

# Towards a Theory of Data Entanglement <sup>\*,\*\*</sup>

James Aspnes<sup>1</sup>,

*Department of Computer Science, Yale University*

Joan Feigenbaum<sup>2</sup>,

*Department of Computer Science, Yale University*

Aleksandr Yampolskiy<sup>3</sup>,

*Oracle Corporation*

Sheng Zhong<sup>4</sup>,

*Department of Computer Science and Engineering, SUNY Buffalo*

---

## Abstract

We give a formal model for systems that store data in entangled form. We propose a new notion of entanglement, called **all-or-nothing integrity** (AONI) that binds the users' data in a way that makes it hard to corrupt the data of any one user without corrupting the data of all users. AONI can be a useful defense against negligent or dishonest storage providers who might otherwise be tempted to discard documents belonging to users without much clout. We show that, if all users use a fixed standard recovery algorithm, we can implement AONI using a MAC, but, if some of the users adopt instead a non-standard recovery algorithm provided by the dishonest storage provider, AONI can no longer be achieved. However, even for the

latter scenario, we describe a simple entangling mechanism that provides AONI for a restricted class of **destructive adversaries**.

*Key words:* data entanglement, untrusted storage, all-or-nothing integrity, upgrade attacks.

---

## 1 Introduction

Suppose that I provide you with remote storage for your most valuable information. I may advertise various desirable properties of my service: underground disk farms protected from nuclear attack, daily backups to chiseled granite monuments, replication to thousands of sites scattered across the globe. But what assurance do you have that I will not maliciously delete your data as soon as your subscription check clears?

If I consider deleting the data of a rich or powerful customer, I may be deterred by economic, social, or legal repercussions. The small secret joy I might

---

\* An extended abstract of this paper appeared in the proceedings of 9th European Symposium on Research in Computer Security.

\*\*This work was supported by the DoD University Research Initiative (URI) administered by the Office of Naval Research under Grant N00014-01-1-0795.

*Email addresses:* `aspnes@cs.yale.edu` (James Aspnes), `jf@cs.yale.edu` (Joan Feigenbaum), `alex.yampolskiy@oracle.com` (Aleksandr Yampolskiy), `szhong@cse.buffalo.edu` (Sheng Zhong).

<sup>1</sup> Supported in part by NSF grants CCR-0098078, CNS-0305258, and CNS-0435201.

<sup>2</sup> Supported in part by ONR grant N00014-01-1-0795 and NSF grants ITR-0219018 and ITR-0331548.

<sup>3</sup> Supported by NSF grants CCR-0098078 and ANI-0207399.

<sup>4</sup> Supported by NSF grant CCR-0208972.

experience from the thought of the loss will not compensate me for losing a posted bond, destroying my reputation, or being imprisoned. But if you are an ordinary customer who does not have much clout, and I see a lucrative opportunity in altering—or simply neglecting to keep—your data, then deterrence loses its effectiveness. Consequently, data of powerful customers end up being more protected than data of average customers. To convince an average customer that she will not lose her data at my random whim, I might offer stronger technical guarantees that I cannot destroy her data without serious costs. One way to do this would be to link the fate of her documents to the documents of enough other users that I cannot hope to offend them all with impunity. We shall call such documents **entangled**.

Data entanglement was initially suggested as a mechanism for increasing censorship resistance in document-storage systems, *e.g.*, Dagster [20] and Tangler [14]. These systems split data into blocks in such a way that a single block becomes part of several documents. New documents are represented using some number of existing blocks, chosen randomly from the pool, combined with new blocks created using exclusive-or (Dagster) or 3-out-of-4 secret sharing [18] (Tangler). Dagster and Tangler use entanglement as one of many mechanisms to prevent a censor from tampering with unpopular data; others involve disguising the ownership and contents of documents and (in Tangler) storing documents redundantly.

It is not clear that data entanglement is actually useful for censorship resistance. Instead of having to specifically attack a target document, a censor only needs to damage any document entangled with the target to achieve his goal. Instead, we consider data entanglement for a different purpose: protecting the data from an untrusted storage provider that might be tempted to damage or

destroy the data through negligence or malice. Entanglement provides an incentive for the storage provider to take extra care in protecting average users' documents by increasing the cost of errors.

We begin in Section 2 by analyzing the intuitive notion of entanglement provided by Dagster and Tangler. We show that entanglement as provided by Dagster and Tangler is not by itself sufficiently strong to deter a dishonest storage provider from tampering with data, because not enough documents get deleted on average when destroying a block of a typical document. This motivates our efforts to obtain a stronger form of an entanglement than the ones provided by these systems.

In Section 3, we define our general model of entanglement in an untrusted storage system. Our goal here is to model the entanglement operation itself, and we do not address the question of where in the system entanglement occurs. However, we do assume that the storage provider does *not* carry out the entangling operation itself, as giving it the users' raw data would allow it to store copies that it could selectively return later, even if the entangled store were lost or damaged. Instead, some trusted third party is assumed to carry out the entangling operation, and a negligent or malicious storage provider is modeled separately as a "tamperer" that has access only to the entangled store.

Section 4 contains our definitions of **document dependency**, where a document cannot be recovered if any document it depends on is lost, and **all-or-nothing integrity**, where no document can be recovered if any document is lost. These definitions allow a system-independent description of the binding between entangled documents. We then consider how different levels of attacks

on the common data store do or do not prevent enforcement of document dependency or all-or-nothing integrity.

In particular, we show that, if all clients use a standard algorithm to recover their data, then all-or-nothing integrity requires only the ability to detect tampering using a MAC (Section 5.1); in this model, the standard recovery algorithm is too polite to return any user’s data if any other user’s data have been lost. Relying on such fastidiousness provides only a weak guarantee; what we really want is that all data become irretrievable even by non-standard algorithms if any is lost. We show that this goal is impossible if an adversary is allowed to both tamper with the common store arbitrarily and provide a replacement recovery algorithm (Section 5.2). Despite such **upgrade attacks**, it is still possible to provide a weaker guarantee that we call **symmetric recovery**, in which each document is equally likely to be destroyed (Section 5.3). Furthermore, if we restrict the adversary to **destructive tampering**, which reduces the amount of information in the common store, all-or-nothing guarantees are possible even with upgrade attacks (Section 5.4).

These results provide a first step toward understanding document dependency. We discuss the strengths and limitations of our approach and offer suggestions for future work in Section 6.

### *1.1 Related Work*

Entanglement is motivated by the goal of deterring data tampering by untrusted servers, a more general problem that has been studied extensively. Entanglement has been used specifically in Dagster [20] and Tangler [14], as

we describe in Section 2. Other approaches to preventing or deterring tampering include replication across global networks of tamper-resistant servers [1, 3, 4, 8, 16, 21] or tamper detection [5–7, 11–13, 19] using digital signatures and Merkle hash trees [15]. Replication protects against data loss if a small number of servers are compromised; tamper detection prevents data loss from going unnoticed. Both techniques complement the entanglement approach considered here.

All-or-nothing integrity as defined in the present work is related to the guarantee provided by the **all-or-nothing transform** proposed by Rivest [17]. An all-or-nothing transform is an invertible transform that guarantees that no bits of the preimage can be recovered if  $\ell$  bits of the image are lost. All-or-nothing transforms are not directly applicable to our problem, because we consider the more general case in which the image may be corrupted in other ways, such as by superencryption<sup>5</sup> or alteration of part or all of the common store.

## 2 Dagster and Tangler

In Sections 2.1 and 2.2, we review how Dagster [20] and Tangler [14] work, concentrating on their operations at a block level. For simplicity, we assume that each user’s document is small enough to fit into a single server block and needs not be split up into pieces. In Section 2.3, we analyze the resulting entangling effects and show their shortcomings for protecting data from a

---

<sup>5</sup> In a **superencryption attack**, the adversary encrypts the data store under some key and then blackmails the users into paying him money in exchange for the key. We discuss different kinds of adversarial attacks in Section 3.2.

negligent storage provider. This analysis motivates us to introduce stronger notions of entanglement in Section 4.

## 2.1 *Dagster*

The Dagster storage system may run on a single server or on a P2P overlay network. Each document in Dagster consists of  $c + 1$  server blocks:  $c$  blocks of older documents and one new block, an exclusive-or of previous blocks with the document. The  $c + 1$  blocks that must be XORed to recover the document are listed in a **Dagster Resource Locator** (DRL). The storage protocol proceeds as follows:

**Initialization** Upon startup, the server creates approximately 1,000 random server blocks. Each user  $i$  generates a private key  $k_i$ .

**Entanglement** To publish document  $d_i$ , user  $i$  chooses  $c$  random server blocks  $C_{i_1}, \dots, C_{i_c}$  and computes a new block

$$C_i = \mathcal{E}_{k_i}(d_i) \oplus \bigoplus_{j=1 \dots c} C_{i_j},$$

where  $\mathcal{E}$  is a secure encryption function and  $\oplus$  is bitwise exclusive-or. Hashes of blocks needed to reconstruct  $d_i$  are recorded in a DRL:

$$\left( k_i, \pi [H(C_i), H(C_{i_1}), H(C_{i_2}), \dots, H(C_{i_c})] \right).$$

Here  $H(\cdot)$  is a cryptographic hash function, and  $\pi$  is a random permutation.

**Recovery** To recover  $d_i$ , the user asks the server for blocks in  $d_i$ 's DRL. If hashes of the blocks returned by the server match the ones in the DRL, the user computes:

$$\mathcal{D}_{k_i} \left( C_i \oplus \bigoplus_{j=1 \dots c} C_{i_j} \right),$$

where  $\mathcal{D}$  is a decryption function. Otherwise, the user exits.

## 2.2 Tangler

The Tangler storage system uses (3, 4) Shamir secret sharing [18] to entangle the data: Each document is represented by four server blocks, any three of which are sufficient to reconstruct the original document. The blocks get replicated across a subset of Tangler servers. Hashes of blocks are recorded in a data structure, similar to Dagster Resource Locator, called an **inode**. Here is the storage protocol:

**Initialization** As in Dagster, the server is jump-started with a bunch of random blocks.

**Entanglement** Each server block in Tangler is a pair  $(x, y)$ , where  $x \in GF(2^{16})$ , and  $y$  is the value of a polynomial at  $x$ . To publish  $d_i$ , user  $i$  downloads two random server blocks,  $C_{i_1}$  and  $C_{i_2}$ , and interpolates them with  $(0, d_i)$  to form a quadratic polynomial  $f(\cdot)$ . He evaluates  $f(\cdot)$  at two random (non-zero) integers to get new blocks,  $C'_{i_1}$  and  $C'_{i_2}$ , uploads the new blocks, and records the hashes of blocks comprising  $d_i$  (*viz.*,  $C_{i_1}, C_{i_2}, C'_{i_1}, C'_{i_2}$ ) in the inode.

**Recovery** To recover his document, user  $i$  sends a request for blocks in  $d_i$ 's inode to a subset of Tangler servers. Upon receiving three of  $d_i$ 's blocks, the user can reconstruct  $f(\cdot)$  and compute  $d_i = f(0)$ .



### 2.3 Analysis of Entanglement

At a given point in time, a Dagster or Tangler server contains a set of blocks  $\{C_1, \dots, C_m\}$  comprising documents  $\{d_1, \dots, d_n\}$  of a set of users. (Here  $m, n \in \mathbb{N}$  and  $m \geq n$ .) Data are partitioned in a way that each block becomes a part of several documents. We can draw an **entanglement graph** (see Figure 1), which has an edge  $(d_j, C_k)$  if block  $C_k$  belongs to document  $d_j$ . This connection is rather tenuous—even if  $(d_j, C_k)$  is in the graph, it may still be possible to reconstruct  $d_j$  from blocks excluding  $C_k$ . Document nodes in Dagster’s entanglement graph have out-degree  $c + 1$ , and those in Tangler’s have out-degree 4. Entangled documents share one or more server blocks. In Figure 1, documents  $d_1$  and  $d_n$  are entangled because they share server block  $C_1$ ; meanwhile, documents  $d_1$  and  $d_2$  are not entangled.

This shared-block notion of entanglement has several drawbacks. Even if document  $d_j$  is entangled with a specific document, it may still be possible to delete  $d_j$  from the server without affecting that particular document. For example, knowing that  $d_n$  is entangled with  $d_1$  (as in Figure 1) and that  $d_1$  is owned by some Very Important Person may give solace to the owner of  $d_n$ , who might assume that no adversary would dare incur the wrath of the VIP merely to destroy  $d_n$ . But in the situation depicted in the figure, the adversary can still delete server blocks  $C_2$  and  $C_m$  corrupt  $d_n$  but not  $d_1$ .

The resulting dependence between documents is thus very weak. We now derive an upper bound on how many documents get destroyed if we delete a random document from a Dagster or Tangler server. One might expect that the earlier the document was uploaded to the server, the more documents




Fig. 1. An entanglement graph is a bipartite graph from the set of documents to the set of server blocks. Edge  $(d_j, C_k)$  is in the graph if server block  $C_k$  can be used to reconstruct document  $d_j$ .

it will be entangled with and the more other documents will be destroyed. This turns out to be the case, as we prove for any specific target document in Lemmas 1 and 2.

### 2.3.1 Deleting a Targeted Document

We consider a restricted adversary, who randomly chooses several blocks of an arbitrary document (one block in Dagster; two in Tangler) and overwrites them with zeroes. How many other documents will it destroy?

Without loss of generality, we assume that documents are numbered in the order in which they were uploaded; namely, for all  $1 \leq j < n$ , document  $d_j$  was uploaded to the server before  $d_{j+1}$ . Then, we have:

**Lemma 1** *In a Dagster server with  $n_0 = O(1)$  initial blocks and  $n$  documents, where each document is linked with  $c$  pre-existing blocks, deleting a random block of document  $d_j$  ( $1 \leq j \leq n$ ) destroys on average*

$$O\left(c \log\left(\frac{n}{j}\right)\right)$$

*other documents.*

**PROOF.** Altogether, there are  $n_0 + n$  blocks stored on the server:  $n_0$  initial blocks and  $n$  data blocks. We label the data blocks  $C_1, \dots, C_n$ . The initial blocks exist on the server before any data blocks have been added. We label them  $C_{-n_0+1}, \dots, C_0$ .

Every document  $d_j$  consists of  $c$  pre-existing or “old” blocks<sup>6</sup> and a “new” data block  $C_j$  that is computed during the entanglement stage. Consider an

<sup>6</sup> These may be either initial blocks or data blocks of documents added earlier than

adversary who destroys a random block  $C_i$  of  $d_j$ . This will destroy  $d_j$ , but it will also destroy any documents with outgoing edges to  $C_i$  in the entanglement graph. We would like to compute the number of such documents,  $N_i$ .

If  $C_i$  is a data block (*i.e.*,  $i \geq 1$ ), then

$$\begin{aligned}
E[N_i] &= \sum_{k=i}^n \Pr[d_k \text{ has an edge to } C_i] \\
&= 1 + \sum_{k=i+1}^n \left( 1 - \binom{k-2+n_0}{c} / \binom{k-1+n_0}{c} \right) \\
&= 1 + c \sum_{j=i+n_0}^{n+n_0-1} \frac{1}{j} \\
&= 1 + c(H_{n+n_0-1} - H_{i+n_0-1}) \\
&= 1 + O\left(c \log\left(\frac{n}{i}\right)\right) \text{ under the assumption that } n_0 \text{ is a constant.}
\end{aligned}$$

Meanwhile, if  $C_i$  is an initial block (*i.e.*,  $i < 1$ ), it can be linked by any of the documents:

$$\begin{aligned}
E[N_i] &= \sum_{k=1}^n \Pr[d_k \text{ has an edge to } C_i] \\
&= O(c \log n).
\end{aligned}$$

The number of documents deleted on average when the adversary destroys a *random* block of  $d_j$  is

$$\begin{aligned}
N_{avg} &= \frac{1}{j+n_0} \left( \sum_{i=-n_0}^j E[N_i] \right) \\
&< \frac{1}{j} \left( O(c \log n) + \sum_{i=1}^j O\left(c \log\left(\frac{n}{i}\right)\right) \right) \tag{1}
\end{aligned}$$

We use Stirling's formula to bound the leading term in (1):

$\overline{d_j}$  (*i.e.*,  $d_k$  with  $k < j$ ).

$$\begin{aligned}
\sum_{i=1}^j O\left(c \log\left(\frac{n}{i}\right)\right) &= O\left(c \log\left(\prod_{i=1}^j \frac{n}{i}\right)\right) \\
&= O\left(cj \log\left(\frac{n}{j/e}\right) - \frac{1}{2}c \log(2\pi j)\right) \\
&= O\left(cj \log\left(\frac{n}{j}\right)\right).
\end{aligned}$$

The lemma follows.  $\square$

We also have:

**Lemma 2** *In a Tangler server with  $n_0 = O(1)$  initial blocks and  $n$  documents, deleting two random blocks of document  $d_j$  ( $1 \leq j \leq n$ ) destroys on average*

$$O\left(\frac{1}{j}\right)$$

*other documents.*

**PROOF.** The server contains  $n_0 + 2n$  blocks:  $n_0$  initial blocks and  $2n$  data blocks. We label the blocks as in the proof of Theorem 3. The initial blocks are  $C_{-n_0+1}, \dots, C_0$ , and the data blocks are  $C_1, C_2, \dots, C_{2n-1}, C_{2n}$ .

In Tangler, every document  $d_j$  consists of two *old* blocks of pre-existing documents and two *new* blocks  $C_{2j-1}$  and  $C_{2j}$ , computed during the entanglement stage. Suppose an adversary deletes any two out of four blocks comprising  $d_j$ ; call these blocks  $C_i, C_t$ . Then any document  $d_k$  ( $k \neq j$ ) that contains both  $C_i$  and  $C_t$  (*viz.*, has edges outgoing to  $C_i$  and  $C_t$  in the entanglement graph), will also get destroyed. We would like to compute the number of such documents,  $N_{avg}$ .

In our analysis, we consider whether deleted blocks  $C_i, C_t$  are new or old to  $d_j$  and  $d_k$ . We consider five cases:

**Case 1:**  $C_i, C_t$  are old to both  $d_j$  and  $d_k$ . For  $k < j$ , block  $k$  is lost if  $j$  chooses from the  $2j - 2 + n_0$  old blocks available precisely the two old blocks of  $k$ . For  $k > j$ , the same occurs if  $k$  chooses precisely the two old blocks of  $j$ . Summing over all  $k$ , the expected number of deleted documents is

$$\sum_{k=1}^{j-1} \frac{1}{\binom{2j-2+n_0}{2}} + \sum_{k=j+1}^n \frac{1}{\binom{2k-2+n_0}{2}}.$$

**Case 2:**  $C_i, C_t$  are old to  $d_j$ . However, only one of them is old to  $d_k$ , while the other is new to  $d_k$ . In this case, we must have  $k < j$ . There are four ways  $j$  can choose an old and a new block of  $k$ , giving

$$\sum_{k=1}^{j-1} \frac{4}{\binom{2j-2+n_0}{2}}.$$

**Case 3:**  $C_i, C_t$  are old to  $d_j$ , but new to  $d_k$ . In this case, we must also have  $k < j$ . The expected loss is

$$\sum_{k=1}^{j-1} \frac{1}{\binom{2j-2+n_0}{2}}.$$

**Case 4:** One of  $C_i, C_t$  is old to  $d_j$ , the other is new to  $d_j$ . Then both  $C_i$  and  $C_t$  must be old to  $d_k$  (because otherwise we would have  $k < j$ , which implies that the block in  $C_i, C_t$  that is new to  $d_j$  is not linked to  $d_k$ , which further implies that  $d_k$  will not get deleted). As in case 2, there are four ways to choose one old and one new block for  $j$ , giving

$$\sum_{k=j+1}^n \frac{4}{\binom{2k-2+n_0}{2}}.$$

**Case 5:**  $C_i, C_t$  are new to  $d_j$ . In this case, both  $C_i$  and  $C_t$  must be old to  $d_k$  (for the same reason as in Case 4). The contribution of this case is

$$\sum_{k=j+1}^n \frac{1}{\binom{2k-2+n_0}{2}}.$$

Summing up the five cases gives us the total number of documents destroyed:

$$\begin{aligned} N_{avg} &= \frac{6(j-1)}{\binom{2j-2+n_0}{2}} + \sum_{k=j+1}^n \frac{6}{\binom{2k-2+n_0}{2}} \\ &< \frac{6}{2j-3} + \sum_{k=j+1}^n \frac{3}{(k-\frac{3}{2})^2} \\ &< \frac{6}{2j-3} + \int_j^n \frac{3}{(x-\frac{3}{2})^2} dx \\ &= \frac{12}{2j-3} - \frac{6}{2n-3} \\ &= O\left(\frac{1}{j}\right) \text{ for large } n. \end{aligned}$$

□

### 2.3.2 Deleting a Random Document

The effects of deleting a document chosen uniformly at random, stated in Theorems 3 and 4, are computed by averaging bounds (given in Lemmas 1 and 2) over all documents:

**Theorem 3** *In a Dagster server with  $n$  documents, where each document is linked with  $c$  pre-existing blocks, deleting a block of a random document destroys on average  $O(c)$  other documents.*

**PROOF.** By Lemma 1, deleting document  $d_j$  ( $1 \leq j \leq n$ ), affects  $O(c \log(n/j))$  other documents. Averaging over all  $j$ , we get:

$$\begin{aligned}
\frac{1}{n} \sum_{j=1}^n O\left(c \log\left(\frac{n}{j}\right)\right) &= O\left(\frac{c}{n} \log\left(\prod_{j=1}^n \frac{n}{j}\right)\right) \\
&= O\left(\frac{c}{n} \log\left(\frac{n^n}{(n/e)^n}\right)\right) \\
&= O(c)
\end{aligned}$$

as claimed.  $\square$

One might expect the theorem to follow immediately from the fact that each document depends on  $c$  blocks, which means (when  $n$  is large enough that the number of blocks and documents are proportional) that the average block is used by  $\Theta(c)$  documents. While such an argument does tell us what happens if the adversary deletes a block chosen uniformly at random, choosing a document uniformly at random and then choosing a block from that document biases the choice of block toward earlier blocks that are used in more documents; nonetheless, Theorem 3 shows that this bias provides at most a constant increase in the destructive effect.

Meanwhile, Lemma 2 tells us that removing document  $d_j$  from the Tangler server affects  $O(1/j)$  other documents on average. It is immediate that

**Theorem 4** *In a Tangler server with  $n$  documents, deleting two blocks of a random document destroys on average  $O\left(\frac{\log n}{n}\right)$  other documents.*

Even a small chance of destroying an important document will deter tampering to some extent, but some tamperers might be willing to run that risk. Still more troubling is the possibility that the tamperer might first flood the system with junk documents so that almost all real documents were entangled only with junk. Since our bounds show that destruction of a typical document will on average affect only a handful of others in Dagster and almost none in



Tangler, we will need stronger entanglement mechanisms if entanglement is to deter tampering by itself.

### 3 Our Model

In Section 3.1, we start by giving a basic framework for systems that entangle data. Specializing the general framework gives specific system models, differentiated by the choice of recovery algorithms and restrictions placed on the adversary. We discuss them in Section 3.2.

Our model abstracts away many details of storage and recovery processes. It concentrates on a single entanglement operation that takes documents of a finite set of users and intertwines these documents to form a common store. In practice, the server contents would be computed as an aggregation of common stores from multiple entanglement operations. We defer analyzing this more complex case to later work; see the discussion of possible extensions to the model in Section 6.

#### 3.1 Basic Framework

The model consists of an **initialization phase**, in which keys are generated and distributed to the various participants in the system; an **entanglement phase**, in which the individual users' data are combined into a common store; a **tampering phase**, in which the adversary corrupts the store; and a **recovery phase**, in which the users attempt to retrieve their data from the corrupted store using one or more recovery algorithms. For simplicity of notation, we number the users  $\{1, \dots, n\}$ , where every user  $i$  possesses a document

$d_i$  that he wants to publish.

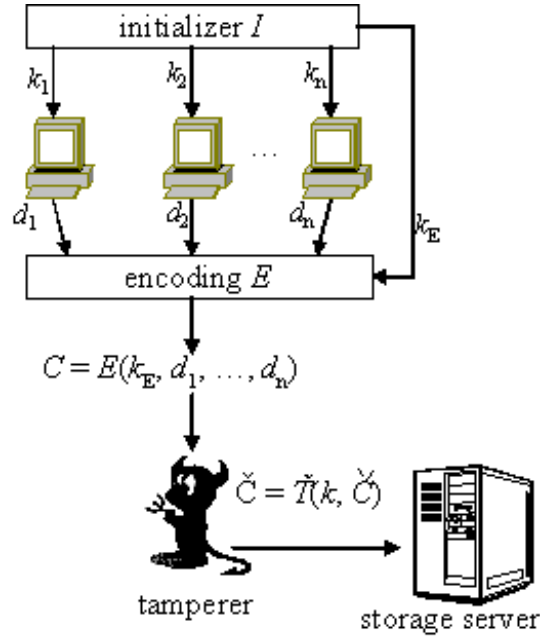


Fig. 2. Initialization, entanglement, and tampering stages.

An encoding scheme consists of three probabilistic Turing machines  $(I, E, R)$  that run in time polynomial in the size of their inputs and a security parameter  $s$ . The first of these, the **initialization algorithm**  $I$ , hands out the keys used in the encoding and recovery phases. The second, the **encoding algorithm**  $E$ , combines the users' data into a common store using the encoding key. The third, the **recovery algorithm**  $R$ , attempts to recover each user's data using the appropriate recovery key. The initialization algorithm takes as input the number of users  $n$  and a security parameter  $s$ . We assume that the number of users and the total size of users' data are polynomial in  $s$ .

Acting against the encoding scheme is an adversary  $(\check{I}, \check{T}, \check{R})$ , which also consists of three probabilistic polynomial-time Turing machines. The first is an **adversary-initialization algorithm**  $\check{I}$ ; like the good initializer  $I$ , the evil  $\check{I}$  is responsible for generating keys used by other parts of the adversary during

the protocol. The second is a **tampering algorithm**  $\check{T}$  that modifies the common store. The third is a **non-standard recovery algorithm**  $\check{R}$  that may be used by some or all of the users to recover their data from the modified store.

We assume that  $\check{I}$ ,  $\check{T}$  and  $\check{R}$  are chosen after  $I$ ,  $E$ , and  $R$  are known but that a fixed tripe  $(\check{I}, \check{T}, \text{and } \check{R})$  is used for arbitrarily large values of  $s$  and  $n$ . This is necessary for polynomial-time bounds on  $\check{T}$  and  $\check{R}$  to have any effect.

Given an encoding scheme  $(I, E, R)$  and an adversary  $(\check{I}, \check{T}, \check{R})$ , the **storage protocol** proceeds as follows (see also Figure 2):

- (1) **Initialization.** The initializer  $I$  generates a combining key  $k_E$  used by the encoding algorithm and recovery keys  $k_1, k_2, \dots, k_n$ , where each key  $k_i$  is used by the recovery algorithm to recover the data for user  $i$ . At the same time, the adversary initializer  $\check{I}$  generates the shared key  $\check{k}$  for  $\check{T}$  and  $\check{R}$ .

$$k_E, k_1, k_2, \dots, k_n \leftarrow I(1^s, n),$$

$$\check{k} \leftarrow \check{I}(1^s, n).$$

- (2) **Entanglement.** The encoding algorithm  $E$  computes the combined store  $C$  from the combining key  $k_E$  and the data  $d_i$ :

$$C \leftarrow E(k_E, d_1, d_2, \dots, d_n).$$

- (3) **Tampering.** The tamperer  $\check{T}$  transforms the combined store  $C$  into  $\check{C}$ :

$$\check{C} \leftarrow \check{T}(\check{k}, C).$$

- (4) **Recovery.** The users attempt to recover their data. User  $i$  applies his

recovery algorithm  $R_i$  to  $k_i$  and the changed store  $\check{C}$ . Each  $R_i$  could be either the standard recovery algorithm  $R$ , supplied with the encoding scheme, or the non-standard algorithm  $\check{R}$ , supplied by the adversary, depending on the choice of the model.

$$d'_i \leftarrow R_i(k_i, \check{C})$$

We say that user  $i$  **recovers** his data if the output of  $R_i$  equals  $d_i$ .

### 3.2 Adversary Classes

In our model, the adversary corrupts the data store and distributes a new recovery algorithm, which users may or may not use. It cannot otherwise affect the operation of the storage server. The model is thus divided on two axes, one bounding the users' choices of reconstruction algorithms and the other bounding the adversary's power to modify the data store. With respect to recovery algorithms, we consider three variants on the basic framework (listed in order of increasing power given to the adversary):

- In the **standard-recovery-algorithm model**, the users are restricted to a single standard recovery algorithm  $R$ , supplied by the system designer. Formally, this means  $R_i = R$  for all users  $i$ ; the adversary's recovery algorithm  $\check{R}$  is not used. This is the model used to analyze Dagster and Tangler.
- In the **public-recovery-algorithm model**, the adversary not only modifies the combined store, but also supplies a single non-standard recovery algorithm  $\check{R}$  to all of the users. Formally, we have  $R_i = \check{R}$  for each  $i$ . The original recovery algorithm  $R$  is not used.<sup>7</sup>

---

<sup>7</sup> Though it may seem unreasonable to prevent users from choosing the original

We call this an **upgrade attack** by analogy to the real-life situation of a company changing the data format of documents processed by its software and distributing a new version of the software to read them. We believe such an attack is a realistic possibility, because most self-interested users will be happy to adopt a new recovery algorithm if it offers new features or performance or if the alternative is losing their data.

- In the **private-recovery-algorithm model**, the adversary may choose to supply the non-standard recovery algorithm  $\check{R}$  to only a subset of the users. The rest continue to use the standard algorithm  $R$ . Formally, this model is a mix of the previous two models:  $R_i = R$  for some  $i$  and  $R_i = \check{R}$  for others.

We also differentiate between two types of tamperers:

- An **arbitrary tamperer** can freely corrupt the data store and is not restricted in any way. Most real-life systems fit into this category, because they place no restrictions on the tamperer.
- A **destructive tamperer** can only apply to the store a transformation whose range of possible outputs is substantially smaller than the set of inputs. The destructive tamperer can superimpose its own encryption on the common store, transform the store in arbitrary ways, and even add additional data, provided that the cumulative effect of all these operations is to decrease the entropy of the data store. Though a destructive tampering assumption may look like an artificial restriction, it subsumes natural models of block deletion or corruption, and either it or some similar assumption is needed to achieve all-or-nothing integrity in the private-recovery-algorithm

---

recovery algorithm  $R$ , any  $R$  can be rendered useless in practice by superencrypting the data store and distributing the decryption key only with the adversary's  $\check{R}$ . We discuss this issue further in Section 5.2.

model.

An **adversary class** specifies what kind of tamperer  $\check{T}$  is and which users, if any, receive  $\check{R}$  as their recovery algorithm. Altogether, we consider six adversary classes, each corresponding to a combination of constraints on the tamperer and the recovery algorithms.

## 4 Dependency and All-Or-Nothing Integrity

We now give our definition of document dependency for a particular encoding scheme and adversary class. We first discuss some basic definitions and assumptions in Section 4.1. Our strong notions of entanglement, called **dependency** and **all-or-nothing integrity**, are defined formally in Section 4.2.

### 4.1 Preliminaries

Because we consider protocols involving probabilistic Turing machines, we must be careful in talking about probabilities. Fix an encoding  $(I, E, R)$ , an adversary  $A = (\check{I}, \check{T}, \check{R})$ , and the recovery algorithm  $R_i$  for each user  $i$ . An **execution** of the resulting system specifies the inputs  $k_i$  and  $d_i$  to  $E$ , the output of  $E$ , the tamperer's input  $\check{k}$  and output  $\check{C}$ , and the output of the recovery algorithm  $R_i$  ( $R(k_i, \check{C})$  or  $\check{R}(\check{k}, k_i, \check{C})$  as appropriate) for each user. The set of possible executions of the storage system is assigned probabilities in the obvious way: the probability of an execution is taken over the inputs to the storage system and the coin tosses of the encoding scheme and the adversary. It will be convenient to consider multiple adversaries with a fixed encoding scheme. In this case, we use  $\Pr_A(Q)$  to denote the probability that

an event  $Q$  occurs when  $A$  is the adversary.

During an execution of the storage system, the tamperer alters the combined store from  $C$  into  $\check{C}$ . As a result, some users end up recovering their documents while others do not. We summarize which users recover their documents in a **recovery vector**, which is a vector-valued random variable  $\vec{r}$  in which  $r_i = 1$  if  $R_i(k_i, \check{C}) = d_i$  (i.e., if user  $i$  recovers his document) and 0 otherwise. For example, if the server contains documents  $d_1, d_2$ , and  $d_3$  and in an execution we recover only  $d_1$  and  $d_2$ , then  $\vec{r} = 110$ .

#### 4.2 Our Notions of Entanglement

In Section 2, we observed that the block-sharing notion of entanglement provided by Dagster and Tangler is not adequate for our purposes. This motivates us to propose the notion of **document dependency**, which formalizes the idea that “if my data depends on yours, I can’t get my data back if you can’t.” In this way, the fates of specific documents become linked together: specifically, if document  $d_i$  depends on document  $d_j$ , then whenever  $d_j$  cannot be recovered neither can  $d_i$ .

Given just one execution, either users  $i$  and  $j$  each get their data back or they don’t. So how can we say that the particular outcome for  $i$  depends on the outcome for  $j$ ? Essentially, we are saying that we are happy with executions in which either  $j$  recovers its data (whether or not  $i$  does) or in which  $j$  does not recover its data and  $i$  does not either. Executions in which  $j$  does not recover its data but  $i$  does are bad executions in this sense. We will try to exclude these bad executions by saying that they either never occur or occur

only with very low probability. Formally:

**Definition 5** A document  $d_i$  **depends on** a document  $d_j$  with respect to a class of adversaries  $\mathcal{A}$ , denoted  $d_i \stackrel{\mathcal{A}}{\hookrightarrow} d_j$ , if, for all adversaries  $A \in \mathcal{A}$ ,

$$\Pr_A [(r_i = 1) \Rightarrow (r_j = 1)] \geq 1 - \epsilon.$$

Here and after,  $\epsilon$  refers to a negligible function of the security parameter  $s$ .<sup>8</sup>

The ultimate form of dependency is **all-or-nothing integrity**. Intuitively, a storage system is all-or-nothing if either every user  $i$  recovers his data or no user does:

**Definition 6** A storage system is **all-or-nothing** with respect to a class of adversaries  $\mathcal{A}$  if, for all  $A \in \mathcal{A}$ ,

$$\Pr_A [\vec{r} = 0^n \vee \vec{r} = 1^n] \geq 1 - \epsilon.$$

It is easy to show that

**Theorem 7** A storage system is all-or-nothing with respect to a class of adversaries  $\mathcal{A}$  if and only if, for all users  $i, j$ ,  $d_i \stackrel{\mathcal{A}}{\hookrightarrow} d_j$ .

**PROOF.** Fix an adversary in  $\mathcal{A}$ . Let  $E$  be the event that an execution of the storage system is not all-or-nothing, and  $F_{ij}$  the event that document  $d_i$  was recovered in an execution and  $d_j$  was not. Then  $E = \{\vec{r} \neq 0^n \wedge \vec{r} \neq 1^n\}$  and  $F_{ij} = \{r_i = 1 \wedge r_j = 0\}$ .

<sup>8</sup> A function  $\epsilon : \mathbb{N} \mapsto (0, 1)$  is **negligible** if for every  $c > 0$ , for all sufficiently large  $s$ ,  $\epsilon(s) < 1/s^c$ . See any standard reference, such as [9], for details.



( $\Rightarrow$ ) : If the system is all-or-nothing, then  $\Pr[E] < \epsilon$ . Clearly, for all  $i, j$ , we have  $F_{ij} \subseteq E$ , which means  $\Pr[F_{ij}] \leq \Pr[E] < \epsilon$ . This in turn implies  $d_i \stackrel{\mathcal{A}}{\hookrightarrow} d_j$ .

( $\Leftarrow$ ) : If for all  $i, j$ ,  $d_i \stackrel{\mathcal{A}}{\hookrightarrow} d_j$ , then  $\Pr[F_{ij}] < \epsilon$ . We can choose  $\epsilon < \epsilon'/n^2$  for a negligible  $\epsilon'$ . Notice that  $E \subseteq \bigcup_{i,j} F_{ij}$ . Therefore,  $\Pr[E] \leq \sum_{i,j} \Pr[F_{ij}] < n^2\epsilon < \epsilon'$ . Hence,  $\Pr[E^c] \geq 1 - \epsilon'$  and so the storage system is all-or-nothing.

□

All-or-nothing integrity is a very strong property. In some models, we may not be able to achieve it, and we will accept a weaker property called **symmetric recovery**. Symmetric recovery requires that all users recover their documents with equal probability:

**Definition 8** *A storage system has **symmetric recovery** with respect to a class of adversaries  $\mathcal{A}$  if, for all  $A \in \mathcal{A}$  and all users  $i$  and  $j$ ,*

$$\Pr_A[r_i = 1] = \Pr_A[r_j = 1].$$

Symmetric recovery says nothing about what happens in particular executions. For example, it is consistent with the definition for exactly one of the data items to be recovered in every execution, as long as the adversary cannot affect which data item is recovered. This is not as strong a property as all-or-nothing integrity, but it is the best that can be done in some cases.

## 5 Possibility and Impossibility Results

The possibility of achieving **all-or-nothing integrity** (abbreviated AONI) depends on the class of adversaries we consider. In Sections 5.1 through 5.3,

we consider adversaries with an **arbitrary tamperer**. We show that AONI cannot always be achieved in this case. Then in Section 5.4, we look at adversaries with a **destructive tamperer**. We give a simple interpolation scheme that achieves all-or-nothing integrity for a destructive tamperer in all three recovery models.

### 5.1 Possibility of AONI for Standard-Recovery-Algorithm Model

In the standard-recovery-algorithm model, all users use the standard recovery algorithm  $R$ ; that is  $R_i = R$  for all users  $i$ .

This model allows a very simple mechanism for all-or-nothing integrity based on Message Authentication Codes (MACs). The intuition behind this mechanism is that the encoding algorithm  $E$  simply tags the data store with a MAC using a key known to all the users, and the recovery algorithm  $R$  returns an individual user's data only if the MAC on the entire database is valid.

We begin by recalling some standard definitions (see [10]). A **MAC** consists of three algorithms ( $GEN, TAG, VER$ ):

- (1) A key generator  $GEN$  on input  $1^s$  outputs an  $s$ -bit key  $k_{MAC}$ .
- (2) A tagging algorithm  $TAG$  on input  $k_{MAC}$  and message  $m$  ( $|m| \leq s^c$ ) computes a signature  $\sigma$ .
- (3) A verification algorithm  $VER$  can be used to check if  $\sigma$  is a valid signature on  $m$ . It has the property that  $VER(k_{MAC}, m, TAG(k_{MAC}, m)) = accept$  for all  $m$ .

We require a MAC to be **existentially unforgeable** under **chosen message**

**attacks.** This means there is no polynomial time forger  $F$  that generates a new message-signature pair  $(m', \sigma')$  that is accepted by  $VER$  with probability exceeding  $O(s^{-c})$  for any  $c > 0$ , even if  $F$  is given a sample of valid message-signature pairs  $(m_i, \sigma_i)$ , where  $m_i$  is chosen by the adversary.

Our encoding scheme  $(I, E, R)$  is based on a MAC scheme  $(GEN, TAG, VER)$  as follows:

**Initialization** The initialization algorithm  $I$  computes  $k_{MAC} = GEN(1^s)$ . It then returns an encoding key  $k_E = k_{MAC}$  and recovery keys  $k_i = (i, k_{MAC})$ .

**Entanglement** The encoding algorithm  $E$  generates an  $n$ -tuple

$m = (d_1, d_2, \dots, d_n)$  and returns  $C = (m, \sigma)$  where  $\sigma = TAG(k_{MAC}, m)$ .

**Recovery** The standard recovery algorithm  $R$  takes as input a key  $k_i = (i, k_{MAC})$  and the (possibly modified) store  $\check{C} = (\check{m}, \check{\sigma})$ . It returns  $\check{m}_i$  if  $VER(k_{MAC}, \check{m}, \check{\sigma}) = \text{accept}$  and returns a default value  $\perp$  otherwise.

The following theorem states that this encoding scheme achieves all-or-nothing integrity with standard recovery algorithms:

**Theorem 9** *Let  $(GEN, TAG, VER)$  be a MAC scheme that is existentially unforgeable against chosen message attacks, and let  $(I, E, R)$  be an encoding scheme based on this MAC scheme as above. Let  $\mathcal{A}$  be the class of adversaries that does not provide non-standard recovery algorithms  $\check{R}$ . Then there exists some minimum  $s_0$  such that for any security parameter  $s \geq s_0$  and any inputs  $d_1, \dots, d_n$  with  $\sum |d_i| \leq s$ ,  $(I, E, R)$  is all-or-nothing with respect to  $\mathcal{A}$ .*

**PROOF.** Fix some  $c > 0$ . Recall that the adversary changes the combined store from  $C = (m, \sigma)$  to  $\check{C} = (\check{m}, \check{\sigma})$ . We consider two cases, depending on whether or not  $\check{m} = m$ .

In the first case,  $\check{m} = m$ . Suppose  $R(k_i, \check{C}) = d_i$  but  $R(k_j, \check{C}) \neq d_j$ . Then  $R(k_j, \check{C}) = \perp$ , which implies that  $V(k_{MAC}, m, \check{\sigma}) \neq \text{accept}$  when computed by  $R(k_j, \check{C})$  and thus that  $\check{\sigma} \neq \sigma$ . But  $R(k_i, \check{C}) = d_i$  only if  $V(k_{MAC}, m, \check{\sigma}) = \text{accept}$  when computed by  $R(k_i, \check{C})$ . It follows that  $(m, \check{\sigma})$  is a message-MAC pair not equal to  $(m, \sigma)$  that  $V$  accepts in the execution of  $R(k_i, \check{C})$ ; by the security assumption this occurs for a particular execution of  $V$  only with probability  $O(s^{-c'})$  for any fixed  $c'$ . If we choose  $c'$  and  $s_0$  so that the  $O(s^{-c'})$  term is smaller than  $\frac{1}{2^n} s^{-c}$  for  $s \geq s_0$ , then the probability that *any* of the  $n$  executions of  $V$  in the recovery stage accepts  $(m, \check{\sigma})$  in some case where  $m = \check{m}$ , is bounded by  $\frac{1}{2} s^{-c}$ .

In the second case,  $m \neq \check{m}$ . Now  $(\check{m}, \check{\sigma})$  is a message-MAC pair not equal to  $(m, \sigma)$ . If every execution of  $V$  rejects  $(\check{m}, \check{\sigma})$ , then all  $R(d_i, \check{C})$  return  $\perp$  and the execution has a recovery vector  $0^n$ . The only bad case is when at least one execution of  $V$  erroneously accepts  $(\check{m}, \check{\sigma})$ . But using the security assumption and choosing  $c', s_0$  as in the previous case, we again have that the probability that  $V$  accepts  $(\check{m}, \check{\sigma})$  in any of the  $n$  executions of  $R$  is at most  $\frac{1}{2} s^{-c}$ .

Summing the probabilities of the two bad cases gives us the desired bound:  
 $\Pr_A[\vec{r} = 0^n \vee \vec{r} = 1^n] > 1 - s^{-c}$ .  $\square$

## 5.2 Impossibility of AONI for Public and Private-Recovery-Algorithm Models

In both these models, the adversary modifies the common store and distributes a non-standard recovery algorithm  $\check{R}$  to the users (either to all users or only to a few select accomplices). Let us begin by showing that all-or-nothing integrity cannot be achieved consistently in either case:

**Theorem 10** *For any encoding scheme  $(I, E, R)$ , if  $\mathcal{A}$  is the class of adversaries providing non-standard recovery algorithms  $\check{R}$ , then  $(I, E, R)$  is not all-or-nothing with respect to  $\mathcal{A}$ .*

**PROOF.** Let the adversary initializer  $\check{I}$  be a no-op and let the tamperer  $\check{T}$  be the identity transformation. We will rely entirely on the non-standard recovery algorithm to destroy all-or-nothing integrity.

Let  $\check{R}$  flip a biased coin that comes up tails with probability  $1/n$ , and return the result of running  $R$  on its input if the coin comes up heads and  $\perp$  if the coin comes up tails. Then exactly one document is not returned with probability  $n \cdot (1/n) \cdot (1 - 1/n)^{n-1}$ , which converges to  $1/e$  in the limit. The outcome is all-or-nothing only if all instances of  $\check{R}$  flip the same way, which occurs with probability  $\Pr_{\mathcal{A}}[\vec{r} = 0^n \vee \vec{r} = 1^n] < 1 - 1/e$ .  $\square$

This proof is rather trivial, which suggests that letting the adversary substitute an error-prone recovery algorithm in place of the standard one gives the adversary far too much power. But it is not at all clear how to restrict the model to allow the adversary to provide an improved recovery algorithm without allowing for this particular attack. Allowing users to apply the original recovery algorithm  $R$  can be defeated by superencrypting the data store and burying the decryption key in the error-prone  $\check{R}$ ; defeating this attack would require analyzing  $\check{R}$  to undo the superencryption and remove the errors, a task that is likely to be difficult in practice.<sup>9</sup>

---

<sup>9</sup> Whether it is difficult from a theoretical perspective depends on how well  $\check{R}$  can be obfuscated. Though obfuscation is impossible in general [2], recovering useful information from  $\check{R}$  is likely to be difficult in practice, especially if the random choice

On the other hand, we do not know of any general mechanism to ensure that no useful information can be gleaned from  $\check{R}$ , and it is not out of the question that there is an encoding so transparent that no superencryption can disguise it for sufficiently large inputs, given that both  $\check{R}$  and the adversary's key  $\check{k}$  are public.

### 5.3 Possibility of Symmetric Recovery for Public-Recovery-Algorithm Model

As we have seen, if we place no restrictions on the tamperer, it becomes impossible to achieve all-or-nothing integrity in the public-recovery-algorithm model. We now show that we can still achieve symmetric recovery.

Because we cannot prevent mass destruction of data, we will settle for preventing targeted destruction. The basic intuition is that if the encoding process is symmetric with respect to permutations of the data, then neither the adversary nor its partner, the non-standard recovery algorithm, can distinguish between different inputs. Symmetry in the encoding algorithm is not difficult to achieve and basically requires not including any positional information in the keys or the representation of data in the common store. One example of a symmetric encoding is a trivial mechanism that tags each input  $d_i$  with a random  $k_i$  and then stores a sequence of  $(d_i, k_i)$  pairs in random order.

Symmetry in the data is a stronger requirement. We assume that users' documents  $d_i$  are independent and identically distributed (i.i.d.) random variables. If documents are not i.i.d (in particular, if they are fixed), we can use a simple  

---

to decrypt incorrectly is not a single if-then test but is the result of accumulating error distributed throughout its computation.

trick to make them appear i.i.d.: Each user  $i$  picks a small number  $r_i$  independently and uniformly at random, remembers the number, and computes  $d'_i = d_i \oplus G(r_i)$ , where  $G$  is a pseudorandom generator. The new  $d'_i$  are also uniform and independent (and thus computationally indistinguishable from i.i.d.). The users can then store documents  $d'_i$  ( $1 \leq i \leq n$ ) instead of the original documents  $d_i$ . To recover  $d_i$ , user  $i$  would retrieve  $d'_i$  from the server and compute  $d_i = d'_i \oplus G(r_i)$ .

We shall need a formal definition of symmetric encodings:

**Definition 11** *An encoding scheme  $(I, E, R)$  is **symmetric** if, for any  $s$  and  $n$ , any inputs  $d_1, d_2, \dots, d_n$ , and any permutation  $\pi$  of the indices 1 through  $n$ , the joint distribution of  $k_1, k_2, \dots, k_n$  and  $C$  in executions with user inputs  $d_1, d_2, \dots, d_n$  is computationally indistinguishable from the joint distribution of  $k_{\pi_1}, k_{\pi_2}, \dots, k_{\pi_n}$  and  $C$  in executions with user inputs  $d_{\pi_1}, d_{\pi_2}, \dots, d_{\pi_n}$ .*

Using this definition, we can show that any symmetric encoding gives symmetric recovery:

**Theorem 12** *Let  $(I, E, R)$  be a symmetric encoding scheme. Let  $\mathcal{A}$  be a class of adversaries as in Theorem 10. Fix  $s$  and  $n$ , and let  $d_1, \dots, d_n$  be random variables that are independent and identically distributed. Then  $(I, E, R)$  has symmetric recovery with respect to  $\mathcal{A}$ .*

**PROOF.** Fix  $i$  and  $j$ . From Definition 11 we have that the joint distribution of the  $k_i$  and  $C$  is symmetric with respect to permutation of the user indices; in particular, for any fixed  $d$ ,  $S$  and  $x$ ,

$$\Pr[C = S, k_i = x \mid d_i = d] = \Pr[C = S, k_j = x \mid d_j = d]. \quad (2)$$

We also have, from the assumption that the  $d_i$  are i.i.d.,

$$\Pr[d_i = d] = \Pr[d_j = d]. \quad (3)$$

Using (2) and (3), we get

$$\begin{aligned} & \Pr[\check{R}(\check{k}, k_i, \check{T}(C)) = d_i] \\ &= \sum_{x,S,d} \Pr[\check{R}(\check{k}, x, \check{T}(S)) = d] \Pr[C = S, k_i = x, d_i = d] \\ &= \sum_{x,S,d} \Pr[\check{R}(\check{k}, x, \check{T}(S)) = d] \Pr[C = S, k_i = x \mid d_i = d] \Pr[d_i = d] \\ &= \sum_{x,S,d} \Pr[\check{R}(\check{k}, x, \check{T}(S)) = d] \Pr[C = S, k_j = x \mid d_j = d] \Pr[d_j = d] \\ &= \Pr[\check{R}(\check{k}, k_j, \check{C}) = d_j]. \end{aligned}$$

This is simply another way of writing  $\Pr_A[r_i = 1] = \Pr_A[r_j = 1]$ .  $\square$

#### 5.4 Possibility of AONI for Destructive Adversaries

Given the results of the previous sections, to achieve all-or-nothing integrity we need to place some additional restrictions on the adversary.

A tampering algorithm  $\check{T}$  is **destructive** if the range of  $\check{T}$  when applied to an input domain of  $m$  distinct possible data stores has size less than  $m$ . The amount of destructiveness is measured in bits: if the range of  $\check{T}$  when applied to a domain of size  $m$  has size  $r$ , then  $\check{T}$  destroys  $\lg m - \lg r$  bits of entropy. Note that it is not necessarily the case that the outputs of  $\check{T}$  are smaller than its inputs; it is enough that there be fewer of them.

Below, we describe a particular encoding, based on polynomial interpolation, with the property that after a sufficiently destructive tampering, the prob-



ability that *any* recovery algorithm can reconstruct a particular  $d_i$  is small. While this is trivially true for an unrestrained tamperer that destroys all  $\lg m$  bits of the common store, our scheme requires only that with  $n$  documents the tamperer destroy slightly more than  $n \lg(n/\epsilon)$  bits before the probability that any of the data can be recovered drops below  $\epsilon$  (a formal statement of this result is found in Corollary 14). Since  $n$  counts only the number of users and not the size of the data, for a fixed population of users the number of bits that can be destroyed before all users lose their data is effectively a constant independent of the size of the store being tampered with.

The encoding scheme is as follows. It assumes that each data item can be encoded as an element of  $\mathbb{Z}_p$ , where  $p$  is a prime of roughly  $s$  bits.

**Initialization** The initialization algorithm  $I$  chooses  $k_1, k_2, \dots, k_n$  independently and uniformly at random *without replacement* from  $\mathbb{Z}_p$ . It sets  $k_E = (k_1, k_2, \dots, k_n)$  and then returns  $k_E, k_1, \dots, k_n$ .

**Entanglement** The encoding algorithm  $E$  computes, using Lagrange interpolation, the coefficients  $c_{n-1}, c_{n-2}, \dots, c_0$  of the unique  $(n-1)$ -degree polynomial  $f$  over  $\mathbb{Z}_p$  with the property that  $f(k_i) = d_i$  for each  $i = 1, \dots, n$ . It returns  $C = (c_{n-1}, c_{n-2}, \dots, c_0)$ .

**Recovery** The standard recovery algorithm  $R$  returns  $f(k_i)$ , where  $f$  is the polynomial whose coefficients are given by  $C$ .

Intuitively, the reason the tamperer cannot remove too much entropy without destroying all data is that it cannot identify which points  $d = f(k)$  correspond to actual user keys. When it maps two polynomials  $f_1$  and  $f_2$  to the same corrupted store  $\tilde{C}$ , the best that the non-standard recovery algorithm can do is return one of  $f_1(k_i)$  or  $f_2(k_i)$  given a particular key  $k_i$ . But if too many

polynomials are mapped to the same  $\check{C}$ , the odds that  $\check{R}$  returns the value of the correct polynomial will be small.

A complication is that a particularly clever adversary could look for polynomials whose values overlap; if  $f_1(k) = f_2(k)$ , it doesn't matter which  $f$  the recovery algorithm picks. But here we can use that fact that two degree  $(n-1)$  polynomials cannot overlap in more than  $(n-1)$  places without being equal. This limits how much packing the adversary can do.

As in Theorem 12, we assume that the user inputs  $d_1, \dots, d_n$  are chosen independently and have identical distributions. We make a further assumption that each  $d_i$  is chosen uniformly from  $\mathbb{Z}_p$ . This is necessary to ensure that the resulting polynomials span the full  $p^n$  possibilities.<sup>10</sup>

Under these conditions, sufficiently destructive tampering prevents recovery of any information with high probability. We will show an accurate but inconvenient bound on this probability in Theorem 13 and give a cruder but more useful statement of the bound in Corollary 14.

**Theorem 13** *Let  $(I, E, R)$  be defined as above. Let  $A = (\check{I}, \check{T}, \check{R})$  be an adversary where  $\check{T}$  is destructive: for a fixed input size and security parameter, there is a constant  $M$  such that for each key  $\check{k}$ ,*

$$|\{\check{T}(\check{k}, f)\}| \leq M,$$

*where  $f$  ranges over the possible store values, i.e. over all degree- $(n-1)$  polynomials over  $\mathbb{Z}_p$ . If the  $d_i$  are drawn independently and uniformly from  $\mathbb{Z}_p$ ,*

---

<sup>10</sup>The assumption that the documents are i.i.d. does not constrain the applicability of our results much, because the technique to get rid of it described in Section 5.2 can also be used here.

then the probability that at least one user  $i$  recovers  $d_i$  using  $\check{R}$  is

$$\Pr_A[\vec{r} \neq 0^n] < \frac{n^2 + nM^{1/n}}{p}, \quad (4)$$

even if all users use  $\check{R}$  as their recovery algorithm.

**PROOF.** Condition on  $\check{k}$  and the outcome of all coin-flips used by  $\check{T}$  and  $\check{R}$ . Then, there are exactly  $p^n \binom{p}{n}$  possible executions, each of equal probability, determined by the  $p^n$  choices for the  $d_i$  and the  $\binom{p}{n}$  choices for the  $k_i$ . For each  $i$ , we will show that the number of these executions in which  $\check{R}(\check{k}, k_i, \check{C}) = d_i$  is small.

For each degree- $(n-1)$  polynomial  $f$ , define  $f^*$  to be the function mapping each  $k$  in  $\mathbb{Z}_p$  to  $\check{R}(\check{k}, k, \check{T}(\check{k}, f))$ . Note that  $f^*$  is deterministic given that we are conditioning on  $\check{k}$  and all coin-flips in  $\check{T}$  and  $\check{R}$ . Define  $C_f$ , the **correct inputs** for  $f$ , to be the set of keys  $k$  for which  $f(k) = f^*(k)$ .

The adversary produces a correct output only if at least one of the  $n$  user keys appears in  $C_f$ . For a given  $f$ , the probability that *none* of the keys appear in  $C_f$  is

$$\begin{aligned} \frac{\binom{p-|C_f|}{n}}{\binom{p}{n}} &> \frac{(p-|C_f|-n)^n}{p^n} \\ &= \left(1 - \frac{|C_f|+n}{p}\right)^n \\ &> 1 - \frac{n(|C_f|+n)}{p}, \end{aligned}$$

and so the probability that *at least one* key appears in  $C_f$  is at most  $\frac{n}{p}|C_f| + \frac{n^2}{p}$ .

Averaging over all  $f$  then gives

$$\Pr [f^*(k_i) = d_i \text{ for at least one } i] < \frac{n^2}{p} + \frac{n}{p^{n+1}} \sum_f |C_f|. \quad (5)$$

We will now use the bound on the number of distinct  $f^*$  to show that  $\sum_f |C_f|$  is small.

Consider the set of all polynomials  $f_1, f_2, \dots, f_m$  that map to a single function  $f^*$ , and their corresponding sets of correct keys  $C_{f_1}, C_{f_2}, \dots, C_{f_m}$ . Because any two degree  $(n-1)$  polynomials are equal if they are equal on any  $n$  elements of  $\mathbb{Z}_p$ , each  $n$ -element subset of  $\mathbb{Z}_p$  can be contained in at most one of the  $C_{f_i}$ . On the other hand, each  $C_{f_i}$  contains exactly  $\binom{|C_{f_i}|}{n}$  subsets of size  $n$ . Since there only  $\binom{p}{n}$  subsets of size  $n$  to partition between the  $C_{f_i}$ , we have

$$\sum_i \binom{|C_{f_i}|}{n} \leq \binom{p}{n},$$

and summing over all  $M$  choices of  $f^*$  then gives

$$\sum_f \binom{|C_f|}{n} \leq M \binom{p}{n}.$$

We now wish to bound the maximum possible value of  $\sum_f |C_f|$  given this constraint.

Observe that  $\binom{|C_f|}{n} > \frac{(|C_f| - n)^n}{n!}$  when  $|C_f| \geq n$ , from which it follows that

$$\sum_{f: |C_f| \geq n} (|C_f| - n)^n < n! \sum_f \binom{|C_f|}{n} < n! M \binom{p}{n}. \quad (6)$$

Now,  $(|C_f| - n)^n$  is a convex function of  $|C_f|$ , so the left-hand side is minimized for fixed  $\sum_f |C_f|$  by setting all  $|C_f|$  equal. It follows that  $\sum_f |C_f|$  is *maximized* for fixed  $\sum_{f: |C_f| \geq n} (|C_f| - n)^n$  when all  $|C_f|$  are equal.

Setting each  $|C_f| = c$  and summing over all  $p^n$  values of  $f$ , we get

$$p^n(c - n)^n < n!M \binom{p}{n}$$

from which it follows that

$$c < \frac{1}{p} \left( n!M \binom{p}{n} \right)^{1/n} + n,$$

and thus that

$$\sum_f |C_f| \leq p^n c < p^{n-1} \left( n!M \binom{p}{n} \right)^{1/n} + np^n.$$

Plugging this bound back into (5) then gives

$$\begin{aligned} \Pr_A [\vec{r} \neq 0^n] &= \Pr [f^*(k_i) = d_i \text{ for at least one } i] \\ &< \frac{n^2}{p} + \frac{n}{p^2} \left( n!M \binom{p}{n} \right)^{1/n} \\ &< \frac{n^2}{p} + \frac{n}{p^2} (Mp^n)^{1/n} \\ &= \frac{n^2 + nM^{1/n}}{p}. \end{aligned}$$

□

We can use Theorem 13 to compute the limit on how much information the tamperer can remove before recovering any of the data becomes impossible:

**Corollary 14** *Let  $(I, E, R)$  and  $(\check{I}, \check{T}, \check{R})$  be as in Theorem 13. Let  $\epsilon > 0$  and let  $p > 2n^3/\epsilon$ . If for any fixed  $\check{k}$ , tamperer  $\check{T}$  destroys at least  $n \lg(n/\epsilon) + 1$  bits of entropy, then*

$$\Pr_A [\vec{r} = 0^n] \geq 1 - \epsilon.$$

**PROOF.** Let  $\epsilon' = \epsilon / \left(\frac{1}{2n} + 2^{-1/n}\right)$ . If  $\tilde{T}$  destroys at least  $n \lg(n/\epsilon') + 1$  bits of entropy, then we have

$$\begin{aligned} M &\leq p^n \cdot 2^{-(n \lg(n/\epsilon') + 1)} \\ &= \frac{1}{2} p^n (n/\epsilon')^{-n} \\ &= \frac{1}{2} (p\epsilon'/n)^n. \end{aligned}$$

Plugging this into (4) gives:

$$\begin{aligned} \Pr[\text{some } d_i \text{ is recovered}] &\leq \frac{n^2 + nM^{1/n}}{p} \\ &\leq \frac{n^2 + n \left(\frac{1}{2}(p\epsilon'/n)^n\right)^{1/n}}{p} \\ &= \frac{n^2}{p} + 2^{-1/n} \epsilon' \\ &< \frac{n^2}{2n^3/\epsilon'} + 2^{-1/n} \epsilon' \\ &= \epsilon' \left(\frac{1}{2n} + 2^{-1/n}\right) \\ &= \epsilon. \end{aligned}$$

We thus have:

$$\Pr_A[\vec{r} = 0^n] = 1 - \Pr[\text{some } d_i \text{ is recovered}] \geq 1 - \epsilon.$$

□

According to the corollary, all users will lose their data if the adversary's tampering destroys even a small amount of information. Consequently, the interpolation scheme we have described must be all-or-nothing.

## 6 Conclusion and Future Work

Our results are summarized below:

	Destructive Tamperer	Arbitrary Tamperer
Standard Recovery	all-or-nothing	all-or-nothing
Public Recovery	all-or-nothing	symmetric recovery
Private Recovery	all-or-nothing	no guarantees possible

They show that it is possible in principle to achieve all-or-nothing integrity with only mild restrictions on the adversary. Whether it is possible in practice is a different question. Our model abstracts away most of the details of the storage and recovery processes, which hides undesirable features of our algorithms such as the need to process all data being stored simultaneously and the need to read every bit of the data store to recover any data item. Some of these undesirable features could be removed with a more sophisticated model, such as a round-based model that treated data as arriving over time, allowing combining algorithms that would touch less of the data store for each storage or retrieval operation at the cost of making fewer documents depend on each other. The resulting system might look like a variant of Dagster or Tangler with stronger mechanisms for entanglement. But such a model might permit more dangerous attacks if the adversary is allowed to tamper with data during storage, and finding the right balance between providing useful guarantees and modeling realistic attacks will be necessary. We have made a first step towards this goal in the present work, but much still remains to be done.

## References

- [1] R. J. Anderson. The Eternity Service. In *Proceedings of PRAGOCRYPT 96*, pp. 242-252, 1996.
- [2] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology - Proceedings of CRYPTO 2001*, pp. 1-18, 2001.
- [3] M. Castro and B. Liskov. Practical Byzantine fault tolerance. In *Proceedings of the 3rd Symposium on Operating Systems Design and Implementation*, pp. 173-186, 1999.
- [4] I. Clarke, O. Sandberg, B. Wiley, and T. Hong. Freenet: A distributed information storage and retrieval system. In *Designing Privacy Enhancing Technologies: International Workshop on Design Issues in Anonymity and Unobservability*, volume 2009 of *Lecture Notes in Computer Science*, pp. 46-66, 2000.
- [5] K. Fu, F. Kaashoek, and D. Mazieres. Fast and secure distributed read-only file system. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, pp. 181-196, 2000.
- [6] G. A. Gibson, D. F. Nagle, K. Amiri, J. Butler, F. W. Chang, H. Gobiuff, C. Hardin, E. Riedel, D. Rochberg, and J. Zelenka. A cost-effective, high-bandwidth storage architecture. In *Proceedings of the 8th International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 92-103, 1998.
- [7] E. Goh, H. Shacham, N. Mdadugu, and D. Boneh. SiRiUS: Securing remote untrusted storage. In *Proceedings of the Internet Society (ISOC) Network and Distributed Systems Security (NDSS) Symposium*, pp. 131-145, 2003.



- [8] A. Goldberg and P. Yianilos. Towards an archival intermemory. In *Proceedings of the IEEE International Forum on Research and Technology, Advances in Digital Libraries (ADL '98)*, pp. 147-156. IEEE Computer Society, 1998.
- [9] S. Goldwasser and M. Bellare. Lecture Notes on Cryptography. Summer Course “*Cryptography and Computer Security*” at MIT, 1996-1999.
- [10] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen message attack. *SIAM Journal on Computing*, 17(1988), pp. 281-308.
- [11] U. Maheshwari and R. Vingralek. How to build a trusted database system on untrusted storage. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, pp. 135-150, 2000.
- [12] D. Mazieres and D. Shasha. Don't trust your file server. In *Proceedings of the 8th IEEE Workshop on Hot Topics in Operating Systems*, pp. 99-104, 2001.
- [13] D. Mazieres and D. Shasha. Building secure file systems out of Byzantine storage. In *Proceedings of the 21st Annual ACM Symposium on Principles of Distributed Computing*, pp. 108-117, 2002.
- [14] D. Mazieres and M. Waldman. Tangler: A censorship-resistant publishing system based on document entanglements. In *Proceedings of the 8th ACM Conference on Computer and Communications Security*, pp. 126-135, 2001.
- [15] R. Merkle. Protocols for public key cryptosystems. In *IEEE Symposium on Security and Privacy*, pp. 122-134, 1980.
- [16] Mojo Nation. Technology overview. Online at [http://www.mojonation.net/docs/technical\\_overview.shtml](http://www.mojonation.net/docs/technical_overview.shtml), 2000.
- [17] R. Rivest. All-or-Nothing Encryption and the Package Transform. In *Fast Software Encryption*, volume 1267 of *Lecture Notes in Computer Science*, pp.210-218, 1997.

- [18] A. Shamir. How to share a secret. *Communications of the ACM*, 22(1979), pp. 612-613.
- [19] J. Strunk, G. Goodson, M. Scheinholtz, C. Soules, and G. Ganger. Self-securing storage: Protecting data in compromised systems. In *Proceedings of the 4th Symposium on Operating Systems Design and Implementation*, pp. 165-180, 2000.
- [20] A. Stubblefield and D. S. Wallach. Dagster: Censorship-resistant publishing without replication. Technical Report TR01-380, Rice University, 2001.
- [21] M. Waldman, A. D. Rubin, and L. F. Cranor. Publius: A robust, tamper-evident, censorship-resistant, web publishing system. In *Proceedings of the 9th USENIX Security Symposium*, pp. 59-72, August 2000.