

On the Impossibility of Private Key Cryptography with Weakly Random Keys

James L. McInnes*
University of Toronto

Benny Pinkas†
Technion — Israel Institute of Technology

Abstract

The properties of weak sources of randomness have been investigated in many contexts and using several models of weakly random behaviour. For two such models, developed by Santha and Vazirani, and Chor and Goldreich, it is known that the output from one such source cannot be “compressed” to produce nearly random bits. At the same time, however, a single source is sufficient to solve problems in the randomized complexity classes BPP and RP . It is natural to ask exactly which tasks can be done using a single, weak source of randomness and which cannot. The present work begins to answer this question by establishing that a single weakly random source of either model cannot be used to obtain a secure “one-time-pad” type of cryptosystem.

1 Introduction

Secret transmission of information over insecure communication lines is a major issue in cryptography. In the classical setting, two parties A and B , share a secret, private key K . A wishes to send a plaintext message M , to B . A encrypts M using K , and sends the resulting ciphertext C , to B . A listener L can eavesdrop on the communication line and find C (but not alter it). In addition L knows the functions employed by A and B . The goal of the cryptosystem is to enable B to correctly decrypt M , while retaining security against the listener.

In order to operate, the parties A and B need an access to a joint source of randomness. Without such a source, L possesses the same information as B does. As L knows B 's program, B has no advantage over L , and so such a cryptosystem will not be secure.

If A and B share a perfect source of unbiased independent random bits, then they can use this source to generate the private key K , and use this key as a *one-*

*Address: Dept. of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 1A4
Email: jimm@theory.toronto.edu

†Address: Dept. of Computer Science, Technion — Israel Institute of Technology, Haifa 32000
Israel.

Email: bennyp@techunix.bitnet. Research supported in part by US-Israel BSF grant 88-00282.

time pad, achieving maximum security [S]. In practice, however, it seems unrealistic to expect a source to be perfectly random. Most physical sources, such as Zener diodes or Geiger counters, are imperfect; that is, they do not output a uniform distribution. The question that motivates this work is whether there is a secure private key cryptosystem (analogous to one-time pad) if a weaker, non-perfect source of randomness is shared by A and B .

Two widely investigated general models for weak-sources of randomness have been suggested by Santha and Vazirani [SV], and Chor and Goldreich [CG]. They are known as the *SV-source* and *PRB-source*, respectively. The sources they describe maintain some amount of randomness but allow the value of each bit output to depend on the values of all previous bits. For both models, it has been shown in [SV, CG] that a single SV or PRB source cannot be used to extract “almost” unbiased random bits.

We show that a private key cryptosystem in which both parties share a private key, generated by a weak source of randomness, and have no access to any other source of randomness, is not secure. It should be stressed that this is not an immediate corollary of the fact that random bits cannot be extracted from such sources. We also show a *secure* cryptosystem, where the parties share a slightly-random private key, and have access to a public source of perfect randomness.

The paper is organized as follows: sections 2 and 3 give the needed background and definitions in weak-sources of randomness and cryptography. In sections 4 and 5 it is shown that a crypto-system in which A and B share a private-key generated by a PRB or SV source, and have no additional sources of randomness, is not secure. In section 6 we allow them to use a public source of truly random bits in addition to the slightly-random key, and show that then it is possible to combine these two sources to produce a secure one-time pad.

2 Weak Sources of Randomness

Physical sources of randomness are imperfect, that is, they do not output a uniform distribution. Several mathematical models of such sources have been investigated. Von Neumann [N] considered a source which generates a sequence of independent tosses of a coin with a fixed but unknown bias, and suggested a method to extract perfect random bits from it. Blum [B] modeled weak randomness as a finite state Markov chain (with unknown transition probabilities). This model allows each output bit to depend on the previous c bits (for any fixed c). Blum gave an algorithm to extract perfect random bits from such a source.

Next we describe two more recent models, in which each output bit can depend on all previous bits.

2.1 SV-model

Santha and Vazirani [SV] suggested a model (hereafter referred to as the *SV-model*) where each bit in the output sequence is 0 with a probability of at least δ and not

more than $1 - \delta$ (where $0 \leq \delta \leq \frac{1}{2}$ is fixed). The probability that a given n -bit string is output is hence bounded above by $(1 - \delta)^n$, and below by δ^n (therefore each bit sequence is output with some positive probability). A source with $\delta = \frac{1}{2}$ is a perfect random source, and one with $\delta = 0$ can have no randomness at all. This model allows each bit to depend on all previous bits.

2.2 PRB-model

A different model was suggested by Chor and Goldreich [CG], where instead of bounding the probability of each individual bit, they bound the probability that any given string will appear in any position. Such a source is called a *PRobability-Bounded source*, or a *PRB-source*. It has two parameters, l and b . A source is said to be an (l, b) -source if for every prefix $\alpha \in \{0, 1\}^*$ of the output sequence, and every l -bit string β , the conditional probability that the next l bits will equal β (given the prefix α) is at most 2^{-b} . Thus an (l, l) -source is a perfect random source, and an $(l, 0)$ -source can have no randomness at all.

The PRB-model is a strict generalization of the SV-model. Any SV-source with parameter δ is a $(1, \log_2(1 - \delta)^{-1})$ PRB-source. The inclusion is proper since a probability bounded source may output some strings with probability 0. For example, a $(2, 1)$ PRB-source which outputs 11 with probability $\frac{1}{2}$ and 10 with probability $\frac{1}{2}$, cannot be modeled by any SV-source.

2.3 Known Results

It was shown in [SV, CG] that a single SV or PRB source cannot be used to extract “almost” unbiased random bits. On the other hand, in both models, two independent sources suffice for this purpose [V, CG]. A more surprising result is that *BPP* and *RP* problems can be efficiently solved using the output of a single SV or PRB source [VV, CG]. This indicates that some useful randomness can be extracted from these sources, and it leads us to ask how useful a slightly random source would be for cryptography.

Given a boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, define the *density of zeroes* of it to be

$$d = \frac{|\{y \in \{0, 1\}^n \mid f(y) = 0\}|}{2^n}$$

The following technique is due to Gerek and is given in [SV]:

Lemma 1 *for all $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $0 \leq \delta \leq \frac{1}{2}$, and $d \geq \frac{1}{2}$, there is a strategy to set a SV-source with parameter δ , such that if x is an n bit string generated by this source, and the density of zeroes of f is d , then*

$$\text{Prob}(f(x) = 1) \geq 2(1 - \delta)(1 - d)$$

Any extraction scheme to extract a random bit from an n -bit string generated by an SV-source, can be viewed as a function $f : \{0, 1\}^n \rightarrow \{0, 1\}$. The lemma implies that no such scheme can extract a bit with less than $1 - \delta$ bias, from a SV-source with

parameter δ . (More precisely, for any scheme f there is a δ -source which causes the outcome of the scheme to be at least $1 - \delta$ biased.)

3 Cryptographic Background

3.1 Cryptosystems

A simple cryptographic system (A, B) is composed of two communicating parties, A and B , communicating over an insecure channel. Their goal is for A to pass a secret message (or *plaintext*), M , to B . In a *private key cryptosystem*, A and B share a randomly chosen string, K , which is the *private key*. A sends to B a message C (*ciphertext*), which is a deterministic function of M and K . A *listener*, L , is able to examine the communication between A and B . He knows the functions they use, but not the private key K . The listener is passive, he cannot interject anything of his own on the line. The listener attempts to find the plaintext, or at least to extract as much information about it as possible. Figure 1 sums up the scenario described here.




Figure 1: A simple private key cryptographic system.

The simplest scenario is of a one-bit cryptosystem, namely one where the plaintext is composed of a single bit. For a one-bit cryptosystem (A, B) , we define the following requirements:

Correctness: (A, B) is *correct* if, given that b is chosen randomly in $\{0, 1\}$,

$$\text{Prob}(B \text{ outputs } b \text{ correctly}) \geq 1 - o(1)$$

for all k , and sufficiently large n .

Security: (A, B) is *secure* if, for every listener L , and for a bit b that is chosen randomly in $\{0, 1\}$,

$$\text{Prob}(L \text{ outputs the same as } B) \leq \frac{1}{2} + o(1)$$

for all k , and sufficiently large n .

3.2 Cryptographic Setup for Our Problem

Given a source of true randomness, there exists a simple cryptosystem which is secure and correct (the *one-time pad* [S]). In this system an n -bit randomly chosen private key is used to communicate an n -bit plaintext. The ciphertext in this case is just the bitwise exclusive-or of the key with the plaintext. We wish to know whether there is a secure cryptosystem analogous to one-time pad, in which A and B share a key generated by a weak source of randomness. Such a system would be secure regardless of any complexity assumptions we might make, and therefore the computational power of L is not limited (and hence L can be restricted to be deterministic). For convenience we do not limit the power of A and B , as well. (And since we show the impossibility of such a system, these assumptions on A and B do not weaken our results). However, we do not allow A and B to use a source of truly random bits in their computation. Any random bits they need for the computations they make (as probabilistic algorithms) should be taken from the slightly-random key.

For our purposes, we can make the following simplifying assumptions about the cryptosystem (A, B) :

- The communication is one-way. Namely, B only receives communication from A , and sends no messages himself. This is possible since A and B share the same source of randomness, and so B does not have any input that A does not have.
- For a given key length n , A always sends the same number m_n of bits to B .

It can be seen that there is no loss of generality incurred in making these assumptions. This reduces a cryptosystem (A, B) to the following:

1. A slightly random key K of length n .
2. An encryption function $f : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}^{m_n}$
3. A decryption function $g : \{0, 1\}^{m_n} \times \{0, 1\} \rightarrow \{0, 1\}$.

Remarks:

1. The plaintext b is uniformly distributed in $\{0, 1\}$.
2. The system operates as follows: A sends the ciphertext $C = f(K, b)$ to B , and B computes $b' = g(K, C)$. As B is allowed to err (with small probability), b' might be different from b .
3. A listener L is specified by an eavesdropping function $h : \{0, 1\}^{m_n} \rightarrow \{0, 1\}$ that he uses, given the ciphertext, to try and retrieve the plaintext.
4. The listener knows f and g . He also knows the strategy of the weakly random source, namely the a priori probability distribution on the private key K .

5. Other than the weakly random source, there is no other source of randomness in the system.

3.3 A Matrix Representation

We represent the decryption function $g : \{0,1\}^n \times \{0,1\}^{mn} \rightarrow \{0,1\}$ by a matrix D having 2^n rows (one for each possible key) and 2^{mn} columns. (one for each possible value of $f(K,b)$). The entries of D are either 0,1 or blank, and each row contains two non-blank entries. Each entry in D is simply $g(K, f(K,b))$ for K and $f(K,b)$ corresponding to that row and column respectively. Two entries need to appear in each row, since, for each K , there are only two possible values of $f(K,b)$ that occur, since A behaves deterministically.

A listener is completely specified by giving an output value in $\{0,1\}$ for each value of $f(K,b)$. In matrix terms this corresponds to giving a $\{0,1\}$ -labelling of the columns of D .

We are frequently interested in the extent to which a listener's output matches that of B . This corresponds to asking if a given $\{0,1\}$ -labelling of the columns of D (i.e. a given listener) matches the particular row that corresponds to the key K . Since each row has at most two entries there are three possibilities. A row *agrees* with a labelling l if all the entries in the row match the label that l gives for the column in which they appear, and the row *disagrees* if none of the entries matches l . A row *half-agrees* with l if it has two entries, one which matches l and one which doesn't. Let the *weight* of a row with respect to a particular labelling be the probability (taken over the messages 0 and 1) that L outputs the same as B if that row is chosen. Thus a row which agrees with l has weight 1, a row which half agrees has weight $1/2$, and a row which disagrees has weight 0.

4 Using a PRB-source to Communicate a Single Bit

We first consider the case where A and B share an n -bit slightly-random key generated by a PRB-source. We show that no cryptosystem is secure when its only source of randomness is an $(n, n - c)$ PRB source ($c > 0$). Theorem 1 establishes this result quantitatively, by showing that for every value of c , there is a listener whose probability of finding the transmitted bit is higher than a constant which depends only on c (and not on n). We also demonstrate a cryptosystem achieving the lower bound of Theorem 1.

Theorem 1 *If (A, B) is a cryptosystem such that*

1. *A and B share an n -bit private key K .*
2. *A and B have no additional sources of randomness.*

3. *A sends the encryption of a single bit message b to B , and B tries to decrypt this and find b .*

then for every $0 \leq c \leq n$ there exist an $(n, n - c)$ source S and a listener L such that if K is generated by S then

$$\text{Prob}(L \text{ outputs the same as } B) \geq \begin{cases} 1 & \text{for } 2 \leq c \leq n \\ \frac{1}{2} + \frac{2^c}{8} & \text{for } 2 - \log_2 3 \leq c \leq 2 \\ \frac{2^c}{2} & \text{for } 0 \leq c \leq 2 - \log_2 3 \end{cases}$$

Proof. We may think of an $(n, n - c)$ source S as selecting a fraction 2^{-c} of the keys, each of which then occurs with probability $1/2^{n-c}$. The remaining keys do not occur. Denote by e the value 2^{-c} . In the combinatorial setting established in the previous section, this corresponds to choosing a set of e of the rows of D . (without loss of generality, assume that $e2^n$ is an integer). From now on we speak of choosing rows of D instead of choosing keys.

Given a labelling l of the columns of D , we say that its rows satisfy $(d_1, 1 - (d_1 + d_0), d_0)$ (where $0 \leq d_0, d_1$ and $d_0 + d_1 \leq 1$), if a fraction d_1 of the rows have weight 1, d_0 of the rows have weight 0, and $1 - (d_1 + d_0)$ of the rows have weight $\frac{1}{2}$.

The strategy of the $(n, n - c)$ source S we use, is to choose exactly a part e of the 2^n rows and give each one the same probability, $\frac{1}{e2^n}$. The source first chooses 1-rows. If there are less than $e2^n$ of these, it also takes $\frac{1}{2}$ -rows; and only if there are not enough of these, too, it takes 0-rows. The following observations can be made:

1. Without loss of generality, $d_1 \geq d_0$ (otherwise consider \bar{l} , the labelling obtained by reversing the labels of l).
2. $d_1 = d_0$. That is because for any value of e , with a $(d_1, 1 - (d_1 + d_0), d_0)$ matrix (where $d_1 \geq d_0$), the source has more positive weight rows than with a $(d_1, 1 - 2d_1, d_1)$ matrix. Therefore the listener can achieve better success on the former matrix.
3. The following lemma implies that there is a labelling l that satisfies $(d_1, 1 - (d_1 + d_0), d_0)$, with $d_1 \geq \frac{1}{4}$.

Lemma 2 *There is a $\{0, 1\}$ -labelling of the columns of D which agrees with at least $1/4$ of the rows.*

Proof (of lemma 2). D has 2^n rows and 2^{m_n} or $2^{2^{m_n}}$ different labellings of the columns of D . Each row agrees with at least $1/4$ of the labellings, since at most two columns have entries, so there are $\frac{2^n \cdot 2^{2^{m_n}}}{4} = 2^{2^{m_n} + n - 2}$ agreements among $2^{2^{m_n}}$ labellings. By the Pigeon Hole Principle there is some labelling which agrees with at least

$$\frac{2^{2^{m_n} + n - 2}}{2^{2^{m_n}}} = 2^{n-2}$$

of the rows. □

Let us now consider the following cases:

- $\square\square\square\square\square\square$ The source chooses only 1-rows. Therefore the listener has probability 1 of guessing the bit b that is being communicated.

- $\square\square\square\square\square\square$

The source doesn't have to take 0-rows, it has enough positive weight rows to choose from. The average weight achieved is then

$$\frac{d_1 + \frac{1}{2}(e - d_1)}{e} = \frac{1}{2} + \frac{d_1}{2e} \geq \frac{1}{2} + \frac{1}{8e}$$

- $\square\square\square\square\square\square$

There are two possible cases, depending on the value of d_0 :

1. $e \leq 1 - d_0$. As in the former case, the source has enough positive weight rows to choose from. The average weight achieved is then $\frac{1}{2} + \frac{1}{8e}$.
2. $e > 1 - d_0$. All positive weight rows are used by the source, as well as some 0-weight rows. The success is at least

$$\frac{d_1 + \frac{1}{2}(1 - 2d_1)}{e} = \frac{1}{2e}$$

For $e \leq 0.75$, the first bound, $\frac{1}{2} + \frac{1}{8e}$, is lower than the second one, $\frac{1}{2e}$. The resulting lower bound is thus

$$\frac{1}{2} + \frac{1}{8e}$$

- $\square\square\square\square\square$

All positive weight rows, as well as some 0-weight rows, are used, and the value achieved is

$$\frac{d_1 + \frac{1}{2}(1 - 2d_1)}{e} = \frac{1}{2e}$$

This completes the proof of the theorem. □

Interesting cases of the above theorem are for an $(n, n - 1)$ source, where we get a lower bound of 0.75, and for $(n, n - c)$ sources ($c \geq 2$), where the lower bound is 1.

To show that the results of theorem 1 cannot be improved, it is sufficient to exhibit a matrix D for which every labelling agrees with exactly one quarter of the rows. One such matrix is as follows, for a key length of $n = 3$ and with $m_n = 2$.

$$\begin{pmatrix} 1 & 0 & \\ 1 & 1 & \\ 0 & 0 & \\ 0 & 1 & \\ & 1 & 1 \\ & 0 & 1 \\ & & 0 & 0 \\ & & 1 & 0 \end{pmatrix}$$

It can be verified by inspection that every labelling of the columns of D agrees with exactly two rows. It is worth noting that this matrix contains rows which consist of only 0 entries (or 1 entries), and therefore B cannot always decrypt correctly. When one of these rows is selected, B will always output 0 (1), regardless of the value of b , and so will be correct only half of the time. For cryptosystems in which B always decrypts b correctly, we have the following theorem:

Theorem 2 *If (A, B) is as in theorem 1, and in addition B always outputs b correctly, then for every $0 \leq c \leq n$ there exist a $(n, n - c)$ source S and a listener L , such that if the key K is generated by S then*

$$\text{Prob}(L \text{ outputs the same as } B) > \begin{cases} 1 & \text{for } 2 \leq c \leq n \\ \frac{1}{2} + \frac{0.125}{2^{-c}} & \text{for } -\log_2 \frac{3}{4} \leq c \leq 2 \\ \frac{0.5}{2^{-c}} & \text{for } 0 \leq c \leq -\log_2 \frac{3}{4} \end{cases}$$

The only difference from theorem 1 is that here the inequality is strict.

Proof. The proof is the same as in theorem 1, except that in the proof of lemma 2 we now disregard the all-zeroes and all-ones labellings, which here agree with no rows at all. Therefore we get that there is a labelling that agrees with *more* than $\frac{1}{4}$ of the rows. \square

Theorem 2 cannot be improved upon. In [M] there is an example of a series of matrices D_1, D_2, \dots such that D_n represents a cryptosystem in which B always decrypts correctly, and the largest group of rows in agreement with every fixed labelling, is not greater than a $\frac{1}{4} + o(1)$ fraction of the rows.

5 Using an SV-source to Communicate a Single Bit

In this section we show that private key cryptography with private keys generated by an SV-source is not possible. Any correct system (A, B) is not secure. Namely, for any encryption and decryption functions, f and g , there is a δ -source and an eavesdropping function h , such that $\text{Prob}(h(C) = g(K, C)) \geq \frac{1}{2} + p(\delta)$, where p depends on δ (but not on n , the length of the key K , or on m_n , the length of the ciphertext).

We emphasize that this result does not follow from the corresponding one on PRB-sources. The source which was used there to show that correct and secure cryptosystems do not exist, was not a SV-source. On the other hand, the quantitative result here is not as strong: Here we only guarantee an advantage over $\frac{1}{2}$, $p(\delta)$, while there was a constant advantage of $\frac{1}{4}$ for $(l, l - 1)$ -sources, and complete certainty in successful eavesdropping for $(l, l - 2)$ -sources.

5.1 Preliminaries

Let us now consider the case where the private key K is an n bit binary string generated by a SV-source with parameter δ . We may think of it as generated by an adversary who chooses, after each bit is output, a probability between δ and $1 - \delta$ that the next bit will be a 0. The adversary strategy can be represented as a complete binary tree of height n , where each left branch corresponds to a 0 being chosen, and a right branch to a 1 being chosen. The leaves correspond to n bit strings, sorted in dictionary order. The adversary chooses at each node, with which probability (between δ and $1 - \delta$) to continue to the left and right branches. This induces a probability distribution on the leaves.

For our purposes, each leaf corresponds to a different key K , of length n , and hence to a different row of D . A $\{0, 1\}$ -labelling l of the columns of D uniquely assigns a weight to each leaf. This weight is 1 if the leaf *agrees* with l , $\frac{1}{2}$ if it *half-agrees*, and 0 if it *disagrees*. To describe a combination of a matrix D (cryptosystem) and a $\{0, 1\}$ -labelling l (listener), we define an *adversary tree of height n* as a complete binary tree of height n , the leaves of which are labelled with weights from $\{0, \frac{1}{2}, 1\}$.

Let W_n denote the set of all adversary trees of height n . Each decryption matrix D and labelling l determine a particular tree $T \in W_n$. For this tree, T , the optimal strategy of the adversary is to label the edges of T with probabilities (between δ and $1 - \delta$) such as to maximize the the expected value of the leaf reached.

For a node v , define $val_\delta(v)$ as:

- its weight, if v is a leaf.
- if v has descendants v_1 and v_2 , $val_\delta(v) = (1 - \delta) \max(val_\delta(v_1), val_\delta(v_2)) + \delta \min(val_\delta(v_1), val_\delta(v_2))$

For a tree T , define $val_\delta(T)$ as $val_\delta(r)$, where r is the root of T .

Lemma 3 *For a given adversary tree, the source strategy which labels the branch leading to the “heavier” descendant with $1 - \delta$, and the other branch with δ , maximizes the expected value of the leaf reached. The expected value of the leaf reached is then $val_\delta(T)$.*

An adversary tree T is called *balanced* if

- at least one quarter of its leaves are 1-leaves.
- it has the same number of 0-leaves as 1-leaves.

Let Y_n denote the set of all balanced trees of height n .

Lemma 4 *For each decryption matrix D there is a $\{0, 1\}$ -labelling l , such that for the adversary tree $T \in W_n$, corresponding to D and l , the following holds:*

1. *The sum of the weights of all 2^n leaves is not less than 2^{n-1} .*
2. *At least one quarter of the leaves (2^{n-2} leaves) are 1-leaves.*
3. *It has value greater than or equal to some balanced tree.*

Proof: by a counting argument.

Therefore, instead of proving a lower bound on the amount of agreement between a decryption matrix and the optimal $\{0, 1\}$ -labelling of it (that is, the labelling giving a maximal value to $val_\delta(T)$), it suffices to prove a lower bound for the expected value achieved on a *balanced tree* of height n .

Conjecture 1 For every balanced tree $T \in Y_n$, it holds that

$$val_\delta(T) \geq \frac{1}{2} + \left(\frac{1}{2} - \frac{3}{2}\delta + \delta^2\right)$$

This lower bound on balanced trees cannot be improved. It equals an upper bound, which is the value achieved on a tree (see figure 2) in which one main subtree (i.e. a subtree descending from the root) has only $\frac{1}{2}$ -leaves, while the other main subtree has one subtree the leaves of which are all labeled 1, and another subtree the leaves of which are all labeled 0.




Figure 2: A tree achieving the lower bound.

Although the lower bound we conjecture here is the best possible bound for adversary trees, it turns out to be difficult to find a cryptosystem for which the maximum success rate a listener can achieve equals this bound. A cryptosystem for which no listener can guess the transmitted bit with more than $1 - \delta$ success, is given in [M].

5.2 Reducing $\{0, \frac{1}{2}, 1\}$ -trees to $\{0, 1\}$ -trees

We now give a reduction from the problem of finding a lower bound for the value achieved on an adversary tree, to finding a lower bound for the value achieved on a tree that has leaves with weights from $\{0, 1\}$.

Let B_n denote the set of all trees of height n , the leaves of which have weights from $\{0, 1\}$. Define $f_{n,m} : W_n \rightarrow B_{n+m}$ as the transformation which, given an adversary tree $T \in W_n$, does the following:

- replaces each 0-leaf with a subtree of height m , all of its leaves being 0-leaves.
- replaces each 1-leaf with a subtree of height m , all of its leaves being 1-leaves.

- replaces each $\frac{1}{2}$ -leaf with a subtree R , such that

- $R \in B_m$
- $val_\delta(R) \leq \frac{1}{2}$
- $val_\delta(R) = \max\{val_\delta(R') \mid R' \in B_m, val_\delta(R') \leq \frac{1}{2}\}$

That is, among all trees of B_m that have value not greater than $\frac{1}{2}$, the maximum value is achieved on R . (if there are several trees matching this definition, we take the first in some lexicographical order).

It is clear that for $T \in W_n$, any lower bound on $val_\delta(f_{n,m}(T))$ implies the same lower bound on $val_\delta(T)$.

5.3 Analysis

We now describe a method for constructing trees R , which is useful for values of δ close to $\frac{1}{2}$. Specifically, we construct trees for a sequence $\{\delta_m\}_{m=2}^\infty$, so that $\delta_{m+1} > \delta_m$, and δ_m converges to $\frac{1}{2}$. Let $T \in Y_n$ be a balanced tree, with $\frac{1}{4}$ of its leaves being 1-leaves. (Hence half of its leaves are $\frac{1}{2}$ -leaves, and the rest are 0-leaves).

Let us choose R_m as the tree of height m , of which the $2^{m-1} - 1$ left leaves are 1-leaves, all other leaves are 0-leaves.

$$val_\delta(R_m) = (1 - \delta) - (1 - \delta)\delta^{m-1} = (1 - \delta)(1 - \delta^{m-1})$$

This value should not be greater than $\frac{1}{2}$. We define δ_m by the equation

$$val_{\delta_m}(R_m) = (1 - \delta_m)(1 - (\delta_m)^{m-1}) = \frac{1}{2} \tag{1}$$

Let T' denote $f_{n,m}(T)$. The density of 0-s in R_m is $d_{R_m} = \frac{1}{2} + 2^{-m}$, and in T' the density is $d_{T'} = \frac{1}{4} + \frac{1}{2}(\frac{1}{2} + 2^{-m}) = \frac{1}{2} + 2^{-m-1}$. Applying Geréb's bound (Lemma 1), we get the following lower bound for T' :

$$val_{\delta_m}(T) \geq 2(1 - \delta_m)(1 - d_{T'}) = (1 - \delta_m)(1 - 2^{-m}) = 1 - \delta_m - (1 - \delta_m)2^{-m} \tag{2}$$

For values of δ in the sequence $\{\delta_m\}$, it is easy, using equation (1), to describe m as a function of δ_m . Thus we get for these values a lower bound of

$$val_{\delta_m}(T) \geq 1 - \delta - \frac{1 - \delta}{2} \left(\frac{2 - 2\delta}{1 - 2\delta} \right)^{\frac{1}{\log \delta}} = \frac{1}{2} + \left(\frac{1}{2} - \delta - \frac{1 - \delta}{2} \left(\frac{2 - 2\delta}{1 - 2\delta} \right)^{\frac{1}{\log \delta}} \right) \stackrel{\text{def}}{=} \frac{1}{2} + p(\delta_m) \tag{3}$$

$p(m)$ is the advantage over $\frac{1}{2}$, of this lower bound.

5.4

Let $g(\delta)$ be the advantage over $\frac{1}{2}$ stated in conjecture 1, that is, $g(\delta) = \frac{1}{2} - \frac{3}{2}\delta + \delta^2$. The following theorem gives a constant bound for the ratio between the advantages $g(\delta)$ and $p(\delta)$.

Theorem 3 *For a crypto system (A, B) , where A and B share an n -bit private key K , which is generated by a SV -source with parameter δ , $\delta \geq 0.45$, there exists a SV -source S and a listener L , such that if K is generated by S then*

$$\text{Prob}(L \text{ outputs the same as } B) \geq \frac{1}{2} + \frac{g(\delta)}{2.76}$$

Proof: Lemma 4 reduces this problem to finding a lower bound on $\text{val}_\delta(T)$, where T is *balanced*, and $\frac{1}{4}$ of its leaves are 1-leaves. The theorem follows from bound (3) and from showing that

$$\frac{g(\delta)}{p(\delta)} \leq 2.76$$

for all δ -s in $[0.45, 0.5]$.

It is easy to show that this ratio is less than 2.60 for all δ_m ($m \geq 3$).

To prove a similar result for all values of δ , note that for $\delta_{m-1} < \delta \leq \delta_m$, a δ_m -source is also a δ -source. Hence lower bound (3) is the same for all δ -s in $(\delta_{m-1}, \delta_m]$.

We should compare the advantages above $\frac{1}{2}$ of the *conjectured* bound with δ_{m-1} , and *our* bound with δ_m . In order to find δ_m , it is needed to solve equation (1). Difficulties arise in solving it analytically (for a general m), and so we needed to solve it using numerical methods. The theorem follows from calculating the ratio between the advantages of the conjectured bound with δ_4 , and our bound with δ_5 (this ratio is a little less than 2.76), and showing that for values of δ larger than δ_4 , this ratio is not higher than that. The same method can be applied to give a bound for all $\delta \geq \delta_3 = 0.404$ and a ratio of 3.69 instead of 2.76, and for all $\delta \geq \delta_2 = 0.293$ where the ratio is 6.56 instead of 2.76. \square

The lower bound we got is not trivial, yet it is not optimal. The technique we employed is rather coarse since it only uses the density of 0-leaves, and not their location. Furthermore, we used a bound for $\{0, 1\}$ -trees (in Lemma 1) which is also not optimal.

6 Allowing an Additional Source of Randomness

In the previous two sections we have assumed that the two parties A and B have no additional sources of randomness, public or private. In this section we show to what extent these results depend on this assumption by introducing a public source of truly random bits. By *public* we mean that any truly random bits used by A or B are known to the listener, and in the cryptosystems presented in this section this is made explicit by including any truly random strings used as part of the ciphertext. As a practical matter, a public, truly random source of bits could be something like a satellite using

the background radiation left over from the "big bang" as a source of entropy. Mainly, however, this situation is of interest to us from the point of view of investigating the mathematical relationship between weak sources of randomness and cryptography. It seems at first that such a source would not be helpful, but it turns out that in this new situation secure cryptography is possible, and we present a secure system. The complete proofs of the theorems in this section are not given here. The interested reader can find them in [M].

6.1 A Secure System

The following cryptosystem (A, B) is suggested by the construction used initially by Vazirani and Vazirani [VV] and subsequently by Chor and Goldreich [CG] to show that BPP and RP algorithms can be modified to work with just one slightly-random source.

1. (A, B) has a *security parameter* denoted $n \in \mathbf{Z}$.
2. A and B share an n -bit key K generated by an (n, b) -source S , where $K = k_1 \dots k_n$ for $k_i \in \{0, 1\}$.
3. A wishes to send $m \in \{0, 1\}$ to B .
4. A generates random X such that $|X| = n$.
5. A computes $p \in \{0, 1\}$ (for pad) by $p = \vec{X} \cdot \vec{K}$ (the inner product function).
6. A sends $p \oplus m$ and X to B .
7. B is then able to compute $p = \vec{X} \cdot \vec{K}$ and $m = p \oplus (m \oplus p)$.

It is easy to see that this system is correct. It is also secure.

Theorem 4 *Suppose that (A, B) is as above and that the private key K is chosen from an (n, b) -source. If m is randomly chosen in $\{0, 1\}$, then for each listener L which outputs a guess at m ,*

$$\text{Prob}(L \text{ outputs } m) \leq 1/2 + 6 \cdot 2^{-b/4}$$

Proof. The proof will appear in the final version. It depends on definitions and results in [CG] in a critical way. See [M] for details. \square

Theorem 4 is easily extended to include SV sources.

Corollary 1 *If (A, B) is modified so that K is generated by an SV source with parameter δ , then Theorem 4 remains true for $b = n \log_2(1 - \delta)^{-1}$.*

There is a natural extension of (A, B) that communicates many bits by running the system many times in parallel. This system can also be proven secure, using an appropriate definition of security for many-bit systems. For details the reader is again referred to [M].

7 Acknowledgments

The authors wish to thank Benny Chor, Michael Luby and Charles Rackoff for many helpful discussions about this work.

References

- [B] M. Blum, "Independent unbiased coin flips from a correlated biased source: A finite state Markov chain", *25th IEEE Sympos. Found. of Comput. Sci.*, pp. 425-433. 1984.
- [CG] B. Chor and O. Goldreich, "Unbiased bits from weak sources of randomness and probabilistic communication complexity", *SIAM J. Comput.*, Vol. 17, No. 2, pp. 230-261. April 1988.
- [M] J. L. McInnes, "*Cryptography using weak sources of randomness*", Technical Report 194/87, Dept. of Computer Science, University of Toronto. 1987.
- [N] J. von Neumann, "Various techniques used in connection with random digits" (notes by G. E. Forsythe), Applied Math Series, Vol. 12, pp.36-38, 1951; National Bureau of Standards, Washington D.C., reprinted in "*Collected Works*", Vol. 5, pp. 768-770, Pergamon, New York, 1963.
- [S] C. E. Shannon, "Communication theory of secrecy systems", *Bell Sys. Tech. J.*, **28**, pp. 656-715. 1949.
- [SV] M. Santha and U. V. Vazirani, "Generating quasi-random sequences from semi-random sources", *J. Comput. System Sci.*, **33**, pp. 75-87. 1986.
- [V] U. V. Vazirani, "Towards a strong communication complexity theory or generating quasi-random sequences from two communicating slightly-random sources", *Proc. 17th Annual Symposium on Theory of Computing*, pp. 366-378. 1985.
- [VV] U. V. Vazirani and V. V. Vazirani, "Random polynomial time is equal to slightly random polynomial time", *Proc. 26th Annual Symposium of Foundations of Computer Science*, pp. 417-428. 1985.