SANDIA REPORT

SAND2005-6979
Unlimited Release
Printed November 2005

Key Management and Encryption under the Bounded Storage Model

Erik Anderson, Tim Draelos, Andy Lanzone, and William Neumann

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.doe.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

SAND2005-6979
Unlimited Release
Printed November 2005

Key Management and Encryption under the Bounded Storage Model

Erik Anderson, Tim Draelos, Andy Lanzone, and William Neumann
Cryptography and Information Systems Surety Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0785

Abstract

There are several engineering obstacles that need to be solved before key management and encryption under the bounded storage model can be realized. One of the critical obstacles hindering its adoption is the construction of a scheme that achieves reliable communication in the event that timing synchronization errors occur. One of the main accomplishments of this project was the development of a new scheme that solves this problem. We show in general that there exist message encoding techniques under the bounded storage model that provide an arbitrarily small probability of transmission error. We compute the maximum capacity of this channel using the unsynchronized key-expansion as side-channel information at the decoder and provide tight lower bounds for a particular class of key-expansion functions that are pseudo-invariant to timing errors. Using our results in combination with Dziembowski et al. [11] encryption scheme we can construct a scheme that solves the timing synchronization error problem.

In addition to this work we conducted a detailed case study of current and future storage technologies. We analyzed the cost, capacity, and storage data rate of various technologies, so that precise security parameters can be developed for bounded storage encryption schemes. This will provide an invaluable tool for developing these schemes in practice.

Contents

List of Figures

Chapter 1

Introduction

The field of Cryptography has made significant strides in the development of its theoretical infrastructure over the past few decades. The invention of new theoretical tools has paved the way for advanced cryptographic systems that were not conceivable only a few years ago. Unfortunately, even with all the major achievements theoretical cryptography has enjoyed, the security of many of today's cryptographic schemes are based on unproven conjectures. It is unlikely given the current state of the field that any of these conjectures will be proven in the near future.

Many of today's modern cryptosystems are developed with specific types of security in mind. Depending on the type of cryptography being employed, the security of such schemes is often contingent on the degree of provable security one is after. In public key cryptography, security protocols is developed based on the intractability of difficult computational problems. Often these problems are based on mathematical puzzles that have been around for centuries, such as integer factorization or discrete logarithms. The security of these schemes remains in question however, as the difficulty of the puzzles they are based on has not been proven. If a significant breakthrough in computational power or algorithm development takes place, then many of these schemes may be at risk.

In the symmetric key arena the same types of problems exist, albeit in a slightly different nature. For example the security of most symmetric key encryption schemes can be reduced to the pseudorandomness of the block-cipher or function used. The existence of an efficient pseudorandom permutation or function is equivalent to the existence of one-way functions, which still remains an unproven conjecture. Therefore the security of such schemes in practice depends on the security of the functions used in their place, such as the Advanced Encryption Standard (AES).

The problems exhibited above are inherent to these types of cryptosystems. Fortunately there are other ways to achieve absolute security without these obstructions. In *information-theoretic* cryptography it is possible to achieve perfect security using a one-time pad encryption, provided the entropy of the pad is as large as the entropy of the message being encrypted. This was proven in a classical result by Shannon [30] and does not depend on the computational resources of an adversary or an unproven conjecture. This is in stark contrast to the security achieved in public key cryptography. There are problems

though with using information-theoretic cryptography in practice. The main frustration is the ability to obtain a perfectly random source to generate the pads used by the one-time pad. For long plaintexts messages this can be a considerable problem.

In this paper we study information-theoretic cryptography and key-management from a slightly different perspective. The goal of our work is to develop a key agreement scheme using a public random source and short random seed to generate one-time pad encryption schemes. Our approach is different than most, since we will assume an adversary has a bounded storage capacity, presumably smaller than the length of the public random source. Under this assumption we can prove our encryption schemes enjoy the same types of security benefits as those achieved by the information-theoretically secure ones. These schemes solve the problem of obtaining a long, perfectly random source to generate one-time pads. Furthermore, these constructions are explicit and do not rely on the computational abilities of an adversary.

1.1 Motivation

Under the bounded storage model, a public random source is available to all parties including the adversary. The source itself is not assumed to be perfectly random but rather satisfies some notion of weak randomness, which we define later. The public source could be obtained using a satellite broadcast or deep space radio telescope such as the LOFAR system, which can transmit on the order of several tens of Terabits per second [22]. In order to make the scheme practical the source must be transmitted at a very high data rate. A short random seed is shared between the two legitimate parties that wish to communicate. The seed is used to extract a mutual random key that is then used as a one-time pad Z to encrypt a message M, as $C = M \oplus Z$. The adversary is assumed to have storage resources close to the size of the public random source, perhaps on the order of several petabytes (10^{12} Bytes) in size. The adversary stores partial information about the source hoping to extract the encrypted message M at some later time. From a practical perspective we assume the adversary's storage resources depend not only on his total storage capacity, but also on his sustained data transfer rate and cost/budget constraints. These parameters are often overlooked in the literature and should be considered when developing security parameters. We provide an analysis of different storage technologies and their rates in Chapter 4. The key extraction schemes we discuss remain secure even if a party is compromised and the secret seed used to extract the key is obtained. If an adversary achieves an unlimited storage capacity at some point in the future we can show any one-time pad encrypted message remains secure indefinitely. Furthermore the secret seed may be used up to an exponential number of times and still remain secure.

As an illustration of the capabilities of the bounded storage model we consider the example presented by Dziembowski et al. in [11]. In the example we assume an adversary has a storage capacity of one petabit and both legitimate parties share a 6000-bit secret seed. The public random source is transmitted for a day and a half at a rate of 100 gigabits per second. Both parties are able to extract an expanded key of length 10 gigabits which is used as a one-time pad. Under this example, we can show the adversary has virtually no information about the expanded key even when an optimal storage strategy is employed.

The encrypted message remains secure indefinitely.

Several problems exist with implementing these schemes in practice. One of the more difficult problems is the synchronization of the two communicating parties. Since the public random source is being transmitted at such a high data rate, very accurate clocks will need to be employed to ensure bit sampling occurs at the same position. A single bit error could waste many days of computations. This problem is compounded by the rapid development in data storage technologies which will likely force bounded storage encryption schemes to either increase the source data rate or the amount of time needed to calculate an extracted key. Increasing the time for key extraction from days to weeks is not a practical option. Therefore the random source will need to be transmitted at a faster rate to offset the increases in storage technologies.

Unfortunately, the development rate of precision clocks is far behind that of data storage technologies. Some of the most accurate clocks built today are based on measuring the frequency of cesium atoms. These types of atomic clocks can achieve accuracy up to 2×10^{-14} Hz. They are expensive and bulky, and therefore are not viable for small portable applications. In order to mitigate synchronization and timing errors without increasing clock accuracy more sophisticated timing error correcting schemes will need to be developed. We explore developing such schemes in the bounded storage model. Our proposed solutions are simple and follow previously established work on communication error correcting codes with additional side-channel information. We compute the capacity of these channels when synchronization errors are present and develop tight lower bounds for a particular class of pseudo-invariant key expansion functions. The outline of this paper is as follows.

In Chapter 2 we first review the notion of a weak random source. We provide a brief history of weak randomness and show how it has evolved to its current form. In the next section we discuss various randomness extraction techniques and introduce the term extractor as it was first defined by Nisan and Zuckerman [26]. We finish the chapter with an overview of bounded storage encryption and discuss some of the recent work in the field. We show how extractors may be modified to construct encryption schemes under the bounded storage model.

In Chapter 3 we discuss the timing synchronization problem and show how to modify the Dziembowski et al. [11] encryption scheme to correct for synchronization errors. We compute the capacity of the channel and show how to construct error correcting encoding schemes without specific knowledge of the weak random source distribution.

In Chapter 4 we conduct a detailed case study of current and future storage technologies. We analyze the cost, capacity, and storage data rate of various technologies so that precise security parameters can be developed for bounded storage encryption schemes. This technology survey will provide an invaluable tool for developing these schemes in practice.

Chapter 2

Randomness Extraction

In this Chapter we cover the bounded storage model in more detail by reviewing the definition and tools used to formulate security. We begin our discussions with a brief history on weak random sources and show how the present-day definition fits appropriately within this model. We follow with a discussion on the current trends in randomness extraction. We define the term extractor as that used in [26] and show how extractors can extract near-uniform random bits using a source of weak randomness. Our discussion ends with a review of some recent results on bounded storage encryption. These results give insight on how to construct secure schemes using strong extractors [23]. Furthermore we show how such schemes remain secure under key re-use using the result of [23].

2.1 Weak Random Sources

Many applications use a source of randomness to support the computation of various tasks. When designing algorithms utilizing random sources one often assumes that perfect randomness is readily available. The performance of an algorithm depends on the characteristics of the random source, so any source imperfections may cause unexpected behavior in the algorithm. Unfortunately, obtaining a perfect source of randomness can be quite difficult. Instead, natural sources such as a Zener diode or quantum noise may often be substituted in its place. Dealing with the imperfections of the real world requires one to weaken the randomness requirement of the source. This allows one to design better performing algorithms in practice. In order to achieve this goal one needs to develop a model of weak randomness.

One of the first random source models was based on the biased coin problem. In the biased coin problem, a two-sided coin is independently flipped generating a sequence of heads and tails. The coin is biased to either one of the two sides, so that one side is more likely than the other. Von Neumann developed an extraction technique to produce a sequence of truly random bits [39]. In his technique he takes the results of two coin tosses and outputs a 1 if he receives HT and 0 for TH. Any other combination is discarded. It is easy to verify that this technique generates a sequence of truly random bits. The extraction of random bits however is not optimal since randomness is thrown away for the HH and TT combinations. There have been various improvements over the years which improve

upon Von Neumann's extraction technique. For example A. Juels et al. [19] developed an efficient scheme that asymptotically extracts the full entropy of the source. It would seem, based on the simplicity of the model, that the biased coin model is a sufficient model for weak random sources. The problem with this model, however, is that many sources of randomness do not have independent bits and a fixed bias. Therefore a more generalized model needs to be considered.

Santha and Vazirani [31] developed a weaker notion of randomness by relaxing the conditions of the biased coin model. In their model each bit is slightly random so that the conditional probability of each bit, given the past bits, is bounded below by some fixed constant $\delta > 0$. This guarantees that no string is completely deterministic. These types of sources are called SV-sources, and their definition is formalized below.

Definition 1. (Santha-Vazirani [31]) *A random source $X = (X_1, \ldots, X_n)$ over the set of binary strings of length n is said to be a δ SV-source if for any $0 < \delta \leq 1/2$ and every $i = 1, \ldots, n$,*

$$\delta \leq \Pr[X_i = a_i \mid X_1 = a_1, \ldots, X_{i-1} = a_{i-1}] \leq 1 - \delta$$

Using their model Santha and Vazirani showed that a single SV-source does not sufficiently guarantee the generation a single unbiased bit. They were, however, able to prove that any two or more independent SV-sources is sufficient. This is, as we shall show, the reason why randomness extractors use two source of randomness to extract random bits. Vazirani extended these results by showing the inner-product of two SV-sources produces a single unbiased bit [37].

The problem with the SV-source model is that it imposes too much structure on the random sources. To work around this restriction we look at our final model, which is based on the *min-entropy* of a distribution. The min-entropy model was first suggested by [40, 41] and has become the standard model of weak random sources.

Definition 2. *The min-entropy of a source X on $\{0,1\}^n$ is defined as,*

$$H_\infty(X) := \min_{x \in \{0,1\}^n} \log(\Pr[X = x]^{-1})$$

A source has min-entropy at least k if the probability of each element is bounded above by 2^{-k}. Intuitively this implies that a source of min-entropy k contains at least k-bits of randomness. In the Santha-Vazirani model an n-bit SV-source has min-entropy $n \cdot \log(1 - \delta)^{-1}$.

Following the impossibility result of Santha and Vazirani, it is not difficult to to show that for any 1-bit extraction technique over the set of n-bit strings there exists a min-entropy source of $n - 1$ bits that gives a constant output. Therefore extractors need random seeds in order to extract randomness.

In order to measure the quality of randomness in distributions a metric called *statistical distance* is used. Statistical distance measures the distinguishibility between two distributions over the same domain. Ideally we would like to make statements on how close an arbitrary distribution is to uniform. This will allow us to relax the condition that an extraction technique is approximately uniform rather than truly uniform. For this we use the following definition.

Definition 3. *Let X and Y be two distributions over the same domain U. The statistical distance between X and Y is,*

$$\|X - Y\| := \max_{A:U\to\{0,1\}} |\Pr[A(X) = 1] - \Pr[A(Y) = 1]|$$

The above definition can be rewritten in the equivalent form $1/2\sum_{u\in U} |\Pr[X(u)] - \Pr[Y(u)]|$. A distribution X is said to be ϵ-close to uniform over U if $\|X - U\| < \epsilon$. This implies that X is indistinguishable from random by any non-efficient algorithm.

2.2 Extractors

In the last section we showed how the Von Neumann algorithm could be used to extract random bits from the biased coin model. In this section we define the term extractor for weak random sources using min-entropy. If you recall, Santha and Vazirani showed that at least two SV-sources are needed to extract randomness. This motivates the following definition.

Definition 4. *A (k, ϵ)-extractor is a function,*

$$Ext : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$$

such that for any distribution X on $\{0,1\}^n$ with min-entropy $H_\infty(X) \geq k$ the distribution $Ext(X, U_d)$ is ϵ-close to the uniform distribution on $\{0,1\}^m$.

The random d-bit seed is used to extract many random bits from the source. The goal in constructing extractors is to maximize the output length while keeping the input seed small. Tight bounds exist on the seed and output length for non-explicit constructions (i.e. existence proofs). These bounds are formulated in the following Proposition.

Proposition 1. (Radhakrishnan & Ta-Shma [28]) *For each $\epsilon > 0$ and $k \leq n$ there exists a (k, ϵ)-extractor $Ext : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$ with $m = k + d - 2\log 1/\epsilon - O(1)$ and $d = \log(n - k) + 2\log 1/\epsilon + O(1)$.*

A great deal of work has been involved in developing explicit constructions that come close to meeting these bounds. Many of these constructions rely on recursion to amplify randomness. One such method was developed by Nisan and Zuckerman in [26]. Their extraction technique transforms a weak random source X into a sequence of Block-wise sources $X_1, \ldots, X_n$. To amplify the randomness a sequence of extractors is composed recursively yielding $Ext_n(X_n, Ext_{n-1}(X_{n-1}, \ldots, Ext_1(X_1, s)))$. This type of construction aims to minimize the initial seed length by expanding the seed at each recursion step. The sequence of extractors used in [26] is constructed using the Left Over Hash Lemma of [18]. In particular for any universal hash family $\mathcal{H} : \{0,1\}^n \to \{0,1\}^l$ and $k + 2\log(\epsilon^{-1})$ min-entropy source X, the distribution $Ext(X, \mathcal{H}) := (h, h(x))$ is ϵ-close to uniform on the set $\mathcal{H} \times \{0,1\}^l$. Since efficient constructions of universal hash families exist, it follows the extractors of [26] are explicit.

Following the original work of Nisan and Zuckerman several authors have developed improvements to the Block-wise source extraction technique [25, 33, 34, 42]. These new techniques aim to compose extractors in various ways. By composing an extractor that optimizes the seed length with an extractor that optimizes the output length, we can develop a new extractor that optimizes both. For an overview of recent developments in extractors see [29].

2.3 Bounded Storage Model

The bounded storage model was first introduced by Maurer in 1990 [24]. Under the model a public random source X is available to all parties. It is assumed that X satisfies some notion of weak randomness, so that extraction is possible. In practice the source could be transmitted via satellite or ground based network and be based on a natural source.

All parties including the adversary are assumed to have a limited storage capacity, presumably close to the size of X. No other computational restrictions are placed on the adversary; therefore, we assume they have unlimited computational power. The challenge in the bounded storage model is to devise information-theoretic secure encryption when storage capacities are arbitrarily close to the size of X. If an adversary could store X in its entirety, information-theoretic security would not be possible.

Under the model two parties share an initial random seed s that is chosen uniformly. Using the seed s and source X both parties extract a mutual random key which is used as a one-time pad. The adversary can store up to a fraction of the bits of the weak random source X using an arbitrary storage function $h : X \to U$. The adversary loses the ability to access X after the storage function h has been applied. In addition to the above conditions the adversary is able to compromise one of the two parties and extract the secret seed s. It is assumed that the adversary gains unlimited storage capacity after s has been learned. This allows the possibility that an adversary tries to decrypt the message at some point in the future when more storage capacity is available. A key-expansion function f is said to be secure in the bounded storage model if the conditional probability of gaining information about the extracted key K given s and $h(X)$ is ϵ-close to uniform. More precisely we have the following definition.

Definition 5. *Let $f : X \times \{0,1\}^d \to \{0,1\}^m$ be a key-expansion function over a weak public random source X. For any arbitrary storage function $h : X \to U$ storing up to a fraction $|h(X)|/H_\infty(X)$-bits of X, we say f is ϵ-secure under the bounded storage model if for any k min-entropy source X, the random variable $(f(X,S), h(X), S)$ is ϵ-close to $(U_m, h(X), S)$, where $s \overset{R}{\leftarrow} \{0,1\}^d$.*

2.4 Recent Work

The original bounded storage model [24] used a uniform random string as the public random source. Under this model Maurer was able to prove with high probability that the mutual information provided by both that the encrypted plaintext and stored source bits gives

zero information about the plaintext distribution. This holds true for any storage strategy storing up to a fraction of the *original bits* of the public random source. This analysis did not include the more general case of storing an arbitrary function of the original bits. Cachin and Maurer later refined the model to include this generalization [5].

Using concepts derived from semantic security and indistinguishability Aumann et al. [3] were able to show that extracting a single bit of information about the key is negligible. This is a slightly weaker notion of security than the Renyi entropy results obtained by [5]. The authors further proved under the indistinguishability condition that if the secret seed was later compromised, then any encrypted message occurring before the compromise remained secure indefinitely. This important property was given the name *everlasting security.*

Some of the recent work in the bounded storage model has sought to develop extractor conditions so they may be used as key-expanders. The problem with the current definition of extractors is that it does not guarantee indistinguishability if the secret seed is exposed. Lu was able to work around this technicality by defining a new class of extractors called strong extractors that yield secure key-expansion functions [23]. A strong extractor is an extractor that remains ϵ-close to uniform when the secret seed s is exposed. Formally a strong extractor is defined as follows.

Definition 6. (Strong Extractor) *A (k, ϵ)-strong extractor is a function,*

$$Ext : \{0,1\}^n \times \{0,1\}^d \rightarrow \{0,1\}^m$$

such that for every distribution X on $\{0,1\}^n$ with min-entropy $H_\infty(X) \geq k$ the distribution $(U_d, Ext(X, U_d))$ is ϵ-close to uniform.

Using any strong extractor Ext as a key-expansion function, Lu, was able to prove these schemes remain secure under the bounded storage model. Further, following the work of Ding et al. [9] Lu, was able to show that reusing the same key-expansion seed degrades the security of the encryption scheme by only a linear amount. This holds true provided that the sources used in extraction $X_1, \ldots, X_t$ form a Block-wise source of at least min-entropy k. The above two results are summarized in the following Proposition.

Proposition 2. (Lu [23]) *Any strong $(\delta n, 2^{-3k})$-extractor yields a secure encryption scheme with everlasting security, whose key-expansion seed can be securely reused up to 2^k times.*

Chapter 3

Synchronization Error Correction

There are a number of engineering obstacles that need to be solved before key management and encryption under the bounded storage model can be realized. One of the critical obstacles hindering its adoption is the construction of a scheme that achieves reliable communication in the event that timing synchronization errors occur. In this section we provide a general strategy for solving the timing synchronization problem by proving the existence of message encoding schemes that achieve a small probability of error. Using the receiver's unsynchronized key-expansion as side-channel information, we compute the maximum capacity of the channel and provide tight lower bounds for a particular class of key-expansion functions that are pseudo-invariant to timing errors. Using these message encoding techniques in combination with Dziembowski et al. [11] encryption scheme we can construct a scheme that solves the timing synchronization error problem.

We begin this section by reviewing the encryption scheme of [11] in more detail. We discuss the various properties which make the scheme attractive for synchronization error correction and classify these properties in a more general framework. Using tools from information theory we show how the scheme of [11] can be modified to correct for timing errors. We compute lower bounds on its capacity and show how to develop encoding schemes that are independent of the public random source distribution.

3.1 Bounded Storage Model with Errors

In order to make key-management under the bounded storage model practical the public random source will need to maintain a very high data rate. This leads to a number of potential timing and synchronization problems between the two parties sampling the source. Managing clock drift and skew errors between two hardware devices may be very difficult and expensive. In most key extraction schemes a position error of a single bit will make decryption impossible. To mitigate these problems we need to construct a key-extraction scheme that allows for timing synchronization error.

The bounded storage encryption scheme of Dziembowski et al. [11] is based on a simple key-expansion construction. The simplicity of the scheme makes it an ideal candidate for

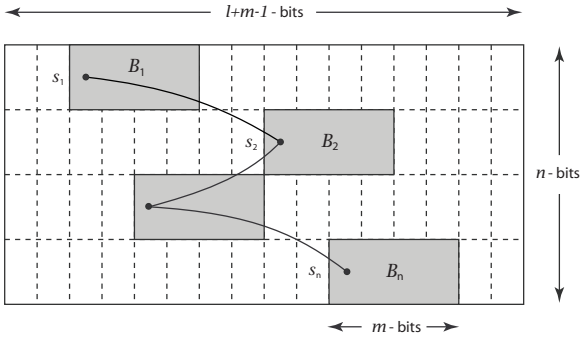

Figure 3.1. Block sampling method of Dziembowski et al. [11] encryption scheme.

use in an actual implementation. Other encryption schemes such as those used in [23, 36] provide near optimal security parameters. These schemes however are far more complex and not suitable for synchronization error correction. The encryption scheme of [11] rearranges a public random source X into a $n \times (l + m - 1)$ matrix with $|X| = n(l + m - 1)$. The secret seed $s := (s_1, \ldots, s_n)$ is randomly generated by choosing bit positions $s_i \in \{1, \ldots, l\}$ in each row of the matrix $i = 1, \ldots, n$. The bit positions s_i are used to denote the starting positions of m-bit blocks B_i. Fig. 3.1 illustrates this sampling method. The key-expansion function $f : X \times S \to \{0,1\}^m$ is defined as the component-wise XOR of the n m-bit blocks B_i,

$$f(x, s) := \bigoplus_{i=1}^{n} B_i \tag{3.1}$$

The security of this key extraction scheme is stated in the following Proposition.

Proposition 3. (Dziembowski & Maurer [11]) *Suppose an adversary has access to an arbitrary storage function $h : X \to \{0,1\}^t$. Then the random variable $(f(X,S), h(X), S)$ is $m \cdot (2^{|X|-H_\infty(X)+t-\Delta} + \epsilon)$-close to $(U_m, h(X), S)$, where $s \xleftarrow{R} \{0,1\}^{n \cdot \log_2 l}$,*

$$\Delta := nl(1-\zeta)\tau^2(\log_2 e)/2 - n(1 + (1-\zeta)(n \log_2 l + m + 1))$$

and

$$\epsilon := \tau^{\zeta m}/2.$$

for any $\tau, \zeta \in [0,1]$.

If a synchronization error occurs in the above scheme, then each of the blocks B_i will shift depending on the timing error. For example, if one party is off by 2-bits then the blocks B_i in Fig. 3.1 will shift to become those in Fig. 3.2. Observe that the extracted keys $K_1 = \bigoplus_{i=1}^{n} B_i$ and $K_2 = \bigoplus_{i=1}^{n} B_i'$ retain 2-bits of information between each other, since B_i and B_i' overlap. We can use this to our advantage by encoding the plaintext messages with a synchronization error correcting code. The extracted key K_2 can be used by the receiver as side-channel information for decryption and decoding of the plaintext.

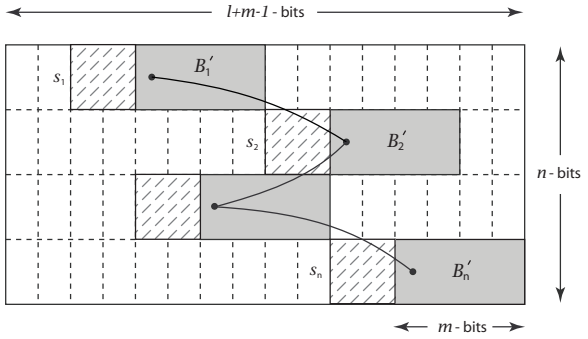

Figure 3.2. A 2-bit synchronization error.

We model the timing synchronization error between two legitimate parties using a probability distribution called *Shift*. The *Shift* distribution accounts for the clock drift and skew errors between two devices and measures the probability that one party starts sampling the source at a different bit position than the other. Instead of the public random source X having a fixed block length we generalize its size by including infinite length strings. This allows for synchronization errors of arbitrary size. The shift distribution *Shift* acts on the public random source $X = \{x_i\}_{-\infty}^{\infty}$ by shifting the source to the left or right, providing each party a separate viewpoint of the same source. If one party has viewpoint X and the other $Shift_i(X)$, then their sampling will be off by i bits. When we say an infinite source X has min-entropy k, then we will assume for some fixed block length n every n-bit window has at least k-bits of min-entropy. This guarantees synchronization errors do not degrade the security of the one-time pad.

3.2 Channel Capacity with Synchronization Errors

Using the encryption scheme of [11], our goal is to compute the maximum data rate at which reliable communication is possible when synchronization errors are present. This maximum rate is called the operational capacity of the channel.

To compute the operational capacity of the channel we use an information-theoretic tool called entropy. The entropy of a distribution is a measure of the average number of random bits needed to describe the source. The more uncertain a source is, the larger its description length will be. A precise definition is given below.

Definition 7. *The entropy $H(X)$ of a discrete random variable X is defined as,*

$$H(X) := -\sum_{x \in \mathcal{X}} \Pr[X = x] \cdot \log_2 \Pr[X = x]$$

The entropy of a source is directly related to its data compressibility. For any lossless encoding scheme the entropy of a source is a lower bound on its compression limit. If a random variable X is uniformly distributed over the set $\{0,1\}^n$, then the entropy of X

is n. This implies that a perfectly random source is already compressed and cannot be compressed further.

We can condition the uncertainty of a random variable Y based on its observation of X. This is called conditional entropy.

Definition 8. *The conditional entropy $H(Y|X)$ over the joint distribution $p(x, y)$ is defined as,*

$$H(Y|X) := \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} \Pr[X = x, Y = y] \cdot \log_2 \Pr[Y = y \mid X = x]$$

Using both entropy and conditional entropy we can combine them to develop a measurement called *mutual information*. Mutual information measures the amount of information a distribution X reveals about a distribution Y. In a noisy communication channel the transition probability $p(y|x)$ describes the uncertainty of Y when transmitting X. The noisier the channel, the less X will reveal about Y, and the more randomness Y will contain. In an ideal channel the output Y will be completely deterministic and invertible with respect to X, since this allows decoding with zero error. In order to send many messages through the channel, the transition probability must be small to minimize decoding error. The mutual information between two random variable X and Y is defined using the following equation $I(X;Y) := H(Y) - H(Y|X)$. The information-theoretic capacity of a channel is derived from mutual information by maximizing over all probability distributions X. In a noisy communication channel this measures the maximum number of bits the output Y retains about the input X.

Definition 9. *The information channel capacity of a discrete memoryless channel is defined as*

$$C := \max_{p(x)} I(X;Y)$$

Shannon showed that the information channel capacity is an upper bound for all achievable rates with error probability approaching 0. This holds true for all discrete memoryless channels.

In order to compute the channel capacity of the encryption scheme of [11] we need to use an extension to the above classical result. We view channel errors under the bounded storage model as a one-time pad encryption of the transmitted message. We allow side-channel information to be available at the receiver in the form of the extracted key $f(\mathit{Shift}(x), s)$ with error distribution *Shift*. Channel encoding with additional state information has been studied extensively by several authors [7, 15, 16, 20]. In [16] it was shown that the channel capacity is,

$$C = \max_{p(x)} I(X;Y|S) \tag{3.2}$$

when additional state information S is available at the receiver. Using a modified version of this result, we are able to compute the maximum data rate in which reliable communication is possible under the bounded storage model.

Proposition 4. *Let $f : X \times S \to \{0,1\}^m$ be a key-expansion function and Shift a synchronization error distribution. Then the information channel capacity of the bounded storage model with synchronization error and side-channel information at the receiver is,*

$$C = m - H(f(X,S)|f(\mathit{Shift}(X), S))$$

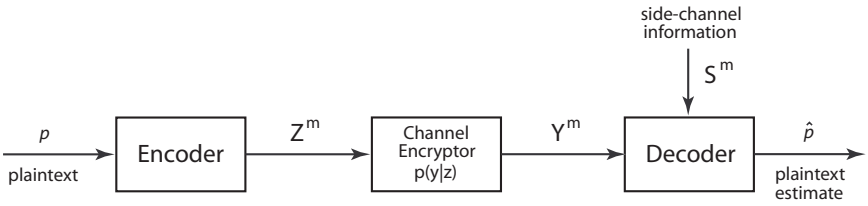

Figure 3.3. Encryption channel with side-channel information at the receiver. Unsynchronized key-expansion=side-channel information.

The capacity of Proposition 4 is an upper-bound on all achievable rates when the joint distribution of the source and encryption output are discrete memoryless channels. Therefore, it follows for any $\epsilon > 0$, $R < C$, and sufficiently large N, there exists synchronization error-correcting codes with size 2^{NR} that have probability of error less than ϵ.

The above Proposition is a general statement for any key-extraction scheme under the bounded storage model. This result can be improved for a particular class of key-extraction schemes by exploiting the properties of [11]. We define the class of all pseudo-invariant key-expansion functions as the set of all $f : X \times S \to \{0,1\}^m$ such that $f(\mathit{Shift}(X), S) = \mathit{Shift}'(f(X,S))$ for some distribution Shift', where $\mathit{Shift}'_i(y) := (y_{i+1}, y_{i+2}, \ldots, y_n, v_1, \ldots, v_i)$ for some i-bit distribution v conditioned on the last $n-i$ bits of y. This implies that a timing error on the input corresponds exactly to the same timing error on the output, with additional random bits added to replace the bits lost by timing.

We can compute a lower bound on the capacity of this particular class of key-expansion functions by looking at a particular instance, which we call the *ideal* case. In the ideal case, both the public random source and key-expansion function f return m perfectly random bits. The synchronization error distribution Shift' adds i-bits of independent randomness to the output string whenever a synchronization error of i-bits occurs, $|i| < m$.

It is not difficult to show that the capacity of the ideal case is $m - H(\mathit{Shift}'(U_m)|U_m)$, when additional side-channel information is available at the receiver. Using this result we have the following Proposition.

Proposition 5. *The class of all pseudo-invariant key-expansion functions $f : X \times S \to \{0,1\}^m$ yield capacities*

$$C \geq m - H(\mathit{Shift}'(U_m)|U_m)$$

with equality in the ideal case.

Designing synchronization error correcting codes without specific knowledge of the public random source distribution can be difficult. Therefore, it would be beneficial to be able to design an encoding scheme for one type of source so that it may be plugged into the other without knowing the particular distribution. Fortunately, as we will show, it is possible to do this for the class of pseudo-invariant key-expansion functions. By designing a code for the ideal case, the same code will also work for the more general pseudo-invariant key-expander.

If we assume our public random source consists of t random sources $(X_1, \ldots, X_t)$ such that any source conditioned on the previous sources still has min-entropy at least n, then we get the following result.

Proposition 6. *Let $f : X \times S \to \{0,1\}^m$ be any pseudo-invariant key expansion function that is also a $(\delta n, \epsilon)$ strong-extractor with $\log_2 \epsilon^{-1} \leq n(1-\delta)-m$. If the maximum probability of error of an encoding scheme for the ideal case over t encryptions is λ_e^t then the maximum probability of error of using the same encoding scheme for f over t encryptions is bounded above by,*

$$P_{error}^t \leq 4t \cdot \epsilon + \lambda_e^t$$

The proof in the above Proposition uses a similar argument as those used in Lemma 3 and 4 of [23]. Therefore we can develop encoding schemes without knowing the particular distribution of the source and still achieve small probability of error.

3.3 Open Problems and Future Work

There are several open problems that are worth exploring that we did not necessarily cover in the previous sections. It would be interesting to see if it is possible to develop a closed-form solution for the capacity of the channel when state information is available both at the encoder and decoder. This problem has been explored before for channel capacities with two-sided state information [7]. We could give the encoder the expanded key $f(X,S)$ and impose the condition that the distribution $(f(X,S) \oplus \text{Encode}_{f(X,S)}(m), h(X), S)$ is ϵ-close to $(U_m, h(X), S)$. We expect, however, that any improvements in the capacity using this scheme will be essentially negligible. Therefore, the one-sided case we considered in the previous section should provide sufficiently large capacity close to the two-sided example.

It would also be interesting to explore the capacity of the channel when the window sampling size of the source and key-expansion outputs are allowed to expand. This should give the maximum data rate at which single encryptions could be transmitted. Furthermore it would be interesting to explore if it is possible to achieve the same capacity results as above when the public random source is a block-wise source over multiple encryptions rather than a discrete memoryless channel.

We have also considered exploring whether the class of pseudo-invariant key-expanders produce the best capacities out of any other key-expander functions. Specifically, it would be interesting to see if it is possible to prove that for any capacity C achieved by a key-expander f, a pseudo-invariant key-expander can achieve equal or better capacity.

To make the bounded storage model practical, we have also considered incorporating fuzzy extractors to correct for channel errors in the public random source. These ideas have been explored in [8, 10]. Finally, we are currently developing explicit synchronization error correcting schemes based on the ideal channel presented in the previous section. These schemes should provide low error probability based on Proposition 6. It will be interesting to see if efficient constructions exist that approach the capacity limit.

Chapter 4

Storage Technology Review

When studying the practical aspects of the Bounded Storage Model of encryption, at some point one needs to consider the question of just how bounded is the storage of the adversary one is attempting to thwart. And when considering the bounds of the adversary, one needs to take into account not just the storage capacity of the adversary, but also their maximum sustained data transfer rate. This is a notion that has consistently been overlooked in practical analyses of the bounded storage model, an odd omission, considering that an adversary with even an exabyte of storage capacity at their disposal will be rendered irrelevant if ten bits pass from the public stream in the time it takes them to store just one.

It is with this idea in mind that this analysis looks at the storage capacity, cost, and recording speed of a number of currently available, and proposed future storage technologies.

4.1 Motivation

The motivation behind this analysis is quite basic: if one wishes to prove the security of a scheme in the bounded storage model, one must first demonstrate that the adversary is unable to have captured the entirety of the public random stream from which the private secrets are derived. Because of this, the security of the system is dependent on the storage capabilities of the adversary. This is analogous to basing the security of a traditional cryptosystem on the amount of computing power available to an adversary. For example the Data Encryption Standard, DES [14], was considered to be a secure cipher when it was developed in the 1970s, as the amount of work required to break DES via brute force, $\mathcal{O}(2^{55})$ operations, was large enough to be considered infeasible for the processing power available at the time. Today, DES can be broken by brute force in just a few hours on relatively inexpensive custom hardware. In a similar sense, one could prove a system secure in the bounded storage model against an adversary who is able to store a given amount of data. However, that same scheme may be breakable by an adversary who can store a few orders of magnitude more data.

Of course, storage capacity is not the only parameter that will determine an adversary's

strength. Other parameters to consider are the speed at which the data can be stored[1], as well as the costs of storage, both in obtaining the storage mechanism and media as well as overhead, such as space, cooling and power requirements. Another issue that will not be considered in this report though it should be analyzed is the rate at which a data stream can be reliably transmitted and received.

4.2 Current Technology

There are a number of different storage technologies available today (and expected in the near future) that an adversary may consider for capturing the contents of a public random source. We will cover the most likely of these choices here. Note that in this section, we will make use of a slight abuse of notation, taking $\mathcal{O}(x)$ to mean "approximately x", rather than its standard meaning in computer science, which would make no distinction in scale between $\mathcal{O}(1)$ and $\mathcal{O}(100)$.

4.2.1 Magnetic storage

Magnetic storage is the most common form of long term storage today, with magnetic hard discs being the standard form of secondary memory on computers, and magnetic tape remain quite popular for storage of extremely large data sets and use as a backup medium. Their ready availability make them obvious choices for use by an adversary [12].

Magnetic tape Magnetic tape has long been the favored format for storing large data sets for archival and batch processing purposes. This was particularly due to tape's large volumetric data density and low cost. While magnetic tape would work well for storing a public random stream because its inherently sequential organization is not a detriment for this purpose, it cannot be considered a viable solution due to its low data transfer rates. Even the best tape systems are limited to write speeds of approximately 30 MiB/second, with speeds of 1-3 MiB/second being more common.

Capacity Individual tape cartridges can hold $\mathcal{O}(0.5)$ TiB depending on the cartridge type. However, large "library" setups can be installed, allowing for $\mathcal{O}(1)$ petabyte of storage for a single site. Areal density of magnetic tape is currently about $\mathcal{O}(2)$ GiB/in^2, and is not expected to rise significantly in the future, with most increases in capacity expected to come from fitting more linear feet of tape in a cartridge.

Speed While magnetic tape has very good data *storage* capabilities, the data transfer rates available for the medium likely remove it from any consideration for use in initially

[1]For example, a compact disc can hold approximately 700MiB, but even the fastest CD writers require a few minutes to transfer this much data to the CD. If the data is passing in as a relatively slow 10MiB/sec stream, it would consume the entirety of an average computer's memory long before a single CD could be written.

capturing a bounded storage public random string. The fastest tape systems available today have a maximum transfer rate of $\mathcal{O}(1)$ MiB/sec and, as would be expected, are priced in the upper ranges of the tape drive market. More common, and less expensive are drives and tapes that allow for transfers in the $\mathcal{O}(2)$ MiB/second range.

Cost Magnetic tape is far and away the least expensive of the storage technologies described in this report. The initial cost of a system can be large, with the tape drive mechanisms costing $\mathcal{O}(5000)$ \$/unit. However, the storage medium itself is quite inexpensive, with tape cartridges costing $\mathcal{O}(0.03)$ \$/GiB of storage.

Hard discs Magnetic disc drives have become the most prevalent form of secondary data storage for computer systems. Because of this, a great deal of resources have been allocated to improving the performance and capacity of this medium, with great advances being made in both areas in the past decade. Unfortunately, the ability to advance much further in either area is in great doubt due to a number of considerations.

The primary reason is that the magnetic particles used to store the data have become so small, they are nearing what is known as the superparamagnetic limit. That is, once the volume of a magnetic particle becomes too small, the energy required to change the direction of the magnetic moment of the particle is on the same order as the ambient thermal energy. When this happens, the information stored in the particle randomly reverses itself at a rate far higher than that needed to provide reliable storage. There are a few techniques being researched as a means of staving off the superparamagnetic effect as long as possible[2], but increases in performance beyond another order of magnitude are not expected to materialize [35].

Capacity Current hard drive capacity is $\mathcal{O}(0.5)$ TiB/unit at the upper end of the spectrum, with areal densities of $\mathcal{O}(15)$ GiB/in^2. Due to the problems associated with the superparamagnetic effect, areal densities are not expected to increase more than an additional order of magnitude, with physical upper limits estimated in the area of $\mathcal{O}(120)$ GiB/in^2.

Speed Serial ATA connectors are capable of transferring $\mathcal{O}(300)$ MiB/second, with double that rate expected within two years. However, the drives themselves are currently limited to $\mathcal{O}(100)$ MiB/second. There are a number of reasons for this, some of which may be overcome by creative engineering. Unfortunately, others, like the superparamagnetic effect which limits areal density, are physical in nature and limit how quickly the magnetic momentum of a particle can be reversed. It is believed that write speeds can be increased by another order of magnitude, but no more.

Cost Costs of magnetic disc drives have fallen greatly over the years decreasing 1000-fold over the past fifteen years to their current cost of less than $\mathcal{O}(1)$ \$/GiB of storage, with

[2]For example, using perpendicular, rather than longitudinal magnetic recording, or thermally assisted recording, as seen in magneto-optical solutions.

price/GiB likely expected to maintain an inverse relationship with areal density.

4.2.2 Optical storage

Perhaps the second most prevalent secondary storage medium is the optical storage provided by compact discs (CDs) and Digital Versatile Discs (DVDs). The high capacity, small footprint, random access nature, and ease of portability these discs provide allow them to fill the space that exists between hard disks and magnetic tape, providing archival storage of mid-sized or loosely related data sets. For the most part, CDs have been superseded by DVDs as they provide an order of magnitude more capacity in the same footprint and at a lower cost, but CDs are still quite prevalent due to backwards compatibility in the DVD drives.

One thing that should be noted about optical storage technology is that, unlike magnetic storage, releases of new standards and products are delayed and centered around applications, such as storing movie-length compressed video with DVDs and high definition video capabilities for the next iteration of optical storage, rather than releasing new products with every incremental evolutionary improvement in technology.

Compact Discs Compact disc writers are nearly ubiquitous in computer systems today due to their low cost and the convenience offered by the medium. However, despite the flexibility offered to the average user by CDs, they aren't appropriate choices for recording bounded-storage public streams, as the capacity and write speeds are both too low to capture any stream likely to be seen in practice.

Capacity At the time compact discs were being developed, the laser technology available limited the areal density of the discs to $\mathcal{O}(0.5)$ GiB/in^2 resulting in a maximal capacity of $\mathcal{O}(700)$ MiB/disc.

Speed Write speeds of optical media such as CDs and DVDs are limited by the ability of the discs themselves to sustain the forces imposed by the drives spinning the medium at high speeds. The upper limit of the current discs is $\mathcal{O}(5)$ MiB/second.

Cost One of the reasons for the popularity of the compact disc medium is the low cost of both the reading/writing mechanism and the discs. Drives that will read and write at the maximum speeds available can be purchased for $\mathcal{O}(50)$ \$/drive, with blank discs priced at $\mathcal{O}(0.3)$ \$/GiB storage.

Digital Versatile Discs As laser technology improved after the development of the compact disc, research efforts focused on the development of a new standard that could take advantage of these developments, leading to increases in capacity and write speeds. The result of this work was the digital versatile disc (DVD).

Capacity Currently, digital versatile discs com in one of two formats, single-layer and double-layer discs. Single-layer discs, which are more common in computer systems can hold $\mathcal{O}(4.5)$ GiB/disc, with an areal density of $\mathcal{O}(2)$ GiB/in^2. Double-layer discs can hold $\mathcal{O}(9)$ GiB/disc, with an effective areal density of $\mathcal{O}(4)$ GiB/in^2.

Speed DVDs, like their predecessor, the compact disc, are limited by the ability of the medium to withstand the physical forces imposed on it. Other technical considerations have also kept DVDs from reaching this limit yet, with an upper bound on write speeds of $\mathcal{O}(20)$ MiB/second.

Cost As with compact discs, both the reader/writer mechanism and the DVD media are inexpensive, with high speed drives available at $\mathcal{O}(50)$ \$/drive, and blank, single-layer discs costing $\mathcal{O}(0.5)$ \$/GiB storage.

4.2.3 Magneto-optical storage

Magneto-optical (MO) storage is a form of storage that targets the mid-size and loosely related data-set archive market. MO technology is a bit of a hybrid between magnetic hard disc and DVD technology, using a magnetic layer to store the data like a hard drive, but utilizing lasers to read the data and to assist in the writing of data. In order to write a bit of information, a laser is used to heat a magnetic particle to the point where it loses its previous information, and a magnet is used to store the new data as the particle cools. This technique allows for (among other things, such as resistance to accidental magnetic erasure) an efficient implementation of perpendicular recording, where the direction of magnetic momentum is oriented perpendicular to the surface of the disc, allowing for higher areal data densities than those possible with longitudinal magnetic recording. Unfortunately, as with tape and optical storage systems, the write speed of the MO discs is quite slow compared to magnetic hard discs.

Capacity Due to their ability to perform perpendicular recording, magneto-optical media are capable of reaching areal densities on the order of magnetic hard discs, with standard media capable of holding $\mathcal{O}(10)$ GiB/disc, and special, write-once discs available in sizes up to $\mathcal{O}(30)$ GiB in a 3.5 inch form factor. Note that these capacities are likely to increase as research in magnetic write-heads and shorter wavelength lasers continues.

Speed Unfortunately, the recording method that allows the MO media to reach such high areal densities also restricts the write speeds to be much slower than those of standard magnetic storage, with current write speeds of $\mathcal{O}(10)$ MiB/second, and few significant improvements expected.

Cost Targeted at the archival storage market rather than the larger consumer market, magneto-optical storage solutions are much more expensive than either pure magnetic or optical solutions, with reader/writer units selling for $\mathcal{O}(200)$ \$/individual drive mechanism,

and "jukebox" style drives selling at prices of $\mathcal{O}(5000)$ \$/unit. The media is also much more expensive, costing $\mathcal{O}(6)$ \$/GiB for the cartridges.

4.2.4 Volatile and non-volatile RAM

Currently, the two main RAM technologies are Dynamic Random Access Memory (DRAM) on the volatile side, and FLASH memory on the non-volatile side. Of the two, DRAM is by far the most common form of primary storage used in computing systems due to its high speed and small footprint. FLASH memory is generally limited to use in external components such as digital cameras and music players due to its slow write speeds and limited number of read/write cycles[3].

Capacity Capacity for RAM is a bit smaller than that of magnetic disc drives, with areal densities of $\mathcal{O}(2)$ GiB/in^2 for DRAM and $\mathcal{O}(5)$ GiB/in^2 for FLASH. A more important consideration, when it comes to DRAM, is the environment in which it is being used. Most modern operating systems and computer architectures limit the amount of useable memory in a computer to 4 GiB of RAM. If one wishes to utilize more memory than that, special purpose hardware must be designed and built.

Speed DRAM is the fastest of the storage technologies discussed in this document, with write speeds of $\mathcal{O}(400)$ MiB/second. In fact the performance of memory on a number of computer systems is limited by the speed of the data bus used to pass information between the processor and the memory. FLASH memory, on the other hand, is much, much slower than DRAM, with write speeds of only $\mathcal{O}(5)$ MiB/second.

Cost The cost of RAM is the primary disincentive of its use, with prices of $\mathcal{O}(100)$ \$/GiB for both FLASH and DRAM, placing them two orders of magnitude above magnetic disc storage.

4.3 Future technology

4.3.1 Next-generation optical storage

With the advent of high-definition television (and general consumer desire for increased storage capacity), a new iteration of optical storage solutions is nearing release. There are two leading products, both evolutions of the DVD format, vying to become the successor to the DVD standard: HD-DVD and Blu-Ray DVD. While there are a number of differences between the two formats, they are both centered around the idea of using a shorter wavelength laser as a way to increase the areal density of the storage medium.

[3]FLASH memory is only reliable for $\mathcal{O}(10000)$ read/write cycles before the memory locations being accessed fail and become unusable. While this limitation is unlikely to affect the performance of devices like cameras, where the number of accesses to any given memory cell is quite limited under normal usage, unlike a computer, which can easily perform hundreds or thousands of reads or writes of any given cell per second.

Capacity The Blu-Ray design uses a different medium than that of the current generation of DVDs, and it is placed much closer to the surface of the disc than in DVDs. This allows for a higher numerical aperture lens to be used for the laser, resulting in smaller pit sizes and therefore, higher areal density than HD-DVD which uses the same medium as current DVDs. The densities of the two standards are $\mathcal{O}(7.5)$ GiB/in^2 for HD-DVD, and $\mathcal{O}(12.5)$ GiB/in^2 for Blu-Ray. However, this design choice for Blu-Ray DVD also limits the data to storage on just one side of the disc, whereas HD-DVD can store data on both sides. This means that the total capacity of a Blu-Ray disc is $\mathcal{O}(50)$ GiB, where an HD-DVD disc can hold $\mathcal{O}(60)$ GiB [4].

Speed The initial write speeds of the two formats are going to be lower than that of current top-of-the-line DVD recorders, with speeds of $\mathcal{O}(7)$ MiB/second for HD-DVD and $\mathcal{O}(9)$ MiB/second for Blu-Ray. However, the write speeds are expected to increase fairly quickly, with maximum write speeds of $\mathcal{O}(40)$ MiB/second for HD-DVD and $\mathcal{O}(50)$ MiB/second for Blu-Ray [13].

Cost Initial costs for either technology is expected to be 2-3 times higher than that of current DVD technology, but is expected to settle quickly to the current prices of DVD technology.

4.3.2 Holographic storage

Both magnetic and optical disc storage record data by making changes on the surface of some recording medium. Holographic data storage, a decades old concept finally made practical by advances in technology, stores data throughout the volume of its medium, offering higher storage capacities in similar amounts of space than current technologies. One large reason for this is that the storage density of a CD or DVD is inversely proportional to the square of the wavelength of the laser being used to record the information, while holographic storage densities are inversely proportional to the wavelength cubed [1].

Recently, six Japanese companies formed the Holographic Versatile Disc Alliance in order to accelerate the development of holographic storage technology, in particular, the holographic versatile disc (HVD), a medium with the same form factor as a DVD [32].

Capacity Using a form factor the same size as a standard DVD, the (HVD) is planned to debut with an equivalent areal density of $\mathcal{O}(120)$ GiB/in^2, with $\mathcal{O}(275)$ GiB discs set to be released by 2007, and $\mathcal{O}(1.5)$ TiB discs (effective areal density of $\mathcal{O}(325)$ GiB/in^2) expected by 2010, and an expected upper limit of $\mathcal{O}(3.5)$ TiB per disc (effective areal density of $\mathcal{O}(865)$ GiB/in^2).

Speed The Holographic Versatile Disc Alliance is claiming that its initial offerings will have write speeds of $\mathcal{O}(20)$ MiB/second, and that the speeds are expected to rise to $\mathcal{O}(120)$ MiB/second by 2010.

Cost No current price projections are available.

4.3.3 MEMS storage

Another avenue of research in storage technologies is focusing on the use of micro-mechanical devices to embed information in a tiny storage surface. The leading technology in this area is the "millipede" project of IBM that is using components from atom force microscopy to create, detect, and remove microscopic depressions in the surface of a polymer medium. Arrays of thousands of components work in parallel to allow high data transfer rates to be realized [38].

While prototypes of this technology have been demonstrated, production versions are not expected for at least another two years.

Capacity In March of 2005, IBM demonstrated a prototype device with an areal density of $\mathcal{O}(60)$ GiB//in^2, and they claim that the technology should yield areal densities exceeding $\mathcal{O}(120)$ GiB/in^2.

Speed Unfortunately, the data transfer rates of these devices is not particularly high, with $\mathcal{O}(1)$ MiB/second read speeds. And while no real projections have been made for its write speeds, it is expected to be slower than reading. The slow speeds, as well as a limited number of read/write cycles ($\mathcal{O}(100,000)$) for any given area of the medium indicate that this technology is targeted as a replacement for FLASH memory, and not magnetic discs or DRAM.

Cost No current projections are available.

4.3.4 Magnetic Random Access Memory

One of the biggest problems with DRAM is that it is volatile, that is, because the state of a bit is represented by the presence or absence of an electrical charge at the corresponding location on the chip, once power to the chip is removed, all of the information stored on the chip is lost. While a number of non-volatile memory solutions such as FLASH memory exist, they have all been far too slow in operation to replace DRAM.

Recently, however, research led by IBM and Infineon into magnetic random access memory (MRAM), a high-speed non-volatile form of memory has begun to show favorable results [17, 27].

Capacity As the technology is still in the laboratory stages of development, the ultimate capacity of MRAM chips is currently unknown. However, $\mathcal{O}(2)$ MiB MRAM chips have been demonstrated, and sizes are expected to be on par with, if not smaller than DRAM.

Speed Much like capacity, the ultimate write speed of MRAM is still unknown, but expected to be at least on par with DRAM. Research done at IBM has shown that the speed of MRAM can be up to 6 times faster than that of DRAM, but it remains to be seen if those results hold for large-scale MRAM systems.

Cost No current projections are available.

4.4 Attack Models

It is clear from the above comparison of data storage technology that any attacker wishing to capture a public random stream is going to have to base the design of their system around magnetic hard drives and DRAM, as these are the only available technologies that can keep pace with any stream of a reasonable rate.

Even with the next generation of storage technology nearing public release, the basic system design is unlikely to change much, with the slightly faster MRAM used to replace DRAM[4], and holographic storage likely replacing magnetic hard drives due to its increased capacity and comparable write speed.

Still, even with these restrictions, there are a few classes of adversaries that one may wish to protect against in the bounded storage model. Below, we try to determine the capabilities of three such classes: an attacker with a single COTS computer system, an attacker with a cluster of COTS computer systems, and an attacker with special-purpose hardware.

4.4.1 An attacker with a single machine

Any attacker with a single system is going to be limited primarily by the speed of primary storage on his system, the main memory. Given the capabilities of DRAM, as well as the memory management abilities of modern operating systems, this places a hard limit of capturing 10 seconds worth of data coming from a 400 MiB/second stream without the aid of secondary storage.

With secondary storage comes additional concerns. Single drives would only lower the maximum data rate that the adversary could handle to $\mathcal{O}(100)$ MiB/second. Even the promise of holographic storage does little to change this. However, one can improve this performance by using a RAID 0 setup across 4 hard drives to raise the overall write speed up to $\mathcal{O}(400)$ MiB/second to allow the data to be stored to disc as quickly as it can be stored in memory, but this too leads to complications. Given the current state of technology, such a high data rate appears to be achievable only through the use of internal SATA drive channels. This effectively limits the total amount of storage to that of the initial disk

[4]Note that the speed benefits of this switch can be partially realized by replacing DRAM with the much more expensive SRAM which can offer data rates about five times that of DRAM. This switch will not, however, give the added bonus of non-volatility.

capacity, $\mathcal{O}(2)$ TiB[5], as once the discs fill up they will need to be removed and replaced with empty discs.

This latter problem can be somewhat ameliorated though the use of external SCSI or Firewire 800 RAID systems that can reach sustained write speeds of $\mathcal{O}(200)$ MiB/second and, with appropriate drivers which would need to be written, could switch between two different external RAID systems when one reaches capacity, allowing adequate time to replace the filled RAID system with a fresh one.

A second approach is to employ two computer systems, each set up with its own 4 drive RAID 0 system. In this case, one system could sit idle until the other nears full drive capacity. At this time, the second computer begins to capture the stream and the first is taken offline, and its drives are replaced so that it can begin to capture the stream again when the second computer nears capacity. The 5000 seconds required to fill the drives on either machine is sufficient for a technician to replace the filled drives with a fresh set. A small amount of redundant data would likely have to be captured at each switch, but this technique would enable the capture of a 400 MiB stream indefinitely, so long as the adversary has enough fresh drives to insert into the system. At this point, cost becomes the limiting factor for the adversary[6], with storage costing $\mathcal{O}(1)$ \$/GiB.

It should be noted that, if the adversary has another set of computer systems available, the contents of each of the drives can be transferred to DVDs in approximately 25000 seconds per drive. Under best case performance conditions, this would allow an adversary with a collection of 22 computers, 20 DVD writers, and 14 TiB of magnetic hard drives to capture a 400 MiB stream indefinitely, while pushing the cost of storage down to that of the DVD media, $\mathcal{O}(0.5)$ \$/GiB. Note that if Blu-Ray DVDs reach their expected maximum write speed, the upfront investment drops to 10 computers, 8 Blu-Ray DVD writers, and 8 TiB of magnetic storage. Holographic storage would drop the investment further to 6 computers, 4 holographic disc writers, and 8 TiB magnetic storage.

4.4.2 An attacker with a cluster of computers

If the attacker has access to a cluster of computers, his ability to capture faster streams is increased. This is because the process of capturing the data is trivially parallelizable, either through the use of a demultiplexer to split the public random stream into smaller portions and distribute them throughout the cluster, or through the use of multiple receivers, one for each node in the cluster[7].

As with the single computer scenario, each node in the cluster is limited to capturing either 2 TiB of data from a 400 MiB/second stream, with a fixed internal RAID 0 system; indefinite amounts of a 200 MiB/second stream with an external RAID 0 system, with enough technicians available to switch the full external raid systems for fresh ones; or if only half the cluster is used to capture data at a time, an indefinite amount of a 400 MiB stream, with enough technicians available to swap out the filled magnetic discs with fresh

[5]Equivalent to capturing 5000 seconds of data from a 400 MiB/second stream.

[6]So long as the stream is slow enough to be captured by the RAID systems.

[7]This method, however, requires the individual nodes to be synchronized with respect to the stream so that each node may capture the appropriate fraction of the bits.

ones.

This means that given a cluster of N computers, an adversary (with a team of technicians) can capture indefinite amounts of a public random stream at data rates up to $200 \cdot N$ MiB/second. So, for example, a moderately sized cluster of 64 nodes would allow an adversary to capture an arbitrary amount of data from a 12.8 GiB/second stream at a cost of $\mathcal{O}(1)$ \$/GiB for storage.

4.4.3 An attacker with a special-purpose multi-stage storage system

Surprisingly, the use of custom hardware buys an adversary little in the bounded storage model. The reason for this is that no matter how well one may improve the flow of data from primary storage to secondary storage on magnetic, optical, or even holographic discs, the overall rate of each individual device is going to be limited to the speed of main memory. If the main memory is DRAM, this leaves an adversary in the same situation as an adversary working with an individual computer[8], albeit a less general purpose computer with a custom designed data bus and controller.

Indeed, it seems as if the only avenue for improvement an adversary with special-purpose hardware has open is the replacement of the DRAM in main memory with the much more expensive SRAM[9] that offers higher data rates. But the nearly 1000 fold increase in cost of memory make such a solution economically nonviable compared to the cluster solutions.

4.5 Summary

Without considering the costs associated with such overhead as space, air conditioning, and power usage, the best an adversary can hope to do in terms of capturing a high-speed public random stream is to employ a cluster of computers to capture portions of the stream in parallel. Given the technology available today or expected in the near future, an adversary with a cluster of N computers and enough technicians to replace filled magnetic discs with fresh ones can capture an arbitrarily large portion of a public random string with a data rate of $200 \cdot N$ MiB/second at a cost of $\mathcal{O}(1)$ \$/GiB stored.

A larger up-front investment of an additional $20 \cdot N$ computers and DVD writers dedicated to transferring data for the magnetic discs to DVDs, along with an additional $12 \cdot N$ TiB of magnetic disc drives can reduce the cost of long term storage to $\mathcal{O}(0.5)$ \$/GiB, which may make financial sense if a large portion of the stream needs to be captured. This investment can be reduced even further upon the release of new technologies, such as Blu-Ray DVDs and holographic storage.

[8] Or an adversary with a cluster of computers if multiple banks of memory are used, each one subsequently passing the data on to an array of secondary storage devices

[9] $\mathcal{O}(100)$ \$/MiB

References

[1] J. Ashley, Maria-Pilar Bernal, G. Burr, H. Coufal, H. Guenther, J. Hoffnagle, C. Jefferson, B. Marcus, R. Macfarlane, R. Shelby, G. Sincerbox, "Holographic data storage technology", *IBM Journal of Research and Development*, 44(3), 341-368, 2000.

[2] Y. Aumann, Y. Z. Ding, M. Rabin, "Everlasting Security in the Bounded Storage Model", *IEEE Trans. on Information Theory*, 48(6), 1668-1680, June 2002.

[3] Y. Aumann, M. Rabin, "Information Theoretically Secure Communication in the Limited Storage Space Model", *CRYPTO 99*, Vol. 1666, 65-79, 1999.

[4] "Blu-ray vs. HD-DVD", `http://www.cdfreaks.com/article/186/1`.

[5] C. Cachin, U. Maurer, "Unconditional Security against Memory-bounded Adversaries", *CRYPTO 97*, Vol. 1294, 292-306, 1997.

[6] B. Chor, O. Goldreich, "Unbiased Bits from Sources of Weak Randomness and Probabilistic Communication Complexity", *SIAM J. on Computing*, 17(2), 230-261, April 1988.

[7] T. Cover, M. Chiang, "Duality Between Channel Capacity and Rate Distortion with Two-Sided State Information", *IEEE Trans. on Information Theory*, Vol. 48(6), 1629-1638, June 2002.

[8] Y. Z. Ding, "Error Correction in the Bounded Storage Model", *Theory of Cryptography TCC 2005*, LNCS 3378, 578-599, 2005.

[9] Y. Z. Ding, M. O. Rabin, "Hyper-encryption and Everlasting Security", *STACS*, Vol. 2285, 1-26, 2002.

[10] Y. Dodis, L. Reyzin, A. Smith, "Fuzzy Extractors and Cryptography, or how to use your Fingerprints", *EUROCRYPT '04*, 2004.

[11] S. Dziembowski, U. Maurer, "Optimal Randomizer Efficiency in the Bounded-Storage Model", *J. of Cryptology*, 17(1), 5-26, 2004.

[12] S. Esener, M. Kryder, W. Doyle, M. Keshner, M. Mansuirpur, D. Thompson, "WTEC Panel Report on The Future of Data Storage Technologies", International Technology Research Institute. World Technology (WTEC) Division, June 1999.

[13] "Everything You Wanted to Know about Blu-ray Disc", `http://www.blu-ray.com`.

[14] FIPS 46-2, "Data Encryption Standard", *Federal Information Processing Standards Publications*, 46-2, U.S. Department of Commerce/National Bureau of Standards, 1993.

[15] S. Gelfand, M. Pinsker, "Coding for Channel with Random Parameters", *IEEE Trans. on Information Theory*, Vol. IT-29, 731-739, Sept. 1983.

[16] C. Heegard and A. El Gamal, "On the Capacity of Computer Memories with Defects", *IEEE Trans. on Information Theory*, Vol. IT-29, 731-739, Sept. 1983.

[17] "IBM, Infineon Develop Most Advanced MRAM Technology to Date", `http://domino.research.ibm.com/comm/pr.nsf/pages/rsc.mramvlsi.html`.

[18] R. Impagliazzo, L.A. Levin, M. Luby, "Pseudorandom Generation from One-Way Functions", *Proc. of the 21st ACM Symposium on Theory of Computation*, 12-24, 1989.

[19] A. Juels, M. Jakobsson, E. Shriver, B. Hillyer, "How to Turn Loaded Dice into Fair Coins", *IEEE Trans. on Information Theory*, 46(3), 911-921, May 2000.

[20] A. Kusnetsov, B. Tsybakov, "Coding in a Memory with Defective Cells", *Probl. Control and Information Theory*, Vol. 10(2), 52-60, 1974.

[21] "The Latest News on HD-DVD and Blu-ray", `http://www.hddvd.org/hddvd/`.

[22] The LOFAR project, `http://www.lofar.org`.

[23] C. Lu, "Encryption against Space-bounded Adversaries from On-line Strong Extractors", *J. of Cryptology*, 17(1), 27-42, 2004.

[24] U. Maurer, "Conditionally-perfect Secrecy and a Provably-secure Randomized Cipher", *J. of Cryptology*, 5(1), 53-66, 1992.

[25] N. Nisan, A. Ta-Shma, "Extracting Randomness: A Survey and New constructions", *Journal of Computer and System Sciences*, 58, 1999.

[26] N. Nisan and D. Zuckerman, "Randomness in Linear in Space", *J. of Computer and System Science*, 52(1),43-52, Feb. 1996.

[27] E. O'Sullivan, "Magnetic Tunnel Junction-Based MRAM and Related Processing Issues", *IBM Research Division*, Research Report RC23525, Feb. 2005.

[28] J. Radhakrishnan, A. Ta-Shma, "Tight Bounds for Depth-two Super-concentrators", *SIAM J. on Discrete Mathematics*, 13(1), 2-24, Feb. 2000.

[29] R. Shaltiel, "Recent Developments in Explicit Constructions of Extractors", *Bulletin of the European Association for Theoretical Computer Science*, 77, 67-95, 2002.

[30] C. Shannon, "Communication Theory of Secrecy Systems", *Bell System Technical Journal*, 28, 656-715, 1949.

[31] M. Santha, U. Vazirani, "Generating Quasi-Random Sequences from Semi-Random Sources", *J. of Computer and System Sciences*, 33(1), 77-87, Aug. 1986.

[32] "Six Companies to form HVD Alliance", http://www.optware.co.jp/english/what_050203.html.

[33] A. Srinivasan, D. Zuckerman, "Computing with very Weak Random Sources", *SIAM J. on Computing*, 28(4), 1433-1459, Aug. 1999.

[34] A. Ta-Shma, "On Extracting Randomness from Weak Random Sources", *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 276-285, 1996.

[35] D. Thompson, J. Best, "The Future of Magnetic Storage Technology", *IBM Journal of Research and Development*, 44(3), 311-322, 2000.

[36] S. P. Vadhan, "Constructing Locally Computable Extractors", *J. Of Cryptology*, 17(1), 43-77, 2004.

[37] U. Vazirani, "Strong Communication Complexity or Generating Quasi-Random Sequences from two Communicating Semi-Random Source", *Combinatorica*, 7, 375-392, 1987.

[38] P. Vettiger, M. Despont, U. Drechsler, U. Dürig, W. Häberle, M. Lutwyche, H. Rothuizen, R. Stutz, R. Widmer, G. Binnig, "The "Millipede"-More than thousand tips for future AFM storage", *IBM Journal of Research and Development*, 44(3), 323-340, 2000.

[39] J. Von Neumann, "Various Techniques used in Connection with Random Digits", *Applied Math Series*, 12, 36-38, 1951.

[40] D. Zuckerman, "General Weak Random Sources", *23rd Annual Symposium on Foundations of Computer Science*, 80-91, Nov. 1982.

[41] D. Zuckerman, "On Unapproximable Versions of NP-Complete Problems", *SIAM J. on Computing*, 25, 1293-1304, 1996.

[42] D. Zuckerman, "Randomness-optimal Oblivious Sampling", *Random Structures and Algorithms*, 11, 345-367, 1997.

DISTRIBUTION:

2 MS 0672
W. E. Anderson, 5614

2 MS 0672
T. J. Draelos, 5614

2 MS 0672
A. Lanzone, 5614

2 MS 0672
W. Neumann, 5614

1 MS 0672
T. McDonald, 5614

1 MS 0671
Gary E. Rivord, 5610

1 MS 0455
S. G. Varnado, 5600

2 MS 9018
Central Technical Files, 8945-1

2 MS 0899
Technical Library, 4536

1 MS 0123
LDRD Program Office, 1011
Attn: Donna Chavez

Sandia National Laboratories