

# Memory Erasability Amplification\*

Jan Camenisch<sup>1</sup>, Robert R. Enderlein<sup>1,2</sup>, and Ueli Maurer<sup>2</sup>

<sup>1</sup> IBM Research – Zurich, Switzerland

<sup>2</sup> Department of Computer Science, ETH Zürich, Switzerland

**Abstract.** Erasable memory is an important resource for designing practical cryptographic protocols that are secure against adaptive attacks. Many practical memory devices, such as solid state drives, hard disks, or file systems, are not perfectly erasable because a deletion operation leaves traces of the deleted data in the system. Various methods for constructing a large erasable memory from a small one, e.g., using encryption, have been proposed. Despite the importance of erasable memory in cryptography, no formal model has been proposed that allows one to formally analyze such memory constructions or cryptographic protocols relying on erasable memory.

The contribution of this paper is three-fold. First, we provide a formal model of erasable memory. A memory device allows a user to store, retrieve, and delete data, and is characterized by a leakage function defining the extent to which erased data is still accessible to an adversary.

Second, we investigate how the erasability of such memories can be amplified. We provide a number of constructions of memories with strong erasability guarantees from memories with weaker guarantees. One of these constructions of perfectly erasable memories from imperfectly erasable ones can be considered as the prototypical application of Canetti et al.’s All-or-Nothing Transform (AoNT). Motivated by this construction, we propose some new and better AoNTs that are either perfectly or computationally secure. These AoNTs are of possible independent interest.

Third, we show (in the constructive cryptography framework) how the construction of erasable memory and its use in cryptographic protocols (for example to achieve adaptive security) can be composed naturally to obtain provable security of the overall protocol.

**Keywords:** secure memory erasure; secure deletion; adaptive corruption; constructive cryptography; all-or-nothing-transforms (AoNT).

---

\* The first and second author were supported by the European Commission through the Seventh Framework Programme under the ERC grant #321310 (PERCY). The third author was supported by the Zurich Information Security & Privacy Center (ZISC). This is the full version of the extended abstract that appeared at SCN 2016.

# Table of Contents

Memory Erasability Amplification .....	1
<i>Jan Camenisch, Robert R. Enderlein, and Ueli Maurer</i>	
1 Introduction .....	3
1.1 Contributions of this Paper .....	3
1.2 Related Work .....	3
2 Preliminaries .....	4
2.1 Notation .....	4
2.2 Constructive Cryptography .....	4
2.3 Cryptographic Building Blocks .....	5
2.4 All-or-Nothing Transform (AoNT) .....	5
3 Modeling Imperfectly Erasable Memory .....	6
3.1 Specification of the General Imperfectly Erasable Memory Resource $M\langle\cdot\rangle$ .....	7
3.2 Instantiations of $M\langle\Sigma, \psi, \rho, \kappa\rangle$ .....	7
4 Constructing Better Memory Resources .....	9
4.1 Admissible Converters for Constructions Using Erasable Memory .....	9
4.2 Memory Erasability Amplification .....	10
4.2.1 Amplifying Memory Leaking Exactly $d$ Symbols .....	10
4.2.2 Amplifying Memory Leaking Symbols with Probability $p$ .....	12
4.2.3 Amplifying Memory with Noisy Leakage .....	13
4.2.4 Amplifying Memory with Limited Leakage Output .....	13
4.3 Constructing a Large Perfectly Erasable Memory from a Small One .....	13
5 New Realizations of All-or-Nothing Transforms .....	16
5.1 AoNT from a Protocol .....	16
5.2 Perfectly Secure AoNT Based on Matrices with Ramp Minimum Distance .....	18
5.3 Realizing a Perfectly Secure AoNT over a Small Field by Combining AoNTs .....	21
5.4 Computationally Secure AoNT over a Large Field from a PRG .....	22
References .....	24
A Known Constructions of AoNT .....	25
A.1 Perfect AoNTs Based on Shamir’s Secret Sharing .....	25
A.2 Perfect AoNTs from Linear Block Codes .....	25
A.3 Statistical AoNTs from Universal Hash Functions .....	26
A.4 Computational AoNTs .....	26

# 1 Introduction

Persistent and erasable memory is a crucial ingredient of many practical cryptographic protocols that are secure against adaptive adversaries. However, for storage devices such as solid state disks, hard disks, and tapes, it is rather difficult to truly erase information written on them. Therefore, constructions have been proposed that use a small amount of memory that is easier to erase (or at least harder for an attacker to tap into), such as smart cards and processor registers, to store a cryptographic key, and then to encrypt the data to be stored on the main medium so that it no longer matters whether or not the ciphertext can be erased [12, 14, 20–22, 24, 25]. This approach is sometimes referred to as crypto paging. Surprisingly, no formal model of erasable memory has been proposed to date, despite the importance of erasable memory for cryptographic protocol design and the cryptographic constructions for it.

## 1.1 Contributions of this Paper

In this paper, we rectify this and first model erasable memory as a general resource in the constructive cryptography framework [17, 18]. Our memory resource defines how a user, an adversary, and the environment can interact with the resource and to what extent stored data can be erased. In particular, different memory resources are characterized in terms of the information about the stored data an adversary will be able to obtain when the environment allows it access to the memory resource. As we discuss, this makes it possible to model many different types of memory, such as hard disks, solid state drives, RAM, and smart cards. Next, we study different constructions of erasable memory from a memory with weaker erasability properties or, in other words, constructions that amplify erasability. These constructions also show how memory resources can be used in protocol design and analysis. We then study the approach of crypto paging in our setting, i.e., constructions of a large erasable memory from a small one and a non-erasable memory. As it turns out, achieving the strongest possible type of erasable memory with this approach requires non-committing encryption and hence is only possible in the random oracle model (else it requires additional communication between sender and receiver, which is not applicable here). We also show which kind of erasable memory can be achieved with this approach in the standard model.

One of our memory constructions uses All-or-Nothing Transforms (AoNT) [5] to obtain a perfectly erasable memory from one that leaks a constant fraction of the erased data. Motivated by this protocol, we study AoNTs and propose several new transforms that achieve better parameters than previously known ones, a result that may be of independent interest. For example, we improve the standard construction of a perfectly-secure AoNT from a Linear Block Code (LBC) by noting that an LBC with a large minimum distance does not yield an AoNT with optimal privacy threshold. Instead, we propose the metric of *ramp minimum distance* and show that LBCs optimized for this metric yield perfectly secure AoNTs with better parameters than what can be achieved with the standard construction. We further propose a computationally secure AoNT that operates over a large alphabet (large enough for one symbol to encode a cryptographic key) and that is optimal: the encoded data is just one symbol longer than the original data, and the transform is secure even if all except one of the symbols of the encoded data are leaked. We show that such an AoNT can be realized from a pseudo-random generator (PRG) with some specific properties.

## 1.2 Related Work

In most security frameworks, unlimited and perfectly erasable memory is available to protocols as part of the framework, with the exception of protocols that are proven to be adaptively secure in the non-erasure model, where no erasable memory is available. However, as mentioned, no security framework explicitly models memory. Consequently, security proofs using the existing

frameworks treat the adversary’s access to the memory of a compromised party only informally. The only exception to this is the work by Canetti et al. and by Lim [6, 16], who model memory as special tapes of the parties’ Turing machines and define how an adversary can access these special tapes. This very specific modeling therefore changes the machine model underlying the UC framework.

Hazay et al. [13] follow a different approach. They introduce the concept of *adaptive security with partial erasures*, where security holds if at least one party of a given protocol can successfully erase. Their model requires a special protocol design and has some restrictions regarding composition.

Both these approaches are rather limited. Indeed, if one wanted to consider different types of memory, one would have to change the modeling framework each time and potentially have to prove all composition theorems all over again. Moreover, these approaches do not allow for the analysis of protocols that construct one type of memory from another type of memory, as we do in this paper. Indeed, one cannot analyze the security of protocols such as Yee’s cryptopaging technique [24, 25] and the constructions of Di Crescenzo et al. [9]. In contrast, we model memory as a resource (or ideal functionality) *within* the security framework (the constructive cryptography framework in our case) and thus do not suffer from these limitations.

## 2 Preliminaries

This section defines the notation used throughout this paper, presents the constructive cryptography model, and recalls various cryptographic building blocks and their security properties.

### 2.1 Notation

Let  $\text{GF}(q)$  denote the Galois field of  $q$  elements, where  $q$  is a prime power. If  $u$  is a vector or a list, let  $u_i$  or  $u[i]$  denote the  $i$ th element of  $u$ . If  $u = (u_1, \dots, u_n)$  and  $v = (v_1, \dots, v_m)$  are lists, then  $(u, e)$  denotes the list  $(u_1, \dots, u_n, e)$  and  $(u, v)$  denotes the list  $(u_1, \dots, u_n, v_1, \dots, v_m)$ ; we write  $e \in u$  to denote  $(\exists i : e = u_i)$ ; we write  $v = (u, \cdot)$  to denote that  $(\forall i \in \{1, \dots, n\} : u_i = v_i)$ . If  $L$  is a set of positive integers, let  $[u]_L$  denote the subvector of  $u$  taken at all positions in  $L$ . If  $S$  is a set, then  $2^S$  denotes the powerset of  $S$  (the set of all subsets of  $S$ ). Let  $\mathbf{I}_r$  denote the identity matrix of size  $r \times r$ , and let  $\mathbf{0}$  denote the zero matrix of appropriate size.

If  $A$  is a deterministic polynomial-time algorithm, then  $y \leftarrow A(x)$  denotes the assignment of variable  $y$  to the output of  $A(x)$ . If  $A$  is a probabilistic polynomial-time (PPT) algorithm, then  $y \xleftarrow{\$} A(x)$  denotes the assignment of  $y$  to the output of  $A(x)$  when run with fresh random coins on input  $x$ . For a set  $\mathbb{A}$ ,  $x \xleftarrow{\$} \mathbb{A}$  denotes the assignment of  $x$  to a value chosen uniformly at random from  $\mathbb{A}$ . For a distribution  $A(x)$ , we denote the ensemble  $\{A(x)\}_{x \in \{1^\eta \mid \eta \in \mathbb{N}, \eta > \eta_0\}}$  by the shorthand  $\{A\}_{1^\eta}$ .

Throughout this paper, we denote the security parameter by  $\eta \in \mathbb{N}$ . Let  $1^\eta$  denote the string consisting of  $\eta$  ones. Unless otherwise noted, all algorithms in this paper are PPT and take  $1^\eta$  as extra (often implicit) input.

### 2.2 Constructive Cryptography

We present our results in the Constructive Cryptography framework [17, 18]. The framework argues about *resources* and how to securely construct a resource from other resources using a protocol that consists of a set of *converters*. Resources and converters are systems that have a set of interfaces. Resources have an interface for each party considered (e.g., *Alice*), one for the adversary (the *Eve* interface), and one for the distinguisher (the *World* interface). The last is an example of what Gaži et al. [11] introduce as a free interface and allows one to model the influence of the distinguisher (environment) on a resource, e.g., to define when a memory becomes readable by the adversary or to model adaptive adversarial behaviour. Converters have only two interfaces, an *inner interface*, which connects to a party interface of a resource, and an




Fig. 1: The constructive statement for a resource with interfaces *Alice*, *Eve*, and *World*. Protocol  $\pi$  constructs  $S$  from  $R$  if there is a simulator  $\sigma$  such that  $R$  with  $\pi$  attached to its *Alice* interface is indistinguishable from resource  $S$  with  $\sigma$  attached to its *Eve* interface (cf. Definition 1).

outer interface, to which a party can connect. A *simulator* is a converter that attaches to the adversary interface of a resource. In this paper, we consider only resources that have a single party interface, i.e., *Alice*. The *security* condition of Constructive Cryptography is as follows (we do not consider the *availability* condition in this paper) [18].

**Definition 1.** A protocol (converter)  $\pi$  constructs resource  $S$  from resource  $R$  with respect to simulator  $\sigma$ , within  $\epsilon$ , denoted

$$R \xrightarrow{\pi, \sigma, \epsilon} S ,$$

if for all distinguishers  $D$  we have  $\Delta^D(\pi^{Alice}R, \sigma^{Eve}S) \leq \epsilon(D)$ , where  $\Delta^D$  is the advantage of  $D$  in distinguishing the two systems [18].

Distinguisher  $D$  is a system itself and has access to all external interfaces of the composition of the resources and converters (cf. Figure 1). With  $\pi^{Alice}R$  we denote the system obtained by attaching the inner interface of  $\pi$  to the interface *Alice* of resource  $R$ , and likewise for  $\sigma^{Eve}$ . In this definition,  $\epsilon$  is a function mapping distinguishers to positive real numbers. Informally, computational security corresponds to  $\pi$  and  $\sigma$  being efficiently implementable and  $\epsilon(D)$  being negligible for all efficiently implementable  $D$ . Constructions are composable, i.e., if  $R \xrightarrow{\pi_1, \sigma_1, \epsilon_1} S$  and  $S \xrightarrow{\pi_2, \sigma_2, \epsilon_2} T$ , then  $R \xrightarrow{\pi_2 \pi_1, \sigma_2 \sigma_1, \epsilon_2 + \epsilon_1} T$ .

### 2.3 Cryptographic Building Blocks

For our constructions, we require pseudo-random generators and exposure-resilient functions, the definitions of which we recall here for convenience.

**Definition 2.** An  $\ell$ -pseudo-random generator (PRG), i.e.,  $\text{prg} : \{0, 1\}^\eta \mapsto \{0, 1\}^{\ell(\eta)}$ , is secure if the following two ensembles are computationally indistinguishable:  $\{\mathbf{b}\}_{1^\eta}$  for  $\mathbf{b} \xleftarrow{\$} \{0, 1\}^{\ell(\eta)}$  and  $\{\text{prg}(\mathbf{a})\}_{1^\eta}$  for  $\mathbf{a} \xleftarrow{\$} \{0, 1\}^\eta$  [15].

**Definition 3.** A  $d$ -exposure-resilient function (ERF)  $\text{erf} : \Phi^n \mapsto \Phi^k$ , also denoted  $(\Phi, n, d, k)$ -ERF, is  $\epsilon$ -secure if for any set  $L \subset \{1, \dots, n\}$  of size at most  $d$ , the following two distributions are  $\epsilon$ -indistinguishable:  $([\mathbf{b}]_L, \mathbf{x}_0)$  for  $\mathbf{b} \xleftarrow{\$} \Phi^n, \mathbf{x}_0 \leftarrow \text{erf}(\mathbf{b})$  and  $([\mathbf{b}]_L, \mathbf{x}_1)$  for  $\mathbf{b} \xleftarrow{\$} \Phi^n, \mathbf{x}_1 \xleftarrow{\$} \Phi^k$  [5].

### 2.4 All-or-Nothing Transform (AoNT)

An all-or-nothing transform (AoNT) [5, 23] is similar to a secret-sharing scheme that requires all shares in order to reconstruct the secret. It consists of two algorithms,  $\text{aenc}$  and  $\text{adec}$ .

**Definition 4.** A  $d$ -AoNT with  $\text{aenc} : \Phi^k \xrightarrow{\$} \Phi^n$  and  $\text{adec} : \Phi^n \mapsto \Phi^k$ , also denoted  $(\Phi, n, d, k)$ -AoNT, is  $\epsilon$ -secure if:

(Completeness) For all messages  $\mathbf{a} \in \Phi^k$ ,  $\mathbf{a} = \text{adec}(\text{aenc}(\mathbf{a}))$ .

(Privacy) For any set  $L \subset \{1, \dots, n\}$  of size at most  $d$  and for any two messages  $\mathbf{a}_0, \mathbf{a}_1 \in \Phi^k$ , the two distributions  $(\mathbf{a}_0, \mathbf{a}_1, [\text{aenc}(\mathbf{a}_0)]_L)$  and  $(\mathbf{a}_0, \mathbf{a}_1, [\text{aenc}(\mathbf{a}_1)]_L)$  are  $\epsilon$ -indistinguishable.

*Computational security.* In the context of computational security, the two functions  $\mathbf{aenc}$  and  $\mathbf{adec}$  take a (usually implicit) security parameter as additional input, and the parameters  $\Phi, n, k$ , and  $d$  may depend on that security parameter. For the privacy condition, it is required that the ensembles  $\{(\mathbf{a}_0, \mathbf{a}_1, [\mathbf{aenc}(1^n, \mathbf{a}_0)]_L)\}_{1^n}$  and  $\{(\mathbf{a}_0, \mathbf{a}_1, [\mathbf{aenc}(1^n, \mathbf{a}_1)]_L)\}_{1^n}$  be indistinguishable. In what follows, we also denote such computationally secure AoNTs as  $(\Phi, n, d, k)$ -AoNTs, where the security parameter is implicit.

*AoNT with public part.* Canetti et al. [5] also consider AoNTs that have a two-part output: a public part that needs no protection and a secret part that has the privacy property. More formally, an  $(\Phi, n + \nu, d, k)$ -AoNT has a  $\nu$ -public part if in the privacy condition above the last  $\nu$  symbols of  $\mathbf{aenc}(\mathbf{a})$  are output in addition to  $[\mathbf{aenc}(\mathbf{a})]_L$ .

*Realisation of AoNTs.* In Appendix A, we summarise a number of known constructions of AoNTs. For instance, one can realise a perfect  $(\Phi, n, d, k)$ -AoNT from any secret sharing scheme over alphabet  $\Phi$  that outputs  $m$  shares, has a reconstruction threshold of  $n$ , a privacy threshold of  $d$ , and encodes messages of size  $k$  shares, by simply ignoring all shares after the first  $n$  ones. This technique also works in the statistical and the computational case.

Given an ERF, it is easy to realise an  $\epsilon$ -secure  $(\Phi, n + k, d, k)$ -AoNT with a  $k$ -public part from any  $\epsilon$ -secure  $(\Phi, n, d, k)$ -ERF:  $\mathbf{aenc}(\mathbf{a}) \mapsto \mathbf{b} \parallel (\mathbf{erf}(\mathbf{b}) + \mathbf{a})$ , where  $\mathbf{b} \leftarrow_{\$} \Phi^n$ ; and  $\mathbf{adec}(\mathbf{b} \parallel \mathbf{x}) \mapsto \mathbf{x} - \mathbf{erf}(\mathbf{b})$  [5]. This technique also works in the computational case.

### 3 Modeling Imperfectly Erasable Memory

We now present our erasable memory resource. Recall that we aim to model memory that is used for persistent storage (such as hard disks, solid state drives, RAM, and smart cards), and not processor registers that store temporary values during computations. To this end, we define how the resource behaves upon inputs on the user (Alice), the adversary, and the world interfaces. It allows a user *Alice* to store a single data item *once*, retrieve it (many times), and erase it. The adversary can get access to the data only if such access is enabled on the *World* interface. That is, the data stored is not initially available to her. Then, once access has been enabled via a *weaken* input on the *World* interface, the adversary can either *read* the data item stored (if the user has not yet deleted it) or *leak* the data, meaning that she will obtain as answer a function of the data once stored. This function determines the information that is still *leaked* although the data has been deleted. The adversary can influence the leakage by providing an additional input to the function (e.g., specify some bits that are leaked).

In reality, there might be many reason why an adversary gains access to the contents of a memory. This might be because the memory device is lost, the adversary controls some malware on the computer that uses the memory, or the adversary runs a cache-timing attack [2] on the computer, etc. Offering a *World* interface, via which it is determined which access is given to the adversary by the memory resource, models any such event. The UC and GNUC frameworks use a similar mechanism for corrupting parties, except that they (ab)use the party interfaces to do so. In UC, it is the adversary who corrupts and the environment is informed of the corruption through the party interfaces. In GNUC, the environment corrupts parties and the adversary is informed thereof.

There seem to be two natural extensions to our erasable memory resource which for simplicity we chose not to consider. First, we assume that inputs at the *World* interface do not impact the user's ability to access the data, which might often not be the case. Although this would be straightforward to model, it is not important for the scope of this paper. Second, the user cannot change the stored data or store many different data items. Again, while it would not be hard to extend the resource to allow that, we chose not to do so for simplicity. Also, this is not a serious restriction as such requirements can be addressed by using several instances of our

<p>The resource <math>M\langle \Sigma, \psi, \rho, \kappa \rangle</math>:</p> <p><b>Internal state and initial values:</b> <math>\text{DATA} = \perp, \text{LDAT} = \perp, \text{HIST} = ()</math>.</p> <p><b>Behavior:</b></p> <ul style="list-style-type: none"> <li>– <math>\text{Alice}(\text{store}, \mu \in \Sigma)</math>: if <math>\text{DATA} = \perp</math>: <math>\text{DATA} \leftarrow \mu; \text{LDAT} \stackrel{\\$}{\leftarrow} \psi(\mu); \text{Alice} \leftarrow ()</math>.</li> <li>– <math>\text{Alice}(\text{retrieve})</math>: if “e” <math>\notin</math> HIST: <math>\text{Alice} \leftarrow \text{DATA}</math>.</li> <li>– <math>\text{Alice}(\text{erase})</math>: if “e” <math>\notin</math> HIST <math>\wedge</math> DATA <math>\neq \perp</math>: HIST <math>\leftarrow</math> (HIST, “e”); <math>\text{Alice} \leftarrow ()</math>.</li> <li>– <math>\text{Eve}(\text{gethist})</math>: <math>\text{Eve} \leftarrow</math> HIST.</li> <li>– <math>\text{Eve}(\text{read})</math>: if <math>\rho(\text{HIST})</math>: <math>\text{Eve} \leftarrow \text{DATA}</math>.</li> <li>– <math>\text{Eve}(\text{leak}, \xi)</math>: if <math>\kappa(\text{HIST}, \xi)</math>: HIST <math>\leftarrow</math> (HIST, “1”  <math>\xi</math>); <math>\text{Eve} \leftarrow \xi(\text{LDAT})</math>.</li> <li>– <math>\text{World}(\text{weaken}, w)</math>: if (“w”  <math>w</math>) <math>\notin</math> HIST: HIST <math>\leftarrow</math> (HIST, “w”  <math>w</math>); <math>\text{World} \leftarrow ()</math>.</li> </ul>
---

Fig. 2: The general (imperfectly) erasable memory resource  $M\langle \cdot \rangle$ .

memory resources.

### 3.1 Specification of the General Imperfectly Erasable Memory Resource $M\langle \cdot \rangle$

We now present our formal specification of the general resource for imperfectly erasable memory  $M\langle \Sigma, \psi, \rho, \kappa \rangle$  that is given in Figure 2 and then discuss in the next subsection a few instantiations of this general resource that match different types of memory. The resource maintains three variables, DATA, LDAT, and HIST. The first one stores the data provided by the user, the second the data that can potentially be leaked to the adversary, and the third one logs the history of events, namely, the erasure event, the parameter of each call on the *World* interface, and the input arguments of each successful leakage query. The resource is parameterized by an alphabet  $\Sigma$ , a conditional probability distribution  $\psi$ , and two predicates  $\rho$  and  $\kappa$ . Alphabet  $\Sigma$  is the set of possible values that can be stored. The conditional distribution  $\psi$  operates on the data and determines which information could potentially leak to the adversary by outputting LDAT. This models the extent to which the resource is able to erase the data. Predicate  $\rho$  takes as input the history of the resource and determines whether the adversary is allowed to read the memory. Finally, predicate  $\kappa$  takes as input the history of the resource and the deterministic function  $\xi$  submitted by the adversary, and determines whether the adversary obtains the leakage  $\xi(\text{LDAT})$ .

Most of the commands that can be submitted to the resource and its behavior should now be clear from Figure 2; however, a few details merit explanation. First, the data that is potentially leaked, LDAT, is determined using  $\psi$  already when the data is stored in the resource. This is without loss of generality, but is useful here because, depending on predicate  $\kappa$ , the adversary may query the resource multiple times with the **leak** command and the answers to these commands need to be consistent. Second, when the adversary queries the resource with a leak command, she can input a parameter  $\xi$  that may influence the leakage she obtains. This models the process of an adversary reading the erased data from a memory device, e.g., an adversary might try to read the data bit by bit, each time influencing the remaining bits in the memory. Third, the adversary is allowed to obtain the history from the resource at any time. This is necessary so that a simulator has enough information to properly simulate a construction. Finally, the *World* interface accepts any value  $w$  for an external event, because these depend on the particular resource that is being modeled and possibly on how it is constructed. This will become clear later when we discuss constructions of one type of memory from other types in Section 4.

### 3.2 Instantiations of $M\langle \Sigma, \psi, \rho, \kappa \rangle$

We now describe special cases of the  $M\langle \Sigma, \psi, \rho, \kappa \rangle$  resource that correspond to memory devices appearing in the real world. We start by describing non-erasable memory, i.e., memory that becomes readable by the adversary once access has been enabled by the *World* interface. This models what happens in a typical file system: files that are unlinked are not actually erased

Perfectly erasable memory	Influence of the <i>World</i> interface
$\text{PM}\langle\Sigma\rangle$	single world event makes memory readable by adversary
$\text{PMW}\langle\Sigma\rangle$	multiple world events are modeled
$\text{PMWa}\langle\Sigma\rangle$	specific version of $\text{PMW}\langle\Sigma\rangle$ (Fig. 6)
$\text{PMWb}\langle\Sigma\rangle$	specific version of $\text{PMW}\langle\Sigma\rangle$ (Fig. 6)
$\text{PMWc}\langle\Sigma\rangle$	specific version of $\text{PMW}\langle\Sigma\rangle$ (Fig. 6)
Imperfectly erasable memory	Information adversary obtains on deleted data
$\text{IM}\langle\Sigma, \psi, \Xi\rangle$	reveals $\xi(\text{LDAT})$ if $\Xi(\xi) = 1$
$\text{IMD}\langle\Phi, n, d\rangle$	reveals $d$ symbols to adversary
$\text{IMDP}\langle\Phi, s_1, s_2, d\rangle$	reveals $d$ symbols of first part, all symbols of second part
$\text{IMI}\langle\Phi, n, d\rangle$	each symbol revealed independently with probability $p$
$\text{IMN}\langle\Phi, n, d\rangle$	reveals through a noisy channel
$\text{IML}\langle\Sigma, v\rangle$	reveals through a length-shrinking function
$\text{IMLP}\langle\Phi, n, a + k, v\rangle$	reveals through a length-shrinking function on first part, full second part

Table 1: Different specializations of  $\text{M}\langle\Sigma, \psi, \rho, \kappa\rangle$  that allow one to erase data.

and can often be completely recovered with specialized tools (at least until the blocks are re-used). We then describe perfectly erasable memory. Such a memory could be implemented by specialized hardware, such as smart cards, but often will have only limited capacity. Large perfectly erasable memories are often not directly available in reality. We are thus interested in the construction of such memories from resources with lesser guarantees. Each of the latter can be influenced through *World* events separately, hence we will describe both a variant of the perfectly erasable memory that accepts a single type of *World* event (easier to describe) and a variant that accepts an arbitrary number of such events. Finally we describe imperfectly erasable memories, i.e., memories with security guarantees between the two extremes just discussed. Such memories leak partial information if the adversary is granted access by *World* after an erasure. In reality, often not all the data is actually removed during an erasure: on magnetic storage, overwritten data can still be partially recovered with specialized equipment [12]. Similarly, often parts of the data were copied to a different medium (swap space, backup, file system journal, etc.) before the erasure and the copies themselves were not fully erased. One can thus easily imagine that the adversary can deduce a constant number of bits that were stored or obtains a noisy version of the data that was stored. For simplicity, we consider imperfectly erasable memories which ignore the parameter of `weaken` (only a single *World* event can be modeled), and only leak once (no adaptive leakage). We now describe these categories of memory in detail. Table 1 provides an overview of these and further specializations that we consider in the sections that follow.

**Non-erasable Memory.** To model non-erasable memory, we let  $\rho$  return true if `weaken` was called irrespective of `erase`. (In fact, the `erase` command could be dropped entirely.) The memory does not leak, hence  $\kappa$  always returns false and  $\psi$  is irrelevant. The only relevant parameter is the alphabet  $\Sigma$  and thus we denote this resource by  $\text{NM}\langle\Sigma\rangle$ .

**Perfectly Erasable Memory.** To model perfectly erasable memory, we let  $\rho$  return true only if `weaken` was called (perhaps multiple times with specific parameters) before `erase` was called.<sup>3</sup> This memory does not leak, hence  $\kappa$  always returns false and  $\psi$  is irrelevant. We describe two

<sup>3</sup> In this paper, we chose to consider monotone  $\rho$ 's. We chose to model the memory resource in such a way that it only responds on the same interface it was activated, hence it is not possible for the adversary to be notified of an event that causes the memory to become readable. To simplify the modeling of simulators, we consider the adversary to be eager and trying to read the memory as soon as possible and then placing the resulting data into an “intermediate buffer,” which can then be collected through the *Eve* interface at a later point in time.



versions of the resource:  $\text{PM}\langle\Sigma\rangle$  fixes  $\rho$  to return true if **weaken** appears in the history earlier than or without **erase**, hence only a single *World* event can be modeled.  $\text{PMW}\langle\Sigma, \rho\rangle$  lets one specify a custom  $\rho$ , allowing the modeling of many *World* events. Figure 6 in the next section shows examples of  $\rho$  in the case when there are two relevant *World* events.

**Imperfectly Erasable Memory.** To model imperfectly erasable memory, we fix  $\rho$  and split  $\kappa$  into two predicates: a fixed predicate that checks only the history and a freely specifiable predicate  $\Xi$  that checks only the adversary’s choice  $\xi$ . The other parameters,  $\Sigma$  and  $\psi$ , can be freely specified. We denote this resource by  $\text{IM}\langle\Sigma, \psi, \Xi\rangle$ . We consider only resources allowing a single *World* event. The predicate  $\rho$  returns true only if the first recorded event in the history is a **weaken** command (as opposed to an **erase** command). The fixed predicate returns true if the first two recorded events in the history are an **erase** command followed by a **weaken** command (if **weaken** was called first, the adversary should call **read** and not **leak**), and no **leak** query succeeded previously. Thus, we consider only resources allowing a single *World* event. The predicate  $\kappa$  returns true if the fixed predicate does so and  $\Xi$  accepts  $\xi$ . In the next section, when we discuss erasability amplification, we specialize this resource further.

## 4 Constructing Better Memory Resources

In this section, we consider constructions of memory resources with stronger security properties from memory resources with weaker ones. We start by showing how to use our memory resources in protocol constructions and then explain the issues that arise when doing so. Thereafter, we describe several specializations of the imperfectly erasable memory resource  $\text{IM}\langle\cdot\rangle$  presented in the previous section and show how to construct memory resources with stronger properties from ones with weaker properties. For example, we show how to construct perfectly erasable memories from memories that leak a certain number of bits. Finally, we consider the construction of a large perfectly erasable memory from a small one plus a large non-erasable memory.

### 4.1 Admissible Converters for Constructions Using Erasable Memory

As stated, one of our reasons to model memory is to be able to analyze cryptographic protocols where the adversary at some point obtains access to the memory. This means that one needs to restrict converters to use only our memory resources for storage. Assuming that an adversary in a real environment may typically not be able to get access to processor registers, we still allow a converter to store temporary values locally and use a memory resource only for persistent storage. Let us now formalize the distinction between persistent and temporary storage and the restrictions we put on converters.

The computation done by a converter is divided in *computation phases*. A phase starts when a converter is activated outside of a computation phase. Informally, a phase ends as soon as the converter responds to that activation or makes a request that is not guaranteed to be answered immediately, i.e., where there is a chance that the adversary is activated before the request completes. For example, a computation phase ends if the converter makes a request that goes over an unreliable communication network, but does not end if the converter asks to store or retrieve data from a memory resource.

In this paper, all our resources always respond on the same interface as the one they were activated on. It is then easy to define the computation phase of a converter: the phase starts as soon as the converter’s outer interface is activated, and stops as soon as the converter writes on its outer interface. That is, activations of the inner interface do not interrupt the phase. However, in a more general setting, resources may respond on a different interface than the one they were activated on, and thereby activate a different party or the adversary. The definition of the computation phase of converters must therefore be adjusted to take this into account.

State that is discarded at the end of a computation phase is *temporary*. State that must persist between two or more computation phases is *persistent*. (Converters must keep all persistent state in memory resources.) This distinction ensures that whenever the adversary has control, the entire internal state of a protocol is inside memory resources, and thus subject to attack.

*Discussion.* Other models, notably the one by Canetti et al. and by Lim [6, 16], also make a distinction between storage needed during computation and persistent storage. However they do it in a way that does not cleanly separate the various layers of abstraction: they assume the existence of a constant number of “processor registers” that are perfectly erasable and place no restriction on the amount of time that data can remain in such a register. For example, their models therefore do not exclude reserving a part of the CPU registers to permanently store a cryptographic key, and use a crypto paging technique [24, 25] to have as much (computationally secure) perfectly erasable memory as required. Thus, to ensure a meaningful analysis, a similar restriction would have to be used in their approach.

## 4.2 Memory Erasability Amplification

We now describe several variants of imperfectly erasable memory that are relevant in practice, namely, memory that leaks a constant number of bits, memory that leaks bits with a certain probability, memory that leaks a noisy version of the data, and memory that leaks the output of a length-shrinking function of the data. We then show how to construct memories that leak less information from each of these variants; in other words, we show how to amplify the erasability of each variant.

### 4.2.1 Amplifying Memory Leaking Exactly $d$ Symbols

On many file systems, unlinked files are not necessarily erased in their entirety immediately. For instance, on most SSDs, deleted data persists until the flash translation layer flashes the corresponding erase block. Furthermore, data may survive erasure if it was copied to another medium, such as cache, swap space, or backups. An adversary could therefore potentially recover parts of data that were believed to be erased. In full generality, the adversary may not obtain the entire data, but still have an influence on which parts of the data she obtains in an attack, e.g., because she can steal just one backup tape, because of the cost of the attack or time constrains force her to choose the most juicy parts of the data, or because she could influence the system beforehand to some degree and ensure that the parts of the data she is interested in were backed up/swapped/cached.

To model such a scenario, we define the memory resource  $\text{IMD}\langle\Phi, n, d\rangle$  storing  $n$  symbols of an alphabet  $\Phi$  and where the adversary can obtain exactly  $d$  symbols of her choice when the memory leaks. This resource is a specialization of  $\text{IM}\langle\Sigma, \psi, \Xi\rangle$ , where  $\Sigma = \Phi^n$ ,  $\psi$  is the identity function, and  $\Xi$  accepts any function that reads at most  $d$  symbols from  $\text{LDAT}$ . In a real setting, depending on the nature of the attack, the adversary may obtain fewer than  $d$  symbols or might not have full control over which symbols she obtains.

A memory resource in such a setting can be perfectly constructed from  $\text{IMD}\langle\Phi, n, d\rangle$  with the identity converter. (A memory resource where the adversary can obtain more than  $d$  symbols with a small probability  $\epsilon$  can also be constructed from  $\text{IMD}\langle\Phi, n, d\rangle$ , albeit with an error probability equal to  $\epsilon$ ; see, e.g., Section 4.2.2.) The converter  $\text{I2P}$  shown in Figure 3 constructs  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$ . This converter is parametrized by an AoNT (cf. Section 2.4). In a nutshell,  $\text{I2P}$  just applies the AoNT encoding algorithm  $\text{aenc}(\cdot)$  to the incoming data before storing it in  $\text{IMD}\langle\cdot\rangle$ , and decodes the encoded data stored in  $\text{IMD}\langle\cdot\rangle$  using  $\text{adec}(\cdot)$  before outputting it. The erasure command is transmitted to  $\text{IMD}\langle\cdot\rangle$  directly. The privacy property of the AoNT guarantees that if the adversary obtains  $d$  symbols of the encoded data, she obtains no meaningful information about the original data. Thus, we obtain the following theorem.

Converter $\text{I2P}\langle\Phi, k, \text{aenc}, \text{adec}\rangle$ :
<b>Behavior:</b> <ul style="list-style-type: none"> <li>– <math>\text{Outer}(\text{store}, \mu \in \Phi^k)</math>: <math>\text{Inner} \stackrel{\\$}{\leftarrow} (\text{store}, \text{aenc}(\mu))</math>. <math>\text{Inner} \rightarrow ()</math>. <math>\text{Outer} \leftarrow ()</math>.</li> <li>– <math>\text{Outer}(\text{retrieve})</math>: <math>\text{Inner} \leftarrow (\text{retrieve})</math>. <math>\text{Inner} \rightarrow \phi</math>.  If <math>\phi \neq ()</math>: <math>\text{Outer} \leftarrow \text{adec}(\phi)</math>. Else: <math>\text{Outer} \leftarrow ()</math>.</li> <li>– <math>\text{Outer}(\text{erase})</math>: <math>\text{Inner} \leftarrow (\text{erase})</math>. <math>\text{Inner} \rightarrow ()</math>. <math>\text{Outer} \leftarrow ()</math>.</li> </ul>

Fig. 3: Converter  $\text{I2P}$  constructing  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$ . The converter is parametrized by a  $(\Phi, n, d, k)$ -AoNT  $(\text{aenc}, \text{adec})$ .

Simulator $\text{SI2P}\langle\Phi, n, d, k, \text{aenc}\rangle$ :
<b>Internal state and initial values:</b> $\text{LEAKED} = 0$ . <b>Behavior:</b> <ul style="list-style-type: none"> <li>– <math>\text{Outer}(\text{gethist})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. <math>\text{Outer} \leftarrow \lambda</math>.</li> <li>– <math>\text{Outer}(\text{read})</math>: upon error in the following, abort with <math>\text{Outer} \leftarrow ()</math>.  <math>\text{Inner} \leftarrow (\text{read})</math>. <math>\text{Inner} \rightarrow \mu \in \Phi^k</math>. <math>\text{Outer} \stackrel{\\$}{\leftarrow} \text{aenc}(\mu)</math>.</li> <li>– <math>\text{Outer}(\text{leak}, \xi \in (x \mapsto [x]_L \text{ with } L \in 2^{\{1, \dots, n\}} \cap \mathbb{N}^d))</math>: upon error in the following, abort with <math>\text{Outer} \leftarrow ()</math>.  If <math>\text{LEAKED} = 0</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>; <math>\text{Inner} \rightarrow (\text{"e"}, \text{"w"})</math>; <math>\text{LEAKED} \leftarrow 1</math>; <math>\text{Outer} \stackrel{\\$}{\leftarrow} \xi(\text{aenc}(0^k))</math>.</li> </ul>

Fig. 4: Simulator  $\text{SI2P}$  for the construction of  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$  using a  $(\Phi, n, d, k)$ -AoNT  $(\text{aenc}, \text{adec})$ .

**Theorem 1.** *If  $(\text{aenc}, \text{adec})$  is an  $\epsilon$ -secure  $(\Phi, n, d, k)$ -AoNT, then*

$$\text{IMD}\langle\Phi, n, d\rangle \xrightarrow{\pi, \sigma, \epsilon} \text{PM}\langle\Phi^k\rangle ,$$

where  $\pi = \text{I2P}\langle\Phi, k, \text{aenc}, \text{adec}\rangle$  (Figure 3) and  $\sigma = \text{SI2P}\langle\Phi, n, d, k, \text{aenc}\rangle$  (Figure 4).

A similar theorem can be stated for the computational case.

*Proof.* We now prove that all distinguishers have at most advantage  $\epsilon$  in distinguishing  $\text{I2P}\langle\Phi, n, \text{aenc}, \text{adec}\rangle\text{IMD}\langle\Phi, n, d\rangle$  (the “real world”) from  $\text{PM}\langle\Phi^k\rangle\text{SI2P}\langle\Phi, k, d, n, \text{aenc}\rangle$  (the “ideal world”).

Consider a system  $W$  that interacts with any distinguisher  $D$  and that behaves like the ideal world except that, instead of outputting  $\xi(\text{aenc}(0^k))$  during a leak,  $W$  plays the AoNT distinguishing game with  $\xi$  and  $0^k$  and the message  $\mu$  that was stored in the memory, obtains  $(0^k, \mu, \omega)$  from that game, and outputs  $\omega$  on *leak*.  $W$  outputs the same value as  $D$ . If  $D$  never triggers a leak,  $W$  outputs a random bit.

It is easy to see that if the leakage output by the AoNT privacy game corresponds to the message  $0^k$ , the behavior of  $W$  is exactly the same as the ideal world for  $D$ , and if the distribution output by the AoNT privacy game corresponds to  $\mu$ , the behavior of  $W$  is exactly the same as the real world for  $D$ .  $W$ ’s advantage in the AoNT distinguishing game is thus not less than the distinguishing advantage of  $D$ .  $\square$

*Multi-part leakage.* It is sometimes the case that the memory is segmented into multiple independent parts, e.g., over two different file systems on different partitions of the same physical disc, and that each part reacts differently to an attack.

We define a multi-part memory resource  $\text{IMDP}\langle\Phi, s_1, s_2, d\rangle$  storing data in  $\Phi^{s_1+s_2}$ . The memory is divided into two parts, the first part consisting of the first  $s_1$  symbols and the second of the other  $s_2$  symbols. The first part of the memory leaks similarly to  $\text{IMD}\langle\Phi, s_1, d\rangle$ , whereas the second one leaks the entire data. When attacking the memory, the adversary must submit the choice of leakage for the first part *before* obtaining the leakage of the second part. We get the following theorem, the proof of which is similar to the one of Theorem 1 and is omitted.

The simulator $\text{SE2R}\langle n, d, p \rangle$ :
<b>Internal state and initial values:</b> $\text{LEAKED} = 0$ .
<b>Behavior:</b>
<ul style="list-style-type: none"> <li>– <math>\text{Outer}(\text{gethist})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. <math>\text{Outer} \leftarrow \lambda</math>.</li> <li>– <math>\text{Outer}(\text{read})</math>: <math>\text{Inner} \leftarrow (\text{read})</math>. <math>\text{Inner} \rightarrow \mu</math>. <math>\text{Outer} \leftarrow \mu</math>.</li> <li>– <math>\text{Outer}(\text{leak})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>.               <ul style="list-style-type: none"> <li>if <math>\text{LEAKED} = 0 \wedge \lambda = (\text{"e"}, \text{"W"})</math>:                   <ul style="list-style-type: none"> <li><math>\text{LEAKED} \leftarrow 1</math>; <math>b \leftarrow 0^n</math>; <math>\omega \leftarrow (\perp, \dots, \perp)</math> with <math> \omega  = n</math>;</li> <li>for <math>i \in \{1, \dots, n\}</math>, set <math>b_i \leftarrow 1</math> with probability <math>p</math>;</li> <li>if <math>\sum_i b_i &gt; d</math>, then abort with <math>\text{Outer} \leftarrow ()</math>;</li> <li><math>\xi \leftarrow \{i \mid b_i = 1\}</math>; add additional indices to <math>\xi</math> until <math>\xi \in (2^{\{1, \dots, n\}} \cap \mathbb{N}^d)</math>;</li> <li><math>\text{Inner} \leftarrow (\text{leak}, \xi)</math>; <math>\text{Inner} \rightarrow \delta</math>;</li> <li>for all <math>i</math> where <math>b_i = 1</math>, set <math>\omega_i</math> to the corresponding value in <math>\delta</math>; <math>\text{Outer} \leftarrow \omega</math>.</li> </ul> </li> </ul> </li> </ul>

Fig. 5: The simulator  $\text{SE2R}$  for the construction of  $\text{IMD}\langle \Phi, n, d \rangle$  from  $\text{IMI}\langle \Phi, n, p \rangle$ .

**Theorem 2.** *If  $(\text{aenc}, \text{adec})$  is an  $\epsilon$ -secure  $(\Phi, s_1 + s_2, d, k)$ -AoNT with public part  $s_2$ , then*

$$\text{IMDP}\langle \Phi, s_1, s_2, d \rangle \xrightarrow{\pi, \sigma, \epsilon} \text{PM}\langle \Phi^k \rangle,$$

where  $\pi = \text{I2P}\langle \Phi, k, \text{aenc}, \text{adec} \rangle$  (Figure 3) and  $\sigma = \text{SI2P}$  (Figure 4).

A similar theorem can be stated for the computational case.

*Choice of alphabet.* The most suitable choice of  $\Phi$  depends on the application. Possible values are  $\text{GF}(2)$  when bits can be leaked independently, e.g., because the adversary must read them one by one from the surface of a disc;  $\text{GF}(2^{512 \cdot 8})$  to  $\text{GF}(2^{4096 \cdot 8})$  when the smallest leakable unit is a file system block; or even  $\text{GF}(2^{128 \cdot 1024 \cdot 8})$  to  $\text{GF}(2^{8192 \cdot 1024 \cdot 8})$  when the smallest leakable unit is an erase blocks of an SSD. In the latter two cases, it is also possible to design the system in such a way that only parts of a block are written to before proceeding with the next one, thereby reducing the alphabet size and limiting the amount of exposure per leaked block.

#### 4.2.2 Amplifying Memory Leaking Symbols with Probability $p$

Above, we modeled an adversary who chooses which symbols leak from the imperfect memory. In practice, the adversary may not have that much power: for example, some parts of a deleted file might still be present in the journal, but the adversary has no control over which ones. To model this, let us now consider an adversary who obtains each symbol of the data uniformly and independently at random with a certain probability  $p$  during a leakage. We denote a memory with such a behavior by  $\text{IMI}\langle \Phi, n, d \rangle$ . This resource is a specialization of  $\text{IM}\langle \Sigma, \psi, \Xi \rangle$  in which  $\Sigma = \Phi^n$ ,  $\psi$  acts like an erasure channel with erasure probability  $(1 - p)$  (i.e., each symbol of the data is transmitted correctly with probability  $p$  and otherwise is replaced with “ $\perp$ ”), and  $\Xi$  accepts only the identity function.

In constructions, one can treat  $\text{IMI}\langle \cdot \rangle$  similarly to  $\text{IMD}\langle \cdot \rangle$ , with just a small statistical error, as the following observation shows. Constructing  $\text{PM}\langle \Phi^k \rangle$  from  $\text{IMI}\langle \Phi, n, p \rangle$  directly without first constructing  $\text{IMD}\langle \Phi, n, d \rangle$  might be more efficient (better parameters, less statistical error), but such a direct construction exceeds the scope of this paper.

**Observation 1** *For all  $(n, d) \in \mathbb{N}^2$ ,  $p \in [0, 1]$ , and fields  $\Phi$ , we have that  $\text{IMI}\langle \Phi, n, p \rangle \xrightarrow{\text{id}, \sigma, \epsilon} \text{IMD}\langle \Phi, n, d \rangle$ , where  $\text{id}$  is the identity converter,  $\sigma = \text{SE2R}$  (Figure 5), and  $\epsilon = (1 - \text{BinomialCDF}(d; n, p)) = \sum_{i=d+1}^n \binom{n}{i} p^i \cdot (1 - p)^{n-i}$ .*

*Proof.* It has to be shown that  $\text{IMI}\langle \Phi, n, p \rangle$  and  $\text{IMD}\langle \Phi, n, d \rangle$  composed with  $\text{SE2R}\langle n, d, p \rangle$  can be distinguished with advantage at most  $(1 - \text{BinomialCDF}(d; n, p))$ .

It is easy to see that if the simulator does not abort, then the simulation is perfect. As the number of bits of  $b$  set to 1 follows a binomial distribution with parameters  $n$  and  $p$ , the probability of an abort is  $(1 - \text{BinomialCDF}(d; n, p))$ . Hence the maximal advantage of any distinguisher is  $(1 - \text{BinomialCDF}(d; n, p))$ .  $\square$

### 4.2.3 Amplifying Memory with Noisy Leakage

Another possible setting is that the data is written to and erased from magnetic storage, and the adversary, who has physical access to the storage medium, must make an educated guess for each bit of the data [12]. One can model this as if the data was transmitted through a noisy binary symmetric channel. We denote such a memory by  $\text{IMN}\langle\Phi, n, d\rangle$ . This resource is a specialization of  $\text{IM}\langle\Sigma, \psi, \Xi\rangle$  in which  $\Sigma = \Phi^n$ ,  $\psi$  acts like a noisy  $|\Phi|$ -ary channel with crossover probability  $(1 - p)/|\Phi|$  (i.e., each symbol of the data is transmitted correctly with probability  $p$  and otherwise is replaced with a symbol drawn uniformly at random from  $\Phi$ ), and  $\Xi$  accepts only the identity function.

**Observation 2** For all  $(n, d) \in \mathbb{N}^2$ ,  $p \in [0, 1]$ , and fields  $\Phi$  we have that  $\text{IMN}\langle\Phi, n, p\rangle \xrightarrow{\pi, \sigma, 0} \text{IMI}\langle\Phi, n, p\rangle$ , where  $\pi = \text{id}$  is the identity converter and  $\sigma$  is the simulator that replaces all erased symbols in the leakage by random symbols.

### 4.2.4 Amplifying Memory with Limited Leakage Output

Another possible setting is that the adversary does not obtain individual symbols of the data but rather a function of the data. For example, with a cache-timing attack [2], she might deduce some information about the data without recovering it completely. In general, one can consider an adversary that obtains any *length-shrinking* function of the contents of the memory. We denote such a memory by  $\text{IML}\langle\Sigma, v\rangle$ . This resource is a specialization of  $\text{IM}\langle\Sigma, \psi, \Xi\rangle$  in which  $\psi$  is the identity function and  $\Xi$  accepts only functions that have at most  $v$  different output values.

For any non-trivial parameters, it is not possible to construct a perfectly erasable memory from  $\text{IML}\langle\cdot\rangle$ , because the adversary can submit a leakage function  $\xi \in \Xi$  that runs the decoding logic of the converter. The reason for this is as follows. Let  $v \geq 2$ ,  $|\Sigma'| \geq 2$ ,  $|\Sigma| \geq 2$ , and let  $\pi$  be a converter that constructs  $\text{PM}\langle\Sigma'\rangle$  from  $\text{IML}\langle\Sigma, v\rangle$ . We now show that this construction has a statistical error of at least  $\frac{1}{2}$ . The distinguisher chooses two distinct messages  $a_0, a_1 \in \Sigma'$ , flips a coin  $b \leftarrow_{\$} \{0, 1\}$ , and stores  $a_b$ . It then makes the memory weak by setting the relevant flags on the *World* interface, and submits a leakage function  $\xi$  that returns 0 iff  $a_0$  was encoded in  $\text{IML}\langle\cdot\rangle$  by using the decoding logic of  $\pi$ —recall that the distinguisher may depend on  $\pi$ . The distinguisher then outputs 1 iff  $\xi$  outputs  $b$ . No simulator will be able to properly simulate that scenario with probability more than  $\frac{1}{2}$  as it does not know whether the distinguisher stored  $a_0$  or  $a_1$ .

However, one can obtain a meaningful construction by starting from a memory resource with multi-part leakage. Let  $\text{IMLP}\langle\Phi, s_1, s_2, v\rangle$  be analogous to  $\text{IMDP}\langle\Phi, s_1, s_2, d\rangle$  defined previously, except that the first part leaks similarly to  $\text{IML}\langle\Phi^{s_1}, v\rangle$ . Here it is crucial to note that the function  $\xi$  submitted by the adversary can read only the first part of the memory. In particular, given a universal hash function  $h : \Phi^a \times \Phi^n \mapsto \Phi^k$ , one can construct the resource  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMLP}\langle\Phi, n, a + k, v\rangle$ , by using  $\text{I2P}$  with an  $\text{AoNT}$  obtained from a universal hash function (see A.3). The construction is  $(2v2^{(k-n)/2})$ -secure [5, 7]. This construction is essentially the one proposed by Canetti et al. [6] and Lim [16].

## 4.3 Constructing a Large Perfectly Erasable Memory from a Small One

We now discuss how a small perfectly erasable memory can be used together with a large, possibly non-erasable memory to construct a large perfectly erasable memory. The basic idea

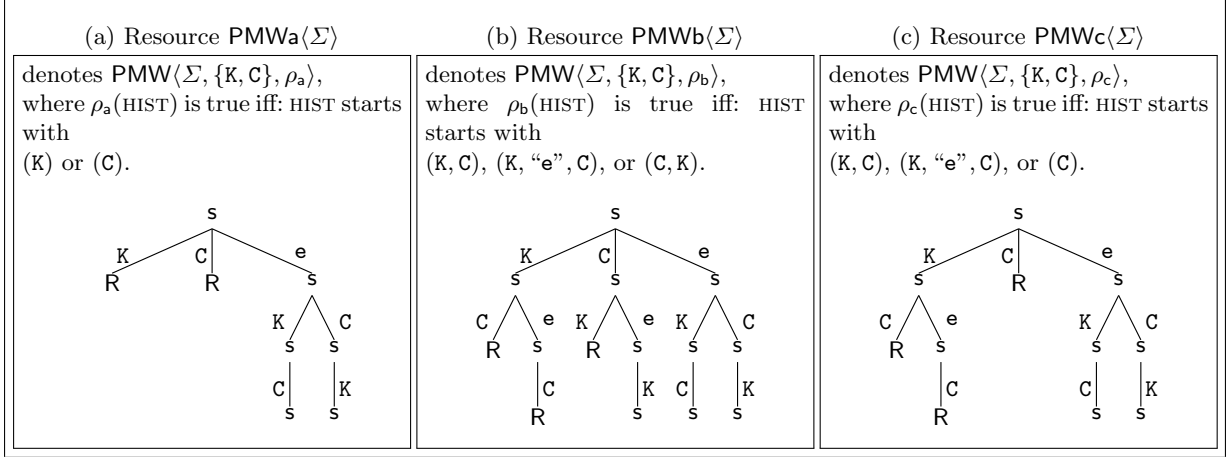


Fig. 6: Several variants of a perfectly erasable memory resource with two *World* flags. The prefix decision trees visualize whether the adversary has read access to the memory depending on the event history HIST. A branch labelled “e” represents an erasure event, and branches labelled “K” (key) or “C” (ciphertext) represent the setting of the corresponding flags on the *World* interface. An “R” node means that the memory is readable (and allows the adversary to collect the data at any time from then on), and an “s” (secure) node means that it does not.




Fig. 7: Converter XPM for constructing a large perfectly erasable memory  $\text{PMWa}\langle\text{GF}(2^{\ell(\eta)})\rangle$  or  $\text{PMWc}\langle\text{GF}(2^{\ell(\eta)})\rangle$  using a small perfectly erasable memory  $\text{PM}\langle\text{GF}(2^\eta)\rangle$  and a large non-erasable memory  $\text{NM}\langle\text{GF}(2^{\ell(\eta)})\rangle$ . The converter is parametrized by an  $\ell$ -PRG  $\text{prg}$ , and the implicit security parameter  $\eta$ .

underlying this construction is that of Yee et al.’s crypto paging [24, 25]: one stores a cryptographic key in the small perfectly erasable memory, encrypts the data with that key, and stores the resulting ciphertext in the large, possibly non-erasable memory. The resulting resource  $\text{PMWa}\langle\text{GF}(2^{\ell(\eta)})\rangle$  will allow the adversary to read the stored data if the resource is weakened by the environment before the user erases the key. The specification of this resource is given in Figure 6a, and the protocol XPM for the construction is provided in Figure 7.

The resource just constructed allows the adversary to read the stored data if either the small erasable *or* the large non-erasable memory becomes weak before the user erases the key. Thus, this resource is weaker than what one would expect, i.e., the adversary should only be able to read the data if *both* underlying resources become weak before the user erases the key. The corresponding resource  $\text{PMWb}\langle\text{GF}(2^{\ell(\eta)})\rangle$  is depicted in Figure 6b. Unfortunately, the realization of this resource would require a non-committing and non-interactive encryption scheme, which can only be constructed in the random oracle model but not in the standard model.

However, it is possible to construct the somewhat better resource  $\text{PMWc}\langle\text{GF}(2^{\ell(\eta)})\rangle$ , shown in Figure 6c. Here the adversary can read the stored data if the memory storing the ciphertext becomes weak before the user calls delete. It is not hard to see that  $\text{PMWc}\langle\text{GF}(2^{\ell(\eta)})\rangle$  implies  $\text{PMWa}\langle\text{GF}(2^{\ell(\eta)})\rangle$ : essentially the simulator attached to the *Eve* interface of  $\text{PMWa}\langle\text{GF}(2^{\ell(\eta)})\rangle$

Simulator $\text{SXPm}\langle\ell, \text{prg}, \rho\rangle$
<b>Internal state and initial values:</b> $\text{CT} = \perp, \text{SK} = \perp$ .
<b>Behavior:</b>
<ul style="list-style-type: none"> <li>– <math>\text{Outer}(\text{PM}, \text{gethist})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. Remove any “C” from <math>\lambda</math>. <math>\text{Outer} \leftarrow \lambda</math>.</li> <li>– <math>\text{Outer}(\text{NM}, \text{gethist})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. Remove any “K” and “e” from <math>\lambda</math>. <math>\text{Outer} \leftarrow \lambda</math>.</li> <li>– <math>\text{Outer}(\text{PM}, \text{read})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. If <math>\lambda</math> does not start with (“K”) nor (“C”, “K”): abort with <math>\text{Outer} \leftarrow ()</math>. If <math>\text{SK} = \perp</math>: <math>\text{SK} \xleftarrow{\\$} \text{GF}(2^n)</math>. <math>\text{Outer} \leftarrow \text{SK}</math>.</li> <li>– <math>\text{Outer}(\text{NM}, \text{read})</math>: <math>\text{Inner} \leftarrow (\text{gethist})</math>. <math>\text{Inner} \rightarrow \lambda</math>. If “C” <math>\notin \lambda</math>: abort with <math>\text{Outer} \leftarrow ()</math>. If <math>\text{SK} = \perp</math>: <math>\text{SK} \xleftarrow{\\$} \text{GF}(2^n)</math>. If <math>\text{CT} = \perp \wedge \rho(\lambda)</math>: <math>\text{Inner} \leftarrow (\text{read})</math>; <math>\text{Inner} \rightarrow \mu</math>; <math>\text{CT} \leftarrow \mu + \text{prg}(\text{SK})</math>. If <math>\text{CT} = \perp</math>: <math>\text{CT} \xleftarrow{\\$} \text{GF}(2^{\ell(n)})</math>. <math>\text{Outer} \leftarrow \text{CT}</math>.</li> </ul>

Fig. 8: Simulator  $\text{SXPm}$  for the proof of the construction of  $\text{PMWc}\langle\text{GF}(2^{\ell(n)})\rangle$  from  $\text{PM}\langle\text{GF}(2^n)\rangle$  and  $\text{NM}\langle\text{GF}(2^{\ell(n)})\rangle$  using converter  $\text{XPM}\langle\ell, \text{prg}\rangle$  and with  $\rho := \rho_c$ . The same simulator with  $\rho := \rho_a$  can be used in the construction of  $\text{PMWa}\langle\text{GF}(2^{\ell(n)})\rangle$ .

has to hold back the leaked data until the non-erasable memory becomes leakable.

**Theorem 3.** *If  $\text{prg}$  is a secure  $\ell$ -PRG, then*

$$[\text{PM}\langle\text{GF}(2^n)\rangle, \text{NM}\langle\text{GF}(2^{\ell(n)})\rangle] \xrightarrow{\pi, \sigma, \epsilon} \text{PMWc}\langle\text{GF}(2^{\ell(n)})\rangle ,$$

where  $\pi = \text{XPM}\langle\ell, \text{prg}\rangle$  (Figure 7),  $\sigma = \text{SXPm}$  (Figure 8), and  $\epsilon$  is a negligible function.

*Proof.* For the sake of contradiction, let us assume the existence of an efficient distinguisher  $D$  that has non-negligible advantage in distinguishing between the “real world”  $\text{XPM}\langle\ell, \text{prg}\rangle[\text{PM}\langle\text{GF}(2^n)\rangle, \text{NM}\langle\text{GF}(2^{\ell(n)})\rangle]$  and the “ideal world”  $\text{PMWc}\langle\text{GF}(2^{\ell(n)})\rangle\text{SXPm}\langle\ell, \text{prg}, \rho_a\rangle$ . (We assume that the difference on the *World* interface is taken care of implicitly.) We show how to construct an efficient distinguisher  $W$  with non-negligible advantage for the PRG distinction game.

$W$  behaves like the ideal world, except that, when asked to leak the non-erasable memory when the “C” *World* flag was set after the data was erased but before the “K” *World*-flag was set (in the sequel, we call this the (“e”, “C”)-event),  $W$  obtains a challenge  $c$  from the PRG distinction game and outputs  $(c + \mu)$ .  $W$  then outputs the same thing as  $D$ . If  $D$  terminates without having provoked the (“e”, “C”)-event,  $W$  outputs a random bit.

Note that  $W$  is constructed in such a way that the following applies:

- If the (“e”, “C”)-event does not happen, then the real and ideal worlds are perfectly indistinguishable.
- If the PRG distinction game outputs  $c = \text{prg}(sk)$ , then  $W$  behaves exactly like the real world to  $D$ .
- If the PRG distinction game outputs a random  $c$ , then  $W$  behaves exactly like the ideal world to  $D$ .

Hence  $WD$  has the same non-negligible advantage in the PRG distinction game as  $D$  has in distinguishing the real and the ideal world.  $\square$

As stated above,  $\text{XPM}$  also constructs  $\text{PMWa}\langle\text{GF}(2^{\ell(n)})\rangle$  from the same resources. Furthermore, in the random oracle model, a protocol that is identical to  $\text{XPM}$  except that calls to

prg are replaced by calls to the random oracle, constructs  $\text{PMWb}\langle\text{GF}(2^{\ell(n)})\rangle$  from the same resources.

Let us discuss now our the memory resources just discussed in the light of some secure memory constructions in the literature. As mentioned, Yee et al. introduce crypto paging [24, 25] to let a secure co-processor encrypt its virtual memory before paging it out to its host’s physical memory or hard disk. Translated to our setting, this means that the non-erasable memory is weak from the beginning. Therefore, to get meaningful guarantees, only the resource  $\text{PMWb}\langle\text{GF}(2^{\ell(n)})\rangle$  can be used in their setting. The other two resources would allow the adversary to always read the data. Thus, to realize their system, Yee et al. require a non-committing and non-interactive encryption scheme (and hence random oracles).

Di Crescenzo et al. [9] consider a memory resource that allows one to update the stored data such that when the resource becomes weak the adversary can only read the data stored last. They then provide a construction for a large such resource from a small one and a large non-erasable memory. Again they assume that for both resources the data can be updated and that the non-erasable one leaks all data ever stored in it. None of our resources allows such updates but, as discussed, resources that allow this can be constructed by using several of our resources in parallel. Thus, their security definition and construction can be indeed modeled and analyzed with the memory resources we define. However, doing this exceeds the scope of this paper.

## 5 New Realizations of All-or-Nothing Transforms

In Section 4 we saw the importance of AoNTs for constructing perfectly erasable memory from certain types of imperfectly erasable ones. In this section, we present several novel AoNTs. We start by showing the dual of the I2P protocol: any protocol that constructs  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$  can be used to realize a  $(\Phi, n, k, d)$ -AoNT. We then present a perfect AoNT with better parameters than what is found in the literature, based on the novel concept of *ramp minimum distance* of a matrix. We next show that one can combine several AoNTs to achieve an AoNT over a small field but with a large message space and a good privacy threshold  $d$ . Finally, we provide a computationally-secure AoNT over a large field that has a very large privacy threshold.

### 5.1 AoNT from a Protocol that Constructs $\text{PM}\langle\Phi^k\rangle$ from $\text{IMD}\langle\Phi, n, d\rangle$

Section 4.2.1 described the protocol I2P, parametrized by an AoNT, that constructs a perfectly erasable memory  $\text{PM}\langle\Phi^k\rangle$  from an imperfectly erasable one  $\text{IMD}\langle\Phi, n, d\rangle$ . As the following theorem states, any protocol  $\pi$  (not necessarily one based on an AoNT) that constructs  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$  can be used to construct an AoNT using algorithm C2A, albeit one where  $\text{adec}$  is a probabilistic algorithm and where decoding might fail with a small probability.

**Theorem 4.** *If  $(\pi, \sigma, \epsilon)$  are such that  $\text{IMD}\langle\Phi, n, d\rangle \xrightarrow{\pi, \sigma, \epsilon} \text{PM}\langle\Phi^k\rangle$ , then the algorithm  $\text{C2A}\langle\Phi, n, k, \pi\rangle$  shown in Figure 9 is a  $6\epsilon$ -secure  $(\Phi, n, d, k)$ -AoNT with a probabilistic  $\text{adec}$  and where decoding may fail with probability less than  $2\epsilon$ .*

One can make an analogous statement in the computational case.

*Proof.* We first consider correctness and then privacy.

*Correctness.* We first show that if there exists a message  $\mu$  such that encoding and decoding fail with probability at least  $2\epsilon$ , then the distinguisher  $D$  shown in Figure 10 distinguishes between the “real world”  $\pi\text{IMD}\langle\Phi, n, d\rangle$  and the “ideal world”  $\text{PM}\langle\Phi^k\rangle\sigma$  with advantage at least  $\epsilon$  for any  $\sigma$ , which would be a contradiction. Let  $Y$  denote the system with which  $D$  is interacting. Let  $x$  be an integer such that  $(\frac{1}{4})^x < \epsilon$ .



Algorithm C2A( $\Phi, n, k, \pi$ )	
<b>Behavior:</b>	<ul style="list-style-type: none"> <li>– <b>aenc</b>(<math>\mu \in \Phi^k</math>): <math>\pi.Outer \leftarrow (\mathbf{store}, \mu)</math>.  While true: <ul style="list-style-type: none"> <li>If <math>\pi.Inner \rightarrow (\mathbf{store}, \phi \in \Phi^n)</math>: return <math>\phi</math>.</li> <li>Else if anything is sent by <math>\pi.Inner</math>: <math>\pi.Inner \leftarrow ()</math>.</li> <li>Else: abort by returning <math>\perp</math>.</li> </ul> </li>   <li>– <b>adec</b>(<math>\phi \in \Phi^n</math>): <math>\pi.Outer \leftarrow (\mathbf{retrieve})</math>.  While true: <ul style="list-style-type: none"> <li>If <math>\pi.Inner \rightarrow (\mathbf{retrieve})</math>: <math>\pi.Inner \leftarrow \phi</math>.</li> <li>Else if <math>\pi.Outer \rightarrow \mu \in \Phi^k</math>: return <math>\mu</math>.</li> <li>Else if <math>\pi.Inner \rightarrow (\mathbf{erase})</math>: abort by returning <math>\perp</math>.</li> <li>Else if anything is sent by <math>\pi.Inner</math>: <math>\pi.Inner \leftarrow ()</math>.</li> <li>Else: abort by returning <math>\perp</math>.</li> </ul> </li> </ul>

Fig. 9: Algorithm C2A realizes a  $(\Phi, n, d, k)$ -AoNT from a converter  $\pi$  which constructs  $\text{PM}\langle\Phi^k\rangle$  from  $\text{IMD}\langle\Phi, n, d\rangle$ .

Distinguisher $D\langle\mu, x\rangle$	
Upon error in any of the following, abort and return 1. <ul style="list-style-type: none"> <li>Do <math>x</math> times: <math>Y.Alice \leftarrow (\mathbf{retrieve})</math>; <math>Y.Alice \rightarrow ()</math>.</li> <li><math>Y.Alice \leftarrow (\mathbf{store}, \mu)</math>. <math>Y.Alice \rightarrow ()</math>.</li> <li>Do <math>x + 1</math> times: <math>Y.Alice \leftarrow (\mathbf{retrieve})</math>; <math>Y.Alice \rightarrow \mu_b</math>.</li> <li><math>Y.Alice \leftarrow (\mathbf{erase})</math>. <math>Y.Alice \rightarrow ()</math>.</li> <li>Do <math>x</math> times: <math>Y.Alice \leftarrow (\mathbf{retrieve})</math>; <math>Y.Alice \rightarrow ()</math>.</li> </ul> Return 0.	

Fig. 10: The distinguisher for the correctness condition of Theorem 4.

Note that  $\sigma$  is never activated, hence  $D$  works for all simulators.

It is clear that if  $D$  interacts with the ideal world,  $D$  always outputs 0. When interacting with the real world, intuitively,  $D$  outputs 1 if encoding or decoding failed, i.e., with probability  $2\epsilon$ . However, one has to take possible “malicious behavior” of  $\pi$  into account: it is possible that  $\text{C2A}\langle\Phi, n, k, \pi\rangle$  fails because  $\pi$  issued an **erase** command during retrieval or because  $\pi$  never stored anything in the memory, but  $D$  outputs 0 anyway.

Recall that  $\pi$  cannot keep state between phases. Hence if  $\pi$  issued an **erase** command during the first retrieval in the second loop,  $\pi$  cannot distinguish between any of the retrieval phases in the second and third loops. Let  $y$  denote the probability that  $\pi$  returns  $\mu$  if faced with an empty memory. The probability that no error happens in  $D$  if the memory was erased during the first retrieve by  $\pi$ , is thus at most  $y^x \cdot (1 - y)^x < (\frac{1}{4})^x < \epsilon$ . A similar argument shows that if  $\pi$  never stored anything in the memory during the **store** command, then  $\pi$  cannot distinguish any of the retrieval phases in the first and second loops. The probability that no error happens is thus also smaller than  $\epsilon$ . Hence the distinguishing advantage is at least  $2\epsilon - \epsilon = \epsilon$ .

*Privacy.* We now show that if there exists a distinguisher  $W$  for the AoNT distinguishing game with advantage at least  $6\epsilon$  on a set  $\xi$  with messages  $\mu_0$  and  $\mu_1$ , then the distinguisher  $D$  shown in Figure 11 that distinguishes between the “real world”  $\pi\text{IMD}\langle\Phi, n, d\rangle$  and the “ideal world”  $\text{PM}\langle\Phi^k\rangle\sigma$  has a distinguishing advantage of at least  $\epsilon$  for any  $\sigma$ , which would be a contradiction. Let  $Y$  denote the system with which  $D$  is interacting. Let  $x$  be an integer such that  $(\frac{1}{4})^x \leq \epsilon$ .

It is clear that  $D$  never aborts in the first part if it interacts with the ideal world (recall that up to now  $\sigma$  was never activated).

It is also clear that if  $D$  interacts with the real world and *if something was stored in the memory and the memory has been erased*  $D$  never aborts in the second part (recall that  $D$  does

<p style="text-align: center;">Distinguisher <math>D(\mu_0, \mu_1, \xi, x, W)</math>:</p> <p>Upon error in the following “first part”, abort and return 1.</p> <p style="padding-left: 20px;">Do <math>x</math> times: <math>Y.Alice \leftarrow (\text{retrieve})</math>; <math>Y.Alice \rightarrow ()</math>.</p> <p style="padding-left: 20px;"><math>b \xleftarrow{\\$} \{0, 1\}</math>. <math>Y.Alice \leftarrow (\text{store}, \mu_b)</math>. <math>Y.Alice \rightarrow ()</math>.</p> <p style="padding-left: 20px;">Do <math>x</math> times: <math>Y.Alice \leftarrow (\text{retrieve})</math>; <math>Y.Alice \rightarrow \mu_b</math>.</p> <p style="padding-left: 20px;"><math>Y.Alice \leftarrow (\text{erase})</math>. <math>Y.Alice \rightarrow ()</math>.</p> <p style="padding-left: 20px;">Do <math>x</math> times: <math>Y.Alice \leftarrow (\text{retrieve})</math>; <math>Y.Alice \rightarrow ()</math>.</p> <p style="padding-left: 20px;"><math>Y.World \leftarrow (\text{weaken})</math>. <math>Y.World \rightarrow ()</math>.</p> <p>Upon error in the following “second part”, abort and return 0.</p> <p style="padding-left: 20px;"><math>Y.Eve \leftarrow (\text{leak}, \xi)</math>. <math>Y.Eve \rightarrow \lambda \in \mathcal{F}^d</math>.</p> <p><math>W \leftarrow (\mu_0, \mu_1, \lambda)</math>. <math>W \rightarrow g</math>. If <math>b = g</math>: return 1. Else: return 0.</p>
---

Fig. 11: The distinguisher for the privacy condition of Theorem 4.

not interact with  $\pi$  in the second part).

Let us calculate the probability that  $\pi$  did not erase the memory and that no error happened during the first part. Recall that  $\pi$  does not keep state between phases. As  $\pi$  does not erase the memory in any of the  $2x$  retrieve queries following the **store** command, the memory behaves identically in all queries, and thus  $\pi$  cannot determine which query number it is currently servicing. Hence for a given  $\mu_b$ , there must be a constant probability  $y$  that  $\pi$  returns  $\mu_b$  in a retrieve phase. Also,  $\pi$  must return  $\mu_b$  as expected in all  $x$  queries of the second loop (probability =  $y^x$ ) and return  $()$  as expected in all  $x$  queries of the third loop (probability at most  $(1 - y)^x$ ). The probability is thus  $(y \cdot (1 - y))^x \leq (\frac{1}{4})^x$ .

If no error occurred in  $D$ , and  $D$  interacted with the ideal world, it is clear that leakage  $\lambda$  is independent of the bit  $b$  chosen, hence  $W$  cannot have any advantage in the AoNT distinction game.

If  $D$  interacted with the real world, and *if the memory contains the value that would have been returned by C2A.aenc*,  $W$  will output the correct guess with probability at least  $\frac{1}{2} + 3\epsilon$ . Similar to the correctness condition, it is possible that  $\pi$  did not store anything in the memory during the **store** command: this happens with probability at most  $(\frac{1}{4})^x$ . Thus if  $D$  interacted with the real world, and no error occurred, then the probability that  $W$  outputs the correct guess will be at least  $\frac{1}{2} + 3\epsilon - (\frac{1}{4})^x$ .

To conclude, the distinguishing advantage of  $D$  is at least  $3\epsilon - 2 \cdot (\frac{1}{4})^x \geq \epsilon$ , as required.  $\square$

## 5.2 Perfectly Secure AoNT Based on Matrices with Ramp Minimum Distance

This subsection shows how one can improve the standard realization of AoNTs based on linear block codes of Canetti et al. [5] by using our novel concept of *ramp minimum distance*.

*The standard realization.* Let  $\mathbf{G}$  be the  $k \times n$  generator matrix with elements in  $\text{GF}(q)$  of a linear block code with minimum distance  $d$ . The encoding function of the perfectly secure  $(\text{GF}(q), (n + k), d, k)$ -AoNT is as follows:

$$\text{aenc}(\mathbf{a} \in \text{GF}(q)^k) : \mathbf{b} \xleftarrow{\$} \text{GF}(q)^n; \mathbf{y} \leftarrow \begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{G} & \mathbf{I}_k \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix}; \text{ return } \mathbf{y}.$$

Further details are given in Appendix A.2.

Let us now show how to use the concept of ramp minimum distance to construct better AoNTs.

**Definition 5.** A  $k \times n$  matrix  $\mathbf{G}$  with elements in  $\text{GF}(q)$  has ramp minimum distance  $d$  if for every  $r \in \{1, \dots, k\}$ , every  $r \times (n - (d - r))$  submatrix of  $\mathbf{G}$  has rank  $r$ .

Note that the concept of (regular) minimum distance comes from coding theory, and requires that all  $k \times (n - (d - 1))$  sub-matrices of  $G$  have rank  $k$  (which is equivalent to saying that

for every  $r \in \{1, \dots, k\}$ , all  $r \times (n - (d - 1))$  sub-matrices of  $\mathbf{G}$  have rank  $r$ , where  $\mathbf{G}$  is the generator matrix of a linear block code. A matrix with minimum distance  $d$  also has a *ramp* minimum distance  $d$  (the converse is obviously not true).

Here is an example of a  $6 \times 24$  matrix over  $\text{GF}(2)$  with Ramp Minimum Distance 12 found using exhaustive search:

$$\mathbf{G} = \begin{bmatrix} 111101101110000010111000 \\ 111111001111010101011111 \\ 011100111010001100001111 \\ 100110010110110011011010 \\ 110001100110100011000111 \\ 111110010100001001100110 \end{bmatrix}.$$

Now, from a generator matrix with ramp minimum distance, we can construct an AoNT and thus obtain the following theorem.

**Theorem 5.** *The standard realization of an AoNT (sketched above and detailed in Appendix A.2), parametrized by a  $k \times n$  matrix  $\mathbf{G}$  with elements in  $\text{GF}(q)$  with ramp (instead of regular) minimum distance  $d$ , is a perfectly secure  $(\text{GF}(q), (n + k), d, k)$ -AoNT.*

*Proof.* We now show that the privacy threshold of the AoNT is at least  $d$ . For any set  $L$  of size  $d$ , let  $r$  denote the number of elements of  $\mathbf{x}$  output by the AoNT distinguishing game (therefore  $d - r$  elements of  $\mathbf{b}$  are output by the game). Let  ${}^k\mathbf{b}$  denote the sub-vector of  $\mathbf{b}$  of size  $d - r$  containing all elements of  $\mathbf{b}$  that are output by the AoNT distinguishing game, and let  ${}^u\mathbf{b}$  denote the sub-vector of size  $n + r - d$  containing the elements that are not output. Similarly, let  ${}^k\mathbf{x}$  denote the sub-vector of  $\mathbf{x}$  of size  $r$  that is output by the AoNT distinguishing game, and let  ${}^u\mathbf{x}$  denote the sub-vector of size  $k - r$  containing the elements that are not output. Let  $\mathbf{P}$  be the permutation matrix such that

$$\mathbf{P} \begin{bmatrix} \mathbf{b} \\ \mathbf{x} \end{bmatrix} = \begin{bmatrix} {}^u\mathbf{b} \\ {}^k\mathbf{b} \\ {}^k\mathbf{x} \\ {}^u\mathbf{x} \end{bmatrix}.$$

Let  ${}^{ku}\mathbf{G}$ ,  ${}^{kk}\mathbf{G}$ ,  ${}^{uu}\mathbf{G}$ ,  ${}^{uk}\mathbf{G}$  be sub-matrices of  $\mathbf{G}$ , and let  ${}^k\mathbf{a}$  and  ${}^u\mathbf{a}$  be the sub-vectors of  $\mathbf{a}$  such that

$$\begin{bmatrix} \mathbf{I}_{n+r-d} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d-r} & \mathbf{0} & \mathbf{0} \\ {}^{ku}\mathbf{G} & {}^{kk}\mathbf{G} & \mathbf{I}_r & \mathbf{0} \\ {}^{uu}\mathbf{G} & {}^{uk}\mathbf{G} & \mathbf{0} & \mathbf{I}_{k-r} \end{bmatrix} := \mathbf{PMP}^\top \quad \text{and} \quad \begin{bmatrix} {}^u\mathbf{b} \\ {}^k\mathbf{b} \\ {}^k\mathbf{a} \\ {}^u\mathbf{a} \end{bmatrix} := \mathbf{P} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix}.$$

We thus have

$$\begin{bmatrix} {}^u\mathbf{b} \\ {}^k\mathbf{b} \\ {}^k\mathbf{x} \\ {}^u\mathbf{x} \end{bmatrix} = \mathbf{P} \begin{bmatrix} \mathbf{b} \\ \mathbf{x} \end{bmatrix} = \mathbf{PM} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} = \mathbf{PMP}^\top \mathbf{P} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_{n+r-d} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d-r} & \mathbf{0} & \mathbf{0} \\ {}^{ku}\mathbf{G} & {}^{kk}\mathbf{G} & \mathbf{I}_r & \mathbf{0} \\ {}^{uu}\mathbf{G} & {}^{uk}\mathbf{G} & \mathbf{0} & \mathbf{I}_{k-r} \end{bmatrix} \begin{bmatrix} {}^u\mathbf{b} \\ {}^k\mathbf{b} \\ {}^k\mathbf{a} \\ {}^u\mathbf{a} \end{bmatrix} = \begin{bmatrix} {}^u\mathbf{b} \\ {}^k\mathbf{b} \\ {}^{ku}\mathbf{G}{}^u\mathbf{b} + {}^{kk}\mathbf{G}{}^k\mathbf{b} + {}^k\mathbf{a} \\ {}^{uu}\mathbf{G}{}^u\mathbf{b} + {}^{uk}\mathbf{G}{}^k\mathbf{b} + {}^u\mathbf{a} \end{bmatrix}.$$

As  $\mathbf{G}$  has ramp minimum distance  $d$ , the  $r \times (n - (d - r))$  sub-matrix  ${}^{ku}\mathbf{G}$  has rank  $r$ , thus  ${}^{ku}\mathbf{G}{}^u\mathbf{b}$  is uniformly distributed. Therefore  $\begin{bmatrix} {}^k\mathbf{b} \\ {}^k\mathbf{x} \end{bmatrix}$ , which is the output of the AoNT distinction game, is uniformly distributed and independent of  $\mathbf{a}$ . This concludes the proof.  $\square$

It remains to find a matrix with a desired ramp minimum distance. One way is to chose a random matrix, as shown by the following theorem.

**Theorem 6.** *For all  $(n, k, d) \in \mathbb{N}^3$ , and all prime powers  $q$ , a  $k \times n$  matrix in which all elements were chosen independently and uniformly at random over  $\text{GF}(q)$ , has ramp minimum distance  $d$  with probability at least*

$$1 - \sum_{i=1}^k \binom{k}{i} (q-1)^i q^{(H_q(\frac{d-i}{n})-1)n} ,$$

$$\text{where } H_q(x) := \begin{cases} 0 & \text{if } x = 0 \text{ or } x = 1; \\ x \log_q(q-1) - x \log_q(x) - (1-x) \log_q(1-x) & \text{if } 0 < x < 1. \end{cases}$$

*Proof.* Let  $\text{HW}(\mathbf{w})$ , the ‘‘Hamming weight,’’ denote the number of non-zero entries of a vector  $\mathbf{w}$ . Let  $\text{HB}_q(n, r)$ , the ‘‘Hamming ball,’’ denote the set of vectors of length  $n$  in  $\text{GF}(q)$  that have a Hamming weight of at most  $r$ .

Further, let  $\text{MC}_q(k)$ , the ‘‘minimal codewords,’’ be the set of row vectors of length  $k$  in  $\text{GF}(q)$  whose leading non-zero entry is 1.

From the definition of the ramp minimum distance, it follows that a  $(\text{GF}(q), n, k)$ -linear block code with generator matrix  $\mathbf{G}$  has ramp minimum distance  $d$  if

$$\forall \mathbf{w} \in \text{MC}_q(k) : \mathbf{w}\mathbf{G} \notin \text{HB}_q(n, d - \text{HW}(\mathbf{w})) .$$

As the coefficients of  $\mathbf{G}$  are chosen independently and uniformly at random, the codeword  $\mathbf{w}\mathbf{G}$  is distributed uniformly. The probability that  $\mathbf{w}\mathbf{G}$  is in some Hamming ball of radius  $r$  thus is

$$\Pr[\mathbf{w}\mathbf{G} \in \text{HB}_q(n, r)] = \frac{|\text{HB}_q(n, r)|}{q^n} = \frac{\sum_{i=0}^r \binom{n}{i} (q-1)^i}{q^n} \leq \frac{q^{H_q(\frac{r}{n}) \cdot n}}{q^n} = q^{(H_q(\frac{r}{n})-1) \cdot n} .$$

Using the union bound (Boole’s inequality), we have

$$\begin{aligned} \Pr[\exists \mathbf{w} \in \text{MC}_q(k) : \mathbf{w}\mathbf{G} \in \text{HB}_q(n, d - \text{HW}(\mathbf{w}))] &\leq \sum_{\mathbf{w} \in \text{MC}_q(k)} \Pr[\mathbf{w}\mathbf{G} \in \text{HB}_q(n, d - \text{HW}(\mathbf{w}))] \\ &\leq \sum_{\mathbf{w} \in \text{MC}_q(k)} q^{(H_q(\frac{d - \text{HW}(\mathbf{w})}{n}) - 1) \cdot n} \\ &= \sum_{i=1}^k \binom{k}{i} (q-1)^i q^{(H_q(\frac{d-i}{n}) - 1) \cdot n} . \end{aligned}$$

The probability that  $\mathbf{G}$  has ramp minimum distance  $d$  is therefore at least

$$\begin{aligned} \Pr[\mathbf{G} \text{ has ramp minimum distance } d] &= \Pr[\forall \mathbf{w} \in \text{MC}_q(k) : \mathbf{w}\mathbf{G} \notin \text{HB}_q(n, d - \text{HW}(\mathbf{w}))] \\ &= 1 - \Pr[\exists \mathbf{w} \in \text{MC}_q(k) : \mathbf{w}\mathbf{G} \in \text{HB}_q(n, d - \text{HW}(\mathbf{w}))] \\ &\geq 1 - \sum_{i=1}^k \binom{k}{i} (q-1)^i q^{(H_q(\frac{d-i}{n}) - 1) \cdot n} . \end{aligned}$$

□

Unfortunately, we do not know any efficient method to check whether a random matrix has a given ramp minimum distance. For practical parameters, however, it is feasible to generate and test such matrices with small values of  $k$  and  $d$  (e.g., smaller than 20).

*Better AoNTs using our realization.* Given a fixed size, it is sometimes possible to find matrices with a given ramp minimum distance, but no matrix with the same (regular) minimum distance. Hence AoNTs based on matrices with a ramp minimum distance can achieve better parameters than previously known realizations. We now illustrate this fact with a numerical example. Let us determine the best message length  $k$  that a perfect AoNT with fixed parameters  $n = 30$ ,  $d = 12$ , and  $q = 2$  can achieve with both our realization and the standard realization. Both realizations will require a matrix with  $(30 - k)$  rows and (ramp or regular, respectively) minimum distance  $d = 12$ . First, observe that there exists a  $6 \times 24$  matrix over  $\text{GF}(2)$  with ramp minimum distance 12 (i.e., the example matrix shown earlier that we have found with exhaustive search). Hence using our realization, we can achieve  $k = 6$ . Plotkin [19] showed that a binary code with block length  $2d$  and distance  $d$  can have at most  $4d$  codewords. Hence there cannot exist a  $6 \times 24$  matrix with (regular) minimum distance  $d = 12$  (as it would generate a code with  $2^6 = 64$  codewords, which is more than  $4d = 48$ ). The best AoNT one can hope for using the standard realization thus has  $k = 5$ .

*Statistical security.* Theorem 6 stated that by choosing a random generator matrix, one can achieve a certain ramp minimum distance with a certain probability  $(1 - \epsilon)$ . If one uses our realization, but without checking that the matrix actually has the required ramp minimum distance, then the resulting AoNT will be perfectly secure with probability  $(1 - \epsilon)$ . (Note that this is different from saying that the AoNT is  $\epsilon$ -secure, as the randomness used to generate the AoNT is not part of the distinguishing experiment.) In practice, one can make  $\epsilon$  very small, e.g.,  $\epsilon < 2^{-\eta}$ , and, to realize an AoNT, it might be acceptable to choose a random matrix and not check its properties.

### 5.3 Realizing a Perfectly Secure AoNT over a Small Field by Combining AoNTs

Designing perfectly-secure AoNTs over very small fields, e.g.,  $\text{GF}(2)$ , is hard. The realization described above does not scale well to large message lengths  $k$  and large privacy thresholds  $d$ ; and realizations based on Shamir’s secret sharing scheme are always over large fields—using such a  $(\text{GF}(2^a), n, d, k)$ -AoNT unmodified over  $\text{GF}(2)$  instead would result in a  $(\text{GF}(2), an, d, ak)$ -AoNT with a poor privacy threshold  $d$ . The leakage of any  $\text{GF}(2)$  element means that the entire original  $\text{GF}(2^a)$  element is compromised. We now show how to combine the two approaches to realize a perfectly secure AoNT over a small field, but with large  $k$  and  $d$ .

Our realization requires two AoNTs, a “fine-grained” one and a “coarse-grained” one, operating over a small field  $S$  and a large field  $L$ , respectively. We require that the number of elements of  $L$  be a power of that of  $S$  and that  $k^s = \log(|L|)/\log(|S|)$  be true. We need to interpret a string of  $k^\ell k^s$  elements from  $S$  as a string of  $k^\ell$  elements of  $L$ , an operation we denote by  $S \triangleright L$ . The converse operation is denoted  $L \triangleright S$ .

The encoding function of our combined AoNT then works as follows. One first applies the coarse-grained AoNT to the whole data vector and then the fine-grained AoNT to each element of the result:

$$\text{aenc}(\mathbf{a} \in S^{k^s k^\ell}) : \mathbf{x} \stackrel{s}{\leftarrow} \text{aenc}^\ell(S \triangleright L(\mathbf{a})); \forall j \in \{1, \dots, n^\ell\} : \mathbf{b}[j] \stackrel{s}{\leftarrow} \text{aenc}^s(L \triangleright S(\mathbf{x}[j])); \text{ return } \mathbf{b}.$$

It is easy to see how the decoding function  $\text{adec}$  of the combined AoNT works and it is thus omitted.

**Theorem 7.** *Given a perfectly secure  $(S, n^s, d^s, k^s)$ -AoNT  $(\text{aenc}^s, \text{adec}^s)$  and a perfectly secure  $(L, n^\ell, d^\ell, k^\ell)$ -AoNT  $(\text{aenc}^\ell, \text{adec}^\ell)$  such that  $k^s = \log(|L|)/\log(|S|)$ , the AoNT  $(\text{aenc}, \text{adec})$  described above is a perfectly secure  $(S, n^s n^\ell, (d^s + 1)(d^\ell + 1) - 1, k^s k^\ell)$ -AoNT.*

*Proof.* We now show that the combined scheme is secure. For any set  $E$  of at most  $(d^s + 1)(d^\ell + 1) - 1$  elements of  $\{1, \dots, k^s k^\ell\}$  and any two challenge messages  $\mathbf{a}_1, \mathbf{a}_2 \in S^{k^s k^\ell}$ , let  $E'$  be the following set:  $E'$  contains the integer  $e \in \{1, \dots, n^\ell\}$  if at least  $(d^s + 1)$  elements of  $\{(e - 1)n^s + 1, \dots, en^s\}$  are contained in  $E$ . Note that  $E'$  has at most  $d^\ell$  elements.

We now show  $2n^\ell + 2$  distributions, where any two consecutive distributions are perfectly indistinguishable. The first and the last distributions correspond to the two possible outputs of the AoNT distinguishing game, hence proving the claim.

Distribution  $i \in \{1, \dots, 2n^\ell + 2\}$  is  $[\mathbf{b}]_E$ , where  $\mathbf{b}$  is calculated as follows:

$$\begin{aligned} \mathbf{x} &\stackrel{\$}{\leftarrow} \begin{cases} \text{aenc}^\ell(\text{S}\triangleright\text{L}(\mathbf{a}_1)) & \text{if } i \leq n^\ell + 1 \\ \text{aenc}^\ell(\text{S}\triangleright\text{L}(\mathbf{a}_2)) & \text{otherwise.} \end{cases} \\ \forall j \in \{1, \dots, n^\ell\} : \mathbf{x}'[j] &\leftarrow \begin{cases} \mathbf{0} & \text{if } j \notin E' \wedge j + 1 \leq i \leq j + 1 + n^\ell \\ \mathbf{x}[j] & \text{otherwise.} \end{cases} \\ \forall j \in \{1, \dots, n^\ell\} : \mathbf{b}[j] &\stackrel{\$}{\leftarrow} \text{aenc}^s(\text{L}\triangleright\text{S}(\mathbf{x}'[j])) . \end{aligned}$$

It is easy to see that any two consecutive distributions except  $(n^\ell + 1)$  to  $(n^\ell + 2)$  are indistinguishable by the security property of the fine-grained scheme and because at most  $d^s$  elements of  $\mathbf{b}$  affect the output. The distributions  $(n^\ell + 1)$  and  $(n^\ell + 2)$  are indistinguishable because of the security property of the coarse-grained scheme and because at most  $d^\ell$  elements of  $\mathbf{x}$  affect the output.  $\square$

*Numerical example.* Let us suppose that we are interested in a perfect AoNT that operates over  $S = \text{GF}(2)$  and that can store a cryptographic key of size  $k = 256$  bits using at most  $n = 8192$  bits (a kilobyte) of memory.

If we use a  $(\text{GF}(2^{10}), 819, 793, 26)$ -AoNT built according to Franklin and Yung [10] unmodified over the field  $\text{GF}(2)$ , we get a  $(\text{GF}(2), 8190, 793, 260)$ -AoNT. This AoNT has a privacy threshold  $d$  of only 793 bits.

By combining a  $(\text{GF}(2), 32, 11, 8)$ -AoNT (which can be found by exhaustive search) with a  $(\text{GF}(2^8), 255, 223, 32)$ -AoNT built according to Franklin and Yung [10], one gets a  $(\text{GF}(2), 8160, 2687, 256)$ -AoNT. This AoNT has a much better privacy threshold  $d$  of 2687, i.e., 2687 arbitrary bits may leak to the adversary.

#### 5.4 Computationally Secure AoNT over a Large Field from a PRG

We now present a realization of a computationally secure AoNTs over a large field  $\text{GF}(2^\eta)$ , where  $\eta$  is the security parameter. Our realization is optimal in the sense that it achieves both an optimal message length  $k = n - 1$  (thus an optimal rate  $(n - 1)/n$ ) and an optimal privacy threshold  $d = n - 1$ . That is, the AoNT needs just a single additional element to encode a message and remains private even if the adversary obtains all except one arbitrary element.

**Definition 6.** An  $\ell$ -PRG where the output length is a multiple of the input length, i.e.,  $\text{prg} : \text{GF}(2^\eta) \mapsto \text{GF}(2^\eta)^{\ell(\eta)/\eta}$ , is KD-secure if for all  $i = 1, \dots, \ell(\eta)/\eta$ , the following ensembles are computationally indistinguishable:

- $\{(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_{\ell(\eta)/\eta})\}_{1^\eta}$ , where  $sk \stackrel{\$}{\leftarrow} \text{GF}(2^\eta)$ ,  $\mathbf{x} \leftarrow \text{prg}(sk)$ , and  $x'_i \leftarrow x_i + sk$ .
- $\{\mathbf{x}\}_{1^\eta}$ , where  $\mathbf{x} \stackrel{\$}{\leftarrow} \text{GF}(2^\eta)^{\ell(\eta)/\eta}$ .

Note that this property is somewhat reminiscent of the KDM-CCA2 security of encryption functions [4].

Our realization, somewhat reminiscent of the OAEP realization of Canetti et al. [5], is as follows:



In Appendix A.4 we observe that Canetti et al.’s computationally-secure AoNT [5] built by combining an exposure resilient function (ERF) with a pseudo-random generator (PRG) can have an essentially arbitrarily high message length  $k$  and message rate  $k/n$ , but cannot achieve a very high privacy threshold  $d$ .

## References

1. C. H. Bennett, G. Brassard, C. Crépeau, and U. M. Maurer. Generalized privacy amplification. *IEEE Transactions on Information Theory*, 41(6):1915–1923, 1995.
2. D. J. Bernstein. Cache-timing attacks on AES. Manuscript, April 2005. <http://cr.yp.to/antiforgery/cachetiming-20050414.pdf>.
3. G. R. Blakley and C. Meadows. Security of ramp schemes. In *CRYPTO’84*, vol. 196 of *LNCS*, pp. 242–268. Springer, Aug. 1984.
4. J. Camenisch, N. Chandran, and V. Shoup. A public key encryption scheme secure against key dependent chosen plaintext and adaptive chosen ciphertext attacks. In *EUROCRYPT 2009*, vol. 5479 of *LNCS*, pp. 351–368. Springer, Apr. 2009.
5. R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-resilient functions and all-or-nothing transforms. In *EUROCRYPT 2000*, vol. 1807 of *LNCS*, pp. 453–469. Springer, May 2000.
6. R. Canetti, D. Eiger, S. Goldwasser, and D.-Y. Lim. How to protect yourself without perfect shredding. In *ICALP 2008, Part II*, vol. 5126 of *LNCS*, pp. 511–523. Springer, July 2008.
7. R. Canetti, D. Eiger, S. Goldwasser, and D.-Y. Lim. How to protect yourself without perfect shredding. Cryptology ePrint Archive, Report 2008/291, 2008. <http://eprint.iacr.org/2008/291>.
8. H. Chen and R. Cramer. Algebraic geometric secret sharing schemes and secure multi-party computations over small fields. In *CRYPTO 2006*, vol. 4117 of *LNCS*, pp. 521–536. Springer, Aug. 2006.
9. G. Di Crescenzo, N. Ferguson, R. Impagliazzo, and M. Jakobsson. How to forget a secret. In *STACS 99*, vol. 1563 of *LNCS*, pp. 500–509. Springer, 1999.
10. M. K. Franklin and M. Yung. Communication complexity of secure computation (extended abstract). In *24th ACM STOC*, pp. 699–710. ACM Press, May 1992.
11. P. Gazi, U. Maurer, and B. Tackmann. Manuscript.
12. P. Gutmann. Secure deletion of data from magnetic and solid-state memory. In *Proceedings of the Sixth USENIX Security Symposium, San Jose, CA*, vol. 14, pp. 77–89, 1996.
13. C. Hazay, Y. Lindell, and A. Patra. Adaptively secure computation with partial erasures. Cryptology ePrint Archive, Report 2015/450, 2015. <http://eprint.iacr.org/2015/450>.
14. S. Jarecki and A. Lysyanskaya. Adaptively secure threshold cryptography: Introducing concurrency, removing erasures. In *EUROCRYPT 2000*, vol. 1807 of *LNCS*, pp. 221–242. Springer, May 2000.
15. J. Katz and Y. Lindell. *Introduction to Modern Cryptography*. CRC Press, 2015.
16. D.-Y. Lim. *The paradigm of partial erasures*. PhD thesis, Massachusetts Institute of Technology, 2008.
17. U. Maurer. Constructive cryptography – a new paradigm for security definitions and proofs. In *Theory of Security and Applications (TOSCA 2011)*, vol. 6993 of *LNCS*, pp. 33–56. Springer-Verlag, April 2011.
18. U. Maurer and R. Renner. Abstract cryptography. In *ICS 2011*, pp. 1–21. Tsinghua University Press, Jan. 2011.
19. M. Plotkin. Binary codes with specified minimum distance. *IRE Transactions on Information Theory*, 6(4):445–450, 1960.
20. J. Reardon, D. A. Basin, and S. Capkun. SoK: Secure data deletion. In *2013 IEEE Symposium on Security and Privacy*, pp. 301–315. IEEE Computer Society Press, May 2013.
21. J. Reardon, S. Capkun, and D. Basin. Data node encrypted file system: Efficient secure deletion for flash memory. In *Proceedings of the 21st USENIX Conference on Security Symposium*, pp. 17–17. USENIX Association, 2012.
22. J. Reardon, H. Ritzdorf, D. A. Basin, and S. Capkun. Secure data deletion from persistent media. In *ACM CCS 13*, pp. 271–284. ACM Press, Nov. 2013.
23. R. L. Rivest. All-or-nothing encryption and the package transform. In *FSE’97*, vol. 1267 of *LNCS*, pp. 210–218. Springer, Jan. 1997.
24. B. Yee. *Using secure coprocessors*. PhD thesis, CMU, 1994.
25. B. Yee and J. D. Tygar. Secure coprocessors in electronic commerce applications. In *Proceedings of The First USENIX Workshop on Electronic Commerce, New York, New York*, pp. 155–170. 1995.



## A Known Constructions of AoNT

In this section we summarize some known constructions of AoNTs that are relevant to this paper.

### A.1 Perfect AoNTs Based on Shamir's Secret Sharing

Blakely and Meadows's secret sharing scheme [3] can be used to directly realize a perfect  $(\Phi, (k+d), d, k)$ -AoNT for all  $(k, d) \in \mathbb{N}^2$ , all fields  $\Phi$ , and where  $(2k+d) < |\Phi|$ . It is based on Shamir's secret sharing scheme.

Franklin and Young's ramp secret sharing scheme [10] can be used to directly realize an AoNT with the same parameters, but the bound on the field size is now improved to  $(k+d) < |\Phi|$ . Their scheme uses polynomial interpolation over  $\text{GF}(|\Phi|)$ . In a nutshell, their scheme works as follows. To encode a message: choose a polynomial of degree  $(k+d-1)$  over  $\text{GF}(|\Phi|)$ . Set the first  $k$  coefficients to be equal to the message, and the other  $d$  coefficients randomly. Evaluate the polynomial at  $(k+d)$  distinct non-zero locations. To decode, use Lagrange interpolation to recover the coefficients of the polynomial.

*Working in small fields.* The two schemes above require the field size  $|\Phi|$  to depend on the parameters  $k$  and  $d$ . If one needs an AoNT that operates on a smaller field  $\Sigma$ , e.g.,  $\text{GF}(2)$ , one can simply encode each element of  $\Phi$  as multiple elements of  $\Sigma$ . Thereby one immediately gets an  $\epsilon$ -secure  $(\Sigma, (\alpha \cdot (k+d)), d, \alpha k)$ -AoNT from an  $\epsilon$ -secure  $(\Phi, (k+d), d, k)$ -AoNT, where  $\alpha = \log(|\Phi|)/\log(|\Sigma|) \in \mathbb{N}$ . Note that the AoNT realized only achieves a privacy threshold of  $d$  and not  $\alpha d$ : intuitively, if parts of a symbol are leaked, one must consider the whole symbol to be compromised.

There exist more complex secret sharing schemes that can be used to realize AoNTs with better parameters, such as Chen and Cramer's ramp secret sharing scheme [8] based on curves of high genus. We do not consider their results further in this paper.

### A.2 Perfect AoNTs from Linear Block Codes

Linear block codes can be used to create perfect ERFs and thus, by using the standard transformation by Canetti et al. [5], can be used to create perfect AoNTs.

Let  $\mathbf{G}$  be the  $k \times n$  matrix with elements in  $\text{GF}(q)$  with minimum distance  $d$  (i.e.,  $\mathbf{G}$  is the generator matrix of a linear block code with minimum distance  $d$ ).

Let  $\mathbf{M}$  be the following  $(n+k) \times (n+k)$  matrix:

$$\mathbf{M} := \begin{bmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{G} & \mathbf{I}_k \end{bmatrix}.$$

To encode the data column-vector  $\mathbf{a} \in \text{GF}(q)^k$ ,  $\text{aenc}(\mathbf{a})$  selects a random column-vector  $\mathbf{b} \leftarrow \text{GF}(q)^n$ , and returns the vector

$$\mathbf{y} \leftarrow \mathbf{M} \begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} = \begin{bmatrix} \mathbf{b} \\ \mathbf{Gb} + \mathbf{a} \end{bmatrix} := \begin{bmatrix} \mathbf{b} \\ \mathbf{x} \end{bmatrix}.$$

To reconstruct the data,  $\text{adec}(\mathbf{y})$  computes

$$\begin{bmatrix} \mathbf{b} \\ \mathbf{a} \end{bmatrix} \leftarrow \mathbf{M}^{-1} \mathbf{y} = \begin{bmatrix} \mathbf{b} \\ \mathbf{x} - \mathbf{Gb} \end{bmatrix}$$

and outputs  $\mathbf{a}$ . Hence this AoNT is a perfect  $(\Phi, (n+k), d, k)$ -AoNT.

### A.3 Statistical AoNTs from Universal Hash Functions

Universal hash functions can be used to create very good statistical ERFs, and hence (using the standard transformation) very good AoNTs [1, 5, 6, 16]. Given a  $\{0, 1\}^\nu \times \{0, 1\}^n \mapsto \{0, 1\}^k$  universal hash function  $h$ , one can realize the following  $(\{0, 1\}, n + (\nu + k), d, k)$ -AoNT with a  $(\nu + k)$ -public part:

- $\text{aenc}(\mathbf{a})$  : Choose  $\mathbf{b} \xleftarrow{\$} \{0, 1\}^n$  and  $\mathbf{k} \xleftarrow{\$} \{0, 1\}^\nu$ . Set  $\mathbf{x} \leftarrow h(\mathbf{k}, \mathbf{b}) + \mathbf{a}$ . Return  $\mathbf{b}||\mathbf{k}||\mathbf{x}$ . (Here  $\mathbf{k}$  and  $\mathbf{x}$  are in the public part.)
- $\text{adec}(\mathbf{b}||\mathbf{k}||\mathbf{x}) \mapsto \mathbf{x} - h(\mathbf{k}, \mathbf{b})$ .

The AoNT is  $(2 \cdot 2^{(k+d-n)/2})$ -secure [5, 7].

### A.4 Computational AoNTs

Canetti et al. [5] showed how to stretch the output of an ERF with a PRG to obtain an ERF with larger output size (but the privacy threshold will remain unchanged). Let  $\text{erf}$  be a computationally secure  $(\Phi, n, d, k)$ -ERF and let  $\text{prg}$  be a secure  $(\Phi, k, m)$ -PRG. Then the ERF:  $\mathbf{b} \mapsto \text{prg}(\text{erf}(\mathbf{b}))$  is a computationally secure  $(\Phi, n, d, m)$ -ERF. From that, one can realize a computationally secure AoNT.