplus \alpha_j \in \mathcal{D} \\ \frac{K}{2^{0.4k}} & \text{otherwise} \end{cases}$$

(where $\mathcal{D}$ is the set of α's with large discrepancy (eq. (4)). Let

$$\delta(i) = \begin{cases} 1 & \alpha_i \in L^+ \\ 0 & \alpha_i \notin L^+ \end{cases}$$

We have,

$$\|V \cdot 1_{L^+}\| = 1_{L^+}^T \cdot V^T V \cdot 1_{L^+} = \sum_{i=1}^{N} \sum_{j=1}^{N} |h_{i,j}| \delta(i)\delta(j) \leq$$

$$\sum_{i=1}^{N} \delta(i) \left(\sum_{j:\alpha_i \oplus \alpha_j \in \mathcal{D}} |h_{i,j}| + \sum_{j:\alpha_i \oplus \alpha_j \notin \mathcal{D}} |h_{i,j}| \delta(j) \right) \leq$$

$$\sum_{i:\alpha_i \in L^+} \left(\sum_{j:\alpha_i \oplus \alpha_j \in \mathcal{D}} |h_{i,j}| + \sum_{j:\alpha_j \in L^+, \alpha_i \oplus \alpha_j \notin \mathcal{D}} |h_{i,j}| \right) \leq$$

$$\left|L^+\right| \left(2^{cn} K + \left|L^+\right| \frac{K}{2^{0.4k}} \right) \tag{8}$$

(for c as in corollary 3). Combining Equations (5), (6), (7), and (8), we get,

$$|L^+| \cdot \frac{2K}{2^{0.4k/2}} \leq \sqrt{K} \left|L^+\right|^{1/2} \left(2^{cn} K + \frac{|L^+| \, K}{2^{0.4k}} \right)^{1/2}$$

Solving for $|L^+|$ gives,

$$3|L^+| \leq 2^{cn+0.4k}$$

Similarly, for L^-, $3|L^-| \leq 2^{cn+0.4k}$.

In all we conclude,

Lemma 5. *For any possible answer vector B, $|L_B| \leq 2^{cn+0.4k}$.*

We now relate the Kolmogorov complexity of α to the success probability of A_2 on α (see [6] for a comprehensive introduction to Kolmogorov Complexity). For a string α denote by $C(\alpha|A_2)$ the Kolmogorov complexity of α, with regards to a description of A_2 .

Corollary 6 *For any algorithm A_2 and $\alpha \in \{0,1\}^n$ if A_2 is good for α, then*

$$C(\alpha|A_2) \leq \frac{n}{5} + cn + 0.4k.$$

Proof. If A_2 is good then, given A_2, the string α can be fully specified by:

1. η - $n/5$ bits, and
2. the index i_α of α in L_{B_η} (in lexicographic order) - at most $cn + 0.4k$ bits,

Given this information, α is constructed as follows. First, using A_2 and η, the answer vector $B_\eta = (A_2(\eta, s_1), \ldots, A_2(\eta, s_K))$ is constructed. Next, all possible strings $\beta \in \{0,1\}^n$ are constructed one by one, in lexicographic order. For each string, β, we check if $|\bar{B} \cdot v(\beta)| \geq \frac{2K}{2^{0.4k/2}}$. If so, it is marked as a member of L_B. The desired string α is the i_α member of L_B. □

Since $c < 0.8$, we have $\frac{n}{5} + cn + 0.4k < n$ for n sufficiently large. Thus, only a negligible fraction ($2^{-\Omega(n)}$) of the α's have $C(\alpha|A_2) \leq \frac{n}{5} + cn + 0.4k$. Thus, since α is chosen at random, there is only a negligible probability that A_2 is good for α, even for an optimal A_1.

3.3 Multi-bit Secrecy - The One-more Bit Problem

Next, we consider the secrecy of the "one-time-pad" X for the case $m > 1$. Thus, we consider the case where $X = (X_1, \ldots, X_m)$. We consider the following problem. Suppose that the Eavesdropper is given all of X except for the last bit, $X_m = s(\alpha^{(m)})$. Her aim is to guess the value of X_m. We call this *the One-more Bit Problem.*

Thus, we consider the following scenario. First, $\alpha = (\alpha_1, \ldots, \alpha_m)$ is transmitted. The Eavesdropper may compute any function $\eta = A_1(\alpha)$, such that $|\eta| = n/5$, and retain η alone. Next, she is provided with s and X_i for all $i = 1, \ldots, m-1$. She must now guess X_m.

As above, let A_2 be the algorithm the Eavesdropper uses to guess X_m, given η, s and the X_i's. Denote $X^- = (X_1, \ldots, X_{m-1})$ and $\alpha^- = (\alpha^{(1)}, \ldots, \alpha^{(m-1)})$. Note that X^- is fully determined by α^- and s. Thus, we write $X^- = X^-(\alpha^-, s)$.

Using A_2 we can construct another algorithm $\hat{A}_2$ which, given $\eta = A_1(\alpha)$, s and α^- guesses X_m. (The difference between A_2 and $\hat{A}_2$ is that $\hat{A}_2$ gets α^- as input, instead of X^-.) Algorithm $\hat{A}_2$ works as follows. First, $\hat{A}_2$ computes $X^- = X^-(s, \alpha^-)$. Then, it runs $A_2(\eta, s, X^-)$. Thus, the success probability of any algorithm A_2 is at most the success of the best algorithm $\hat{A}_2$. We now bound the success probability of $\hat{A}_2$.

For a given η and α^-, let $B = B_{\eta,\alpha^-} = (\hat{A}_2(\eta, s_1, \alpha^-), \ldots, \hat{A}_2(\eta, s_K, \alpha^-))$. In other words, B is the enumeration of answers of $\hat{A}_2$, for the given η and α^-. Set

$$L_B = \left\{ \alpha_m : |\bar{B} \cdot v(\alpha_m)| \geq \frac{2K}{2^{0.4k/2}} \right\}.$$

By Lemma 5, $|L_B| \leq 2^{cn+0.4k}$. Thus,

Lemma 7. *If $\alpha \in L_{B_{\eta,\alpha^-}}$, for some η, then $C(\alpha|\hat{A}_2) \leq (m-1)n + \frac{n}{5} + cn + 0.4k$.*

Proof. The sequence α is composed of α^- together with $\alpha^{(m)}$. Thus, given $\hat{A}_2$ the entire sequence α can be fully characterized by:

1. η - $n/5$ bits,
2. α^- - $(m-1)n$ bits, and
3. The index of α_m in $L_{B_{\eta,\alpha^-}}$ - $cn + 0.4k$ bits,

The construction of α from this information is analogous to that given in the proof of Corollary 6. □

Corollary 8 *For any A_1, A_2,*

$$\Pr\left[A_2(A_1(\alpha), s, X^-) = X_m\right] \le \frac{1}{2} + \frac{2}{2^{k/5}}.$$

Proof. By Lemma 7 and the definition of L_B, if $C(\alpha|\hat{A}_2) > (m-1)n + \frac{n}{5} + cn + 0.4k$, then, for any η

$$\Pr\left[\hat{A}_2(\eta, s, \alpha^-) = X_m)\right] < \frac{1}{2} + \frac{1}{2^{0.4k/2}}.$$

Thus, we also have that for such α and any algorithm A_2 (which gets X^- instead of α^- as input) and $\eta = A_1(\alpha)$,

$$\Pr\left[A_2(A_1(\alpha), s, X^-) = X_m)\right] < \frac{1}{2} + \frac{1}{2^{0.4k/2}}.$$

For $\alpha \in \{0,1\}^{nm}$,

$$\Pr\left[C(\alpha|\hat{A}_2) \le (m-1)n + \frac{n}{5} + cn + 0.4k\right] \ll \frac{1}{2^{0.4k/2}}.$$

Thus, in all,

$$\Pr\left[A_2(A_1(\alpha), s, X^-) = X_m\right] \le \frac{1}{2} + \frac{2}{2^{0.2k}}.$$

□

3.4 Multi-bit Security - Any Two Messages

We consider the case of distinguishing two messages. Let $M^{(0)}, M^{(1)}$ be two distinct messages, $M^{(i)} \in \{0,1\}^m$. We show that if an eavesdropper's algorithm A can distinguish between $M^{(0)}$ and $M^{(1)}$, with probability p, then there is another algorithm $B = (B_1, B_2)$, which solves the One-more Bit Problem with probability $\ge 1/2 + p/2$. Specifically, w.l.o.g. assume that

$$\Pr\left[A(M^{(1)}) = 1\right] - \Pr\left[A(M^{(0)}) = 1\right] = p \ge 0.$$

We construct $B(B_1, B_2)$, such that for $\beta \in \{0,1\}^{mn}$,

$$\Pr\left[B_2(B_1(\beta), s, X^-(\beta, s)) = s(\beta^{(m)})\right] \ge \frac{1}{2} + \frac{p}{2}.$$

Suppose that $M^{(i)} \ne 0$, for $i = 0, 1$. Let P be an $m \times m$ non-singular matrix over F_2 such that $P \cdot M^{(0)} = e_1$ and $P \cdot M^{(1)} = e_1 + e_m$, where e_1 and e_m are the unit vectors with a 1 in the first and last coordinates, respectively. For $\beta = (\beta^{(1)}, \ldots, \beta^{(m)})$ we view β as an $m \times n$ matrix. Thus, $P \cdot \beta$ is another $m \times n$ matrix. A detailed description of $B = (B_1, B_2)$, given $A = (A_1, A_2)$, is provided hereunder. For the case that $M^{(0)} = 0$, an analogous proof works, omitting e_1 in all its appearances.

```
B_1: Input: $\beta$. Output: $\eta$.
1    $\alpha := P^{-1}\beta$.
2    $\eta := A_1(\alpha)$.

B_2: Input: $\eta$, $s$ and $X^- = X^-(s,\beta)$. Output: $X_m = s(\beta^{(m)})$.
1    Choose $r \in \{0,1\}$ at random.
2    $X := X^- \circ 0$ (concatenation).
3    $Z = X \oplus e_1 \oplus re_m$.
4    $Y := P^{-1}Z$.
5    Output $A_2(\eta, s, Y) \oplus r$.
```

First we prove a technical lemma.

Lemma 9. *Suppose that*

$$\Pr\left[A(M^{(1)}) = 1\right] - \Pr\left[A(M^{(0)}) = 1\right] = p \geq 0.$$

Consider choosing $r \in \{0,1\}$ at random and then running A on $M^{(r)}$. Then

$$\Pr\left[A(M^{(r)}) = r\right] = \frac{1}{2} + \frac{p}{2}.$$

Proof.

$$\Pr\left[A(M^{(r)}) = r\right] =$$
$$\Pr\left[A(M^{(1)}) = 1\right]\Pr[r=1] + \Pr\left[A(M^{(0)}) = 0\right]\Pr[r=0] =$$
$$\frac{1}{2}\left(\Pr\left[A(M^{(1)}) = 1\right] + \left(1 - \Pr\left[A(M^{(0)}) = 1\right]\right)\right) = \frac{1}{2} + \frac{p}{2}.$$

□

Lemma 10.

$$\Pr\left[B_2(B_1(\beta), s, X^-(\beta, s)) = s(\beta^{(m)}))\right] = \frac{1}{2} + \frac{p}{2}.$$

Proof. Set $\delta = s(\beta^{(m)})$, and $\beta^- = (\beta^1, \ldots, \beta^{m-1})$. By construction $Z = \big(s(\beta^-) \circ 0\big) \oplus (e_1 \oplus re_m) = (s(\beta) \oplus \delta e_m) \oplus (e_1 \oplus re_m) = s(\beta) \oplus (e_1 \oplus (r \oplus \delta)e_m)$. Also, by construction, $\alpha = P^{-1}\beta$ (B_1 line 1), and $P^{-1}e_1 = M^{(0)}$ and $P^{-1}(e_1 \oplus e_m) = M^{(1)}$. Thus, $P^{-1}(e_1 \oplus ((r \oplus \delta)e_m)) = M^{(r\oplus\delta)}$. Thus,

$$Y = P^{-1}Z = P^{-1}(s(\beta) \oplus (e_1 \oplus (r \oplus \delta)e_m)) = s(P^{-1}\beta) \oplus P^{-1}(e_1 \oplus (r \oplus \delta)e_m)$$
$$= s(\alpha) \oplus M^{(r\oplus\delta)}.$$

By construction, $B_1(\beta) = A_1(\alpha)$, and $B_2(\eta, s, X^-(s,\beta)) = A_2(\eta, s, Y) \oplus r$. Thus,

$$\Pr\left[B_2(B_1(\beta), s, X^-(s,\beta)) = \delta\right] = \Pr\left[A_2(A_1(\alpha), s, s(\alpha) \oplus M^{(r\oplus\delta)}) = r \oplus \delta\right] =$$
$$\Pr\left[A(M^{(r\oplus\delta)}) = r \oplus \delta\right] \leq \frac{1}{2} + \frac{p}{2}.$$

□

Together with Corollary 8 we get

Corollary 11 *For any $M^{(0)}$ and $M^{(1)}$,*

$$\left|\Pr\left[A(M^{(1)})=1\right]-\Pr\left[A(M^{(0)})=1\right]\right|\leq\frac{1}{2^{k/5}}.$$

By Lemma 1 this completes the proof of Theorem 1.

References

1. Y. Aumann and U. Feige. One message proof systems with known space verifies. In D.P. Stinson, editor, *Advances in Cryptology*, pages 85–99. Springer–Verlag, 1993.
2. C. H. Bennett, G. Brassard, C. Crepeau, and U. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41(6), 1995.
3. C. Cachin. *Entropy Measures and Unconditional Security in Cryptography*, volume 1. Hartung-Gorre Verlag, Konstaz, Germany, 1997.
4. C. Cachin and U. M. Maurer. Unconditional security against memory bounded adversaries. In *Proceedings of Crypto '97*, 1997.
5. A. De-Santis, G. Persiano, and M. Yung. One-message statistical zero-knowledge proofs with space-bounded verifier. In *Proceedings of the 19th ICALP*, 1992.
6. M. Li and P. M. B. Vitanyi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer-Verlag, New York, 2nd edition, 1997.
7. U. M. Maurer. Conditionally-perfect secrecy and a provably-secure randomized cipher. *Journal of Cryptology*, 5:53–66, 1992.
8. U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Transactions on Information Theory*, 39:733–742, 1993.
9. R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, 1995.
10. C. E. Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28:656–715, 1949.