

# Locally Decodable and Updatable Non-Malleable Codes and Their Applications

Dana Dachman-Soled<sup>1</sup>, Feng-Hao Liu<sup>2</sup>, Elaine Shi<sup>2</sup>, and Hong-Sheng Zhou<sup>3</sup>

<sup>1</sup> University of Maryland, [danadach@ece.umd.edu](mailto:danadach@ece.umd.edu)

<sup>2</sup> University of Maryland, [{fenghao,elaine}@cs.umd.edu](mailto:{fenghao,elaine}@cs.umd.edu)

<sup>3</sup> Virginia Commonwealth University, [hszhou@vcu.edu](mailto:hszhou@vcu.edu)

**Abstract.** Non-malleable codes, introduced as a relaxation of error-correcting codes by Dziembowski, Pietrzak and Wichs (ICS '10), provide the security guarantee that the message contained in a tampered codeword is either the same as the original message or is set to an unrelated value. Various applications of non-malleable codes have been discovered, and one of the most significant applications among these is the connection with tamper-resilient cryptography. There is a large body of work considering security against various classes of tampering functions, as well as non-malleable codes with enhanced features such as *leakage resilience*.

In this work, we propose combining the concepts of *non-malleability*, *leakage resilience*, and *locality* in a coding scheme. The contribution of this work is three-fold:

1. As a conceptual contribution, we define a new notion of *locally decodable and updatable non-malleable code* that combines the above properties.
2. We present two simple and efficient constructions achieving our new notion with different levels of security.
3. We present an important application of our new tool – securing RAM computation against memory tampering and leakage attacks. This is analogous to the usage of traditional non-malleable codes to secure implementations in the *circuit* model against memory tampering and leakage attacks.

## 1 Introduction

The notion of non-malleable codes was defined by Dziembowski, Pietrzak and Wichs [22] as a relaxation of error-correcting codes. Informally, a coding scheme is **non-malleable** against a tampering function if by tampering with the codeword, the function can either keep the underlying message unchanged or change it to an unrelated message. Designing non-malleable codes is not only an interesting mathematical task, but also has important implications in cryptography; for example, Coretti et al. [11] showed an efficient construction of a multi-bit CCA secure encryption scheme from a single-bit one via non-malleable codes. Agrawal et al. [3] showed how to use non-malleable codes to build non-malleable commitments. Most notably, the notion has a deep connection with security

against so-called *physical attacks*; indeed, using non-malleable codes to achieve security against physical attacks was the original motivation of the work [22]. Due to this important application, research on non-malleable codes has become an important agenda, and drawn much attention in both coding theory and cryptography.

Briefly speaking, physical attacks target implementations of cryptographic algorithms beyond their input/output behavior. For example, researchers have demonstrated that leaking/tampering with sensitive secrets such as cryptographic keys, through timing channel, differential power analysis, and various other attacks, can be devastating [2, 4, 5, 32, 39, 40, 45], and therefore the community has focused on developing new mechanisms to defend against such strong adversaries [12, 13, 15–19, 21, 26–28, 30, 31, 34–36, 38, 43, 44, 46]. Dziembowski, Pietrzak and Wichs [22] showed a simple and elegant mechanism to secure implementations against memory tampering attacks by using non-malleable codes – instead of storing the secret (in the clear) on a device, one instead stores an encoding of the secret. The security of the non-malleable code guarantees that the adversary cannot learn more than what can be learnt via black box access to the device, even though the adversary may tamper with memory.

In a subsequent work, Liu and Lysyanskaya [42] extended the notion to capture **leakage resilience** as well – in addition to non-malleability, the adversary cannot learn anything about the underlying message even while obtaining partial leakage of the codeword. By using the approach outlined above, one can achieve security guarantees against both tampering and leakage attacks. In recent years, researchers have been studying various flavors of non-malleable codes; for example some work has focused on constructions against different classes of tampering functions, some has focused on different additional features, (e.g. continual attacks, rates of the scheme, etc), and some focused on other applications [1, 3, 7–9, 20, 23, 25].

In this paper, we focus on another important feature inspired from the field of coding theory – **locality**. More concretely, we consider a coding scheme that is locally decodable and updatable. As introduced by Katz and Trevisan [37], local decodability means that in order to retrieve a portion of the underlying message, one does not need to read through the whole codeword. Instead, one can just read a few locations at the codeword. Similarly, local updatability means that in order to update some part of the underlying messages, one only needs to update some parts of the codeword. Locally decodable codes have many important applications in private information retrieval [10] and secure multi-party computation [33], and have deep connections with complexity theory; see [49]. Achieving local decodability and updatability simultaneously makes the task more challenging. Recently, Chandran et al. [6] constructed a locally decodable and updatable code in the setting of *error-correcting* codes. They also show an application to dynamic proofs of retrievability. Motivated by the above results, we further ask the following intriguing question:

*Can we build a coding scheme enjoying all three properties, i.e., non-malleability, leakage resilience, and locality? If so, what are its implications in cryptography?*

*Our Results.* In light of the above questions, our contribution is three-fold:

- **(Notions).** We propose new notions that combine the concepts of non-malleability, leakage resilience, and locality in codes. First, we formalize a new notion of *locally decodable and updatable non-malleable codes* (against one-time attacks). Then, we extend this new notion to capture leakage resilience under continual attacks.
- **(Constructions).** We present two simple constructions achieving our new notions. The first construction is highly efficient—in order to decode (update) one block of the encoded messages, only *two* blocks of the codeword must be read (written)—but is only secure against *one-time* attacks. The second construction achieves security against *continual* attacks, while requiring  $\log(n)$  number of reads (writes) to perform one decode (update) operation, where  $n$  is the number of blocks of the underlying message.
- **(Application).** We present an important application of our new notion – achieving tamper and leakage resilience in the random access machine (RAM) model. We first define a new model that captures tampering and leakage attacks in the RAM model, and then give a generic compiler that uses our new notion as a main ingredient. The compiled machine will be resilient to leakage and tampering on the random access memory. This is analogous to the usage of traditional non-malleable codes to secure implementations in the *circuit* model.

## 1.1 Techniques

In this section, we present a technical overview of our results.

*Locally Decodable Non-malleable Codes.* Our first goal is to consider a combination of concepts of non-malleability and local decodability. Recall that a coding scheme is non-malleable with respect to a tampering function  $f$  if the decoding of the tampered codeword remains the same or becomes some unrelated message. To capture this idea, the definition in the work [22] requires that there exists a simulator (with respect to such  $f$ ) who outputs **same\*** if the decoding of the tampered codeword remains the same as the original one, or he outputs a decoded message, which is unrelated to the original one. In the setting of local decodability, we consider encodings of blocks of messages  $M = (m_1, m_2, \dots, m_n)$ , and we are able to retrieve  $m_i$  by running  $\text{DEC}^{\text{ENC}(M)}(i)$ , where the decoding algorithm gets oracle access to the codeword.

The combination faces a subtlety that we cannot directly use the previous definition: suppose a tampering function  $f$  only modifies one block of the codeword, then it is likely that  $\text{DEC}$  remains unchanged for most places. (Recall a  $\text{DEC}$  will only read a few blocks of the codeword, so it may not detect the modification.)

In this case, the (overall) decoding of  $f(C)$  (i.e.  $(\text{DEC}^{f(C)}(1), \dots, \text{DEC}^{f(C)}(n))$ ) can be highly related to the original message, which intuitively means it is highly malleable.

To handle this issue, we consider a more fine-grained experiment. Informally, we require that for any tampering function  $f$  (within some class), there exists a simulator that computes a vector of decoded messages  $\mathbf{m}^*$ , a set of indices  $\mathcal{I} \subseteq [n]$ . Here  $\mathcal{I}$  denotes the coordinates of the underlying messages that have been tampered with. If  $\mathcal{I} = [n]$ , then the simulator thinks that the decoded messages are  $\mathbf{m}^*$ , which should be unrelated to the original messages. On the other hand, if  $\mathcal{I} \subsetneq [n]$ , the simulator thinks that all the messages not in  $\mathcal{I}$  remain unchanged, while those in  $\mathcal{I}$  become  $\perp$ . This intuitively means the tampering function can do only one of the following cases:

1. It destroys a block (or blocks) of the underlying messages while keeping the other blocks unchanged, or
2. If it modifies a block of the underlying messages to some unrelated string, then it must have modified all blocks of the underlying messages to encodings of unrelated messages.

Our construction of locally decodable non-malleable code is simple – we use the idea similar to the *key encapsulation mechanism/data encapsulation mechanism (KEM/DEM) framework*. Let NMC be a regular non-malleable code, and  $\mathcal{E}$  be a secure (symmetric-key) authenticated encryption. Then to encode blocks of messages  $M = (m_1, \dots, m_n)$ , we first sample a secret key  $\text{sk}$  of  $\mathcal{E}$ , and then output  $(\text{NMC.ENC}(\text{sk}), \mathcal{E}.\text{Encrypt}_{\text{sk}}(m_1, 1), \dots, \mathcal{E}.\text{Encrypt}_{\text{sk}}(m_n, n))$ . The intuition is clear: if the tampering function does not change the first block, then by security of the authenticated encryption, any modification of the rest will become  $\perp$ . (Note that here we include a tag of positions to prevent permutation attacks). On the other hand, if the tampering function modified the first block, it must be decoded to an unrelated secret key  $\text{sk}'$ . Then by semantic security of the encryption scheme, the decoded values of the rest must be unrelated. The code can be updated locally: in order to update  $m_i$  to some  $m'_i$ , one just need to retrieve the 1<sup>st</sup> and  $(i + 1)$ <sup>st</sup> blocks. Then he just computes a *fresh* encoding of  $\text{NMC.ENC}(\text{sk})$  and the ciphertext  $\mathcal{E}.\text{Encrypt}_{\text{sk}}(m'_i)$ , and writes back to the same positions.

*Extensions to Leakage Resilience against Continual Attacks.* We further consider a notion that captures leakage attacks in the continual model. First we observe that suppose the underlying non-malleable code is also leakage resilient [42], the above construction also achieves one-time leakage resilience. Using the same argument of Liu and Lysyanskaya [42], if we can refresh the whole encoding, we can show that the construction is secure against continual attacks. However, in our setting, refreshing the whole codeword is not counted as a solution since this is in the opposite of the spirit of our main theme – *locality*. The main challenge is how to refresh (update) the codeword locally while maintaining tamper and leakage resilience.

To capture the local refreshing and continual attacks, we consider a new model where there is an updater  $\mathcal{U}$  who reads the whole underlying messages and decides how to update the codeword (using the local update algorithm). The updater is going to interact with the codeword in a continual manner, while the adversary can launch tampering and leakage attacks between two updates. To define security we require that the adversary cannot learn anything of the underlying messages via tampering and leakage attacks from the interaction.

We note that if there is no update procedure at all, then no coding scheme can be secure against continual leakage attacks if the adversary can learn the whole codeword bit-by-bit. In our model, the updater and the adversary take turns interacting with the codeword – the adversary tampers with and/or gets leakage of the codeword, and then the updater *locally* updates the codeword, and the process repeats. See [Section 2](#) for the formal model.

Then we consider how to achieve this notion. First we observe that the construction above is not secure under continual attacks: suppose by leakage the adversary can get a full ciphertext  $\mathcal{E}.\text{Encrypt}_{\text{sk}}(m_i, i)$  at some point, and then the updater updates the underlying message to  $m'_i$ . In the next round, the adversary can apply a *rewind attack* that modifies the codeword back with the old ciphertext. Under such attack, the underlying messages have been modified to some related messages. Thus the construction is not secure.

One way to handle this type of rewind attacks is to tie all the blocks of ciphertexts together with a “time stamp” that prevents the adversary from replacing the codeword with old ciphertexts obtained from leakage. A straightforward way is to hash all the blocks of encryptions using a collision resistant hash function and also encode this value into the non-malleable code, i.e.,  $C = (\text{NMC}.\text{ENC}(\text{sk}, v), \mathcal{E}.\text{Encrypt}(1, m_1), \dots, \mathcal{E}.\text{Encrypt}(n, m_n))$ , where  $v = h(\mathcal{E}.\text{Encrypt}(1, m_1), \dots, \mathcal{E}.\text{Encrypt}(n, m_n))$ . Intuitively, suppose the adversary replaces a block  $\mathcal{E}.\text{Encrypt}(i, m_i)$  by some old ciphertexts, then it would be caught by the hash value  $v$  unless he tampered with the non-malleable code as well. But if he tampers with the non-malleable code, the decoding will be unrelated to  $\text{sk}$ , and thus the rest of ciphertexts become “un-decryptable”. This approach prevents the rewind attacks, yet it does not preserve the local properties, i.e. to decode a block, one needs to check the consistency of the hash value  $v$ , which needs to read all the blocks of encryptions. To prevent the rewind attacks while maintaining local decodability/updatability, we use the Merkle tree technique, which allows local checks of consistency. The final encoding outputs the following:

$$(\text{NMC}.\text{ENC}(\text{sk}, v), \mathcal{E}.\text{Encrypt}(1, m_1), \dots, \mathcal{E}.\text{Encrypt}(n, m_n), T)$$

where  $T$  is the Merkle tree of  $(\mathcal{E}.\text{Encrypt}(1, m_1), \dots, \mathcal{E}.\text{Encrypt}(n, m_n))$ , and  $v$  is its root (it can also be viewed as a hash value). To decode a position  $i$ , the algorithm reads the 1<sup>st</sup>, and the  $(i+1)$ <sup>st</sup> blocks together with a path in the tree. If the path is inconsistent with the root, then output  $\perp$ . To update, one only needs to re-encode the first block with a new root, and update the  $(i+1)$ <sup>st</sup> block and the tree. We note that Merkle tree allows local updates: if there is only one single change at a leaf, then one can compute the new root given only a path

passing through the leaf and the root. So the update of the codeword can be done locally by reading the 1<sup>st</sup>, the  $(i + 1)$ <sup>st</sup> blocks and the path. We provide a detailed description and analysis in [Section 3.2](#).

*Concrete Instantiations.* In our construction above, we rely on an underlying non-malleable code NMC against some class of tampering functions  $\mathcal{F}$  and leakage resilient against some class of leakage functions  $\mathcal{G}$ . The resulting encoding scheme is a locally decodable and updatable coding scheme which is continual non-malleable against some class  $\overline{\mathcal{F}}$  of tampering functions and leakage resilient against some class  $\overline{\mathcal{G}}$  of leakage functions, where the class  $\overline{\mathcal{F}}$  is determined by  $\mathcal{F}$  and the class  $\overline{\mathcal{G}}$  is determined by  $\mathcal{G}$ . In order to understand the relationship between these classes, it is helpful to recall the structure of the output of the final encoding scheme. The final encoding scheme will output  $2n + 1$  blocks  $x_1, \dots, x_{2n+1}$  such that the first block  $x_1$  is encoded using the underlying non-malleable code NMC. As a first attempt, we can define  $\overline{\mathcal{F}}$  to consist of tampering functions  $f(x_1, \dots, x_{2n+1}) = (f_1(x_1), f_2(x_2, \dots, x_{2n+1}))$ , where  $f_1 \in \mathcal{F}$  and  $f_2$  is any polynomial-sized circuit. However, it turns out that we are resilient against an even larger class of tampering functions! This is because the tampering function  $f_1$  can actually depend on all the values  $x_2, \dots, x_{2n+1}$  of blocks  $2, \dots, (2n + 1)$ . Similarly, for the class of leakage functions, as a first attempt, we can define  $\overline{\mathcal{G}}$  to consist of leakage functions  $g(x_1, \dots, x_{2n+1}) = (g_1(x_1), g_2(x_2, \dots, x_{2n+1}))$ , where  $g_1 \in \mathcal{G}$  and  $g_2$  is any polynomial-sized circuit. However, again we can achieve even more because the tampering function  $g_1$  can actually depend on all the values  $x_2, \dots, x_{2n+1}$ . For a formal definition of the classes of tampering and leakage functions that we handle, see [Theorem 3](#).

Finally, we give a concrete example of what the resulting classes look like using the NMC construction of Liu and Lysyanskaya [[42](#)] as the building block. Recall that their construction achieves both tamper and leakage resilience for split-state functions. Thus, the overall tampering function  $f$  restricted in the first block (i.e.  $f_1$ ) can be any (poly-sized) split-state function. On the other hand  $f$  restricted in the rest (i.e.  $f_2$ ) can be any poly-sized function. The overall leakage function  $g$  restricted in the first block (i.e.  $g_1$ ) can be a (poly-sized) length-bounded split-state function;  $g$  on the other hand, can leak all the other parts. See [Section 3.3](#) for more details.

*Application to Tamper and Leakage Resilient RAM Model.* Whereas regular non-malleable codes yield secure implementations against memory tampering in the circuit model, our new tool yields secure implementations against memory tampering (and leakage) in the RAM model.

In our RAM model, the data and program to be executed are stored in the random access memory. Through a CPU with a small number of (non-persistent) registers<sup>4</sup>, execution proceeds in clock cycles: In each clock-cycle memory addresses are read and stored in registers, a computation is performed, and the

---

<sup>4</sup> These non-persistent registers are viewed as part of the circuitry that stores some transient states while the CPU is computing at each cycle. The number of these registers is small, and the CPU needs to erase the data in order to reuse them,

contents of the registers are written back to memory. In our attack model, we assume that the CPU circuitry (including the non-persistent registers) is secure – the computation itself is not subject to physical attacks. On the other hand, the random access memory, and the memory addresses are prone to leakage and tampering attacks. We remark that if the CPU has *secure* persistent registers that store a secret key, then the problem becomes straightforward: Security can be achieved using encryption and authentication together with oblivious RAM [29]. We emphasize that in our model, persistent states of the CPU are stored in the memory, which are prone to leakage and tampering attacks. As our model allows the adversary to learn the access patterns the CPU made to the memory, together with the leakage and tampering power on the memory, the adversary can somewhat learn the messages transmitted over the *bus* or tamper with them (depending on the attack classes allowed on the memory). For simplicity of presentation, we do not define attacks on the bus, but just remark that these attacks can be implicitly captured by learning the access patterns and attacking the memory<sup>5</sup>.

In our formal modeling, we consider a next instruction function  $\Pi$ , a database  $D$  (stored in the random access memory) and an internal state (using the non-persistent registers). The CPU will interact (i.e., read/write) with the memory based on  $\Pi$ , while the adversary can launch tamper and leakage attacks during the interaction.

Our compiler is very simple, given the ORAM technique and our new codes as building blocks. Informally speaking, given any next instruction function  $\Pi$  and database  $D$ , we first use ORAM technique to transform them into a next instruction function  $\tilde{\Pi}$  and a database  $\tilde{D}$ . Next, we use our local non-malleable code (ENC, DEC, UPDATE) to encode  $\tilde{D}$  into  $\hat{D}$ ; the compiled next instruction function  $\hat{\Pi}$  does the following: run  $\tilde{\Pi}$  to compute the next “virtual” read/write instruction, and then run the local decoding or update algorithms to perform the physical memory access.

Intuitively, the inner ORAM protects leakage of the address patterns, and the outer local non-malleable codes prevent an attacker from modifying the contents of memory to some *different* but *related* value. Since at each cycle the CPU can only read and write at a small number of locations of the memory, using regular non-malleable codes does not work. Our new notion of locally decodable and updatable non-malleable codes exactly solves these issues!

## 1.2 Related Work

Different flavors of non-malleable codes were studied [1, 3, 7–9, 20, 22, 23, 25, 42]. We can use these constructions to secure implementations against memory at-

---

so they cannot be used to store a secret key that is needed for a long term of computation.

<sup>5</sup> There are some technical subtleties to simulate all leakage/tampering attacks on the values passing the bus using memory attacks (and addresses). We defer the rigorous treatment to future work.

tacks in the circuit model, and also as our building block for the locally decodable/updatable non-malleable codes. See also Section 3.3 for further exposition.

Securing circuits or CPUs against physical attacks is an important task, but out of the scope of this paper. Some partial results can be found in previous work [12, 13, 16–19, 21, 26, 27, 30, 31, 34–36, 38, 41, 43, 44, 46–48].

In an independent and concurrent work, Faust et al. [24] also considered securing RAM computation against tampering and leakage attacks. We note that both their model and techniques differ considerably from ours. In the following, we highlight some of these differences. The main focus of [24] is constructing RAM compilers for keyed functions, denoted  $\mathcal{G}_K$ , to allow secure RAM emulation of these functions in the presence of leakage and tampering. In contrast, our work focuses on constructing compilers that transform any dynamic RAM machine into a RAM machine secure against leakage and tampering. Due to this different perspective, our compiler explicitly utilizes an underlying ORAM compiler, while they assume that the memory access pattern of input function  $\mathcal{G}$  is independent of the secret state  $K$  (e.g., think of  $\mathcal{G}$  as the circuit representation of the function). In addition to the split-state tampering and leakage attacks considered by both papers, [24] do not assume that memory can be overwritten or erased, but require the storage of a tamper-proof program counter. With regard to techniques, they use a stronger version of non-malleable codes in the split-state setting (called continual non-malleable codes [23]) for their construction. Finally, in their construction, each memory location is encoded using an expensive non-malleable encoding scheme, while in our construction, non-malleable codes are used only for a small portion of the memory, while highly efficient symmetric key authenticated encryption is used for the remainder.

## 2 Locally Decodable and Updatable Non-Malleable Codes

In this section, we first review the concepts of non-malleable (leakage resilient) codes. Then we present our new notion that combines non-malleability, leakage resilience, and locality.

### 2.1 Preliminary

**Definition 1 (Coding Scheme).** *Let  $\Sigma, \hat{\Sigma}$  be sets of strings, and  $\kappa, \hat{\kappa} \in \mathbb{N}$  be some parameters. A coding scheme consists of two algorithms (ENC, DEC) with the following syntax:*

- *The encoding algorithm (perhaps randomized) takes input a block of message in  $\Sigma$  and outputs a codeword in  $\hat{\Sigma}$ .*
- *The decoding algorithm takes input a codeword in  $\hat{\Sigma}$  and outputs a block of message in  $\Sigma$ .*

*We require that for any message  $m \in \Sigma$ ,  $\Pr[\text{DEC}(\text{ENC}(m)) = m] = 1$ , where the probability is taken over the choice of the encoding algorithm. In binary settings, we often set  $\Sigma = \{0, 1\}^\kappa$  and  $\hat{\Sigma} = \{0, 1\}^{\hat{\kappa}}$ .*



**Definition 2 (Non-malleability [22]).** Let  $k$  be the security parameter,  $\mathcal{F}$  be some family of functions. For each function  $f \in \mathcal{F}$ , and  $m \in \Sigma$ , define the tampering experiment:

$$\mathbf{Tamper}_m^f \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow \text{ENC}(m), \tilde{c} := f(c), \tilde{m} := \text{DEC}(\tilde{c}). \\ \text{Output} : \tilde{m}. \end{array} \right\},$$

where the randomness of the experiment comes from the encoding algorithm. We say a coding scheme  $(\text{ENC}, \text{DEC})$  is non-malleable with respect to  $\mathcal{F}$  if for each  $f \in \mathcal{F}$ , there exists a PPT simulator  $\mathcal{S}$  such that for any message  $m \in \Sigma$ , we have

$$\mathbf{Tamper}_m^f \approx \mathbf{Ideal}_{\mathcal{S}, m} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \tilde{m} \cup \{\text{same}^*\} \leftarrow \mathcal{S}^{f(\cdot)}. \\ \text{Output} : m \text{ if same}^*; \text{ otherwise } \tilde{m}. \end{array} \right\}$$

Here the indistinguishability can be either computational or statistical.

We can extend the notion of non-malleability to leakage resilience (simultaneously) as the work of Liu and Lysyanskaya [42].

**Definition 3 (Non-malleability and Leakage Resilience [42]).** Let  $k$  be the security parameter,  $\mathcal{F}, \mathcal{G}$  be some families of functions. For each function  $f \in \mathcal{F}$ ,  $g \in \mathcal{G}$ , and  $m \in \Sigma$ , define the tamper-leak experiment:

$$\mathbf{TamperLeak}_m^{f,g} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} c \leftarrow \text{ENC}(m), \tilde{c} := f(c), \tilde{m} := \text{DEC}(\tilde{c}). \\ \text{Output} : (\tilde{m}, g(c)). \end{array} \right\},$$

where the randomness of the experiment comes from the encoding algorithm. We say a coding scheme  $(\text{ENC}, \text{DEC})$  is non-malleable and leakage resilience with respect to  $\mathcal{F}$  and  $\mathcal{G}$  if for any  $f \in \mathcal{F}$ ,  $g \in \mathcal{G}$ , there exists a PPT simulator  $\mathcal{S}$  such that for any message  $m \in \Sigma$ , we have

$$\mathbf{TamperLeak}_m^{f,g} \approx \mathbf{Ideal}_{\mathcal{S}, m} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} (\tilde{m} \cup \{\text{same}^*\}, \ell) \leftarrow \mathcal{S}^{f(\cdot), g(\cdot)}. \\ \text{Output} : (m, \ell) \text{ if same}^*; \text{ otherwise } (\tilde{m}, \ell). \end{array} \right\}$$

Here the indistinguishability can be either computational or statistical.

## 2.2 New Definitions – Codes with Local Properties

In this section, we consider coding schemes with extra *local* properties – decodability and updatability. Intuitively, this gives a way to encode blocks of messages, such that in order to decode (retrieve) a single block of the messages, one only needs to read a small number of blocks of the codeword; similarly, in order to update a single block of the messages, one only needs to update a few blocks of the codeword.

**Definition 4 (Locally Decodable and Updatable Code).** Let  $\Sigma, \hat{\Sigma}$  be sets of strings, and  $n, \hat{n}, p, q$  be some parameters. An  $(n, \hat{n}, p, q)$  locally decodable and updatable coding scheme consists of three algorithms (ENC, DEC, UPDATE) with the following syntax:

- The encoding algorithm ENC (perhaps randomized) takes input an  $n$ -block (in  $\Sigma$ ) message and outputs an  $\hat{n}$ -block (in  $\hat{\Sigma}$ ) codeword.
- The (local) decoding algorithm DEC takes input an index in  $[n]$ , reads at most  $p$  blocks of the codeword, and outputs a block of message in  $\Sigma$ . The overall decoding algorithm simply outputs (DEC(1), DEC(2),  $\dots$ , DEC( $n$ )).
- The (local) updating algorithm UPDATE (perhaps randomized) takes inputs an index in  $[n]$  and a string in  $\Sigma \cup \{\epsilon\}$ , and reads/writes at most  $q$  blocks of the codeword. Here the string  $\epsilon$  denotes the procedure of refreshing without changing anything.

Let  $C \in \hat{\Sigma}^{\hat{n}}$  be a codeword. For convenience, we denote  $\text{DEC}^C, \text{UPDATE}^C$  as the processes of reading/writing individual block of the codeword, i.e. the codeword oracle returns or modifies individual block upon a query. Here we view  $C$  as a random access memory where the algorithms can read/write to the memory  $C$  at individual different locations.

*Remark 1.* Throughout this paper, we only consider non-adaptive decoding and updating, which means the algorithms DEC and UPDATE compute all their queries at the same time before seeing the answers, and the computation only depends on the input  $i$  (the location). In contrast, an adaptive algorithm can compute a query based on the answer from previous queries. After learning the answer to such query, then it can make another query. We leave it as an interesting open question to construct more efficient schemes using adaptive queries.

Then we define the requirements of the coding scheme.

**Definition 5 (Correctness).** An  $(n, \hat{n}, p, q)$  locally decodable and updatable coding scheme (with respect to  $\Sigma, \hat{\Sigma}$ ) satisfies the following properties. For any message  $M = (m_1, m_2, \dots, m_n) \in \Sigma^n$ , let  $C = (c_1, c_2, \dots, c_{\hat{n}}) \leftarrow \text{ENC}(M)$  be a codeword output by the encoding algorithm. Then we have:

- for any index  $i \in [n]$ ,  $\Pr[\text{DEC}^C(i) = m_i] = 1$ , where the probability is over the randomness of the encoding algorithm.
- for any update procedure with input  $(j, m') \in [n] \times \Sigma \cup \{\epsilon\}$ , let  $C'$  be the resulting codeword by running  $\text{UPDATE}^C(j, m')$ . Then we have  $\Pr[\text{DEC}^{C'}(j) = m'] = 1$ , where the probability is over the encoding and update procedures. Moreover, the decodings of the other positions remain unchanged.

*Remark 2.* The correctness definition can be directly extended to handle any sequence of updates.

Next, we define several flavors of security about non-malleability and leakage resilience.

*One-time Non-malleability.* First we consider one-time non-malleability of locally decodable codes, i.e., the adversary only tampers with the codeword once. This extends the idea of the non-malleable codes (as in [Definition 2](#)). As discussed in the introduction, we present the following definition to capture the idea that the tampering function can only do either of the following cases:

- It destroys a block (or blocks) of the underlying messages while keeping the other blocks unchanged, or
- If it modifies a block of the underlying messages to some unrelated string, then it must have modified all blocks of the underlying messages to encodings of unrelated messages.

**Definition 6 (Non-malleability of Locally Decodable Codes).** An  $(n, \hat{n}, p, q)$ -locally decodable coding scheme with respect to  $\Sigma, \hat{\Sigma}$  is non-malleable against the tampering function class  $\mathcal{F}$  if for all  $f \in \mathcal{F}$ , there exists some simulator  $\mathcal{S}$  such that for any  $M = (m_1, \dots, m_n) \in \Sigma^n$ , the experiment  $\mathbf{Tamper}_M^f$  is (computationally) indistinguishable to the following ideal experiment  $\mathbf{Ideal}_{\mathcal{S}, M}$ :

- $(\mathcal{I}, \mathbf{m}^*) \leftarrow \mathcal{S}(1^k)$ , where  $\mathcal{I} \subseteq [n]$ ,  $\mathbf{m}^* \in \Sigma^n$ . (Intuitively  $\mathcal{I}$  means the coordinates of the underlying message that have been tampered with).
- If  $\mathcal{I} = [n]$ , define  $\mathbf{m} = \mathbf{m}^*$ ; otherwise set  $\mathbf{m}|_{\mathcal{I}} = \perp, \mathbf{m}|_{\bar{\mathcal{I}}} = M|_{\bar{\mathcal{I}}}$ , where  $\mathbf{x}|_{\mathcal{I}}$  denotes the coordinates  $\mathbf{x}[v]$  where  $v \in \mathcal{I}$ , and the bar denotes the complement of a set.
- The experiment outputs  $\mathbf{m}$ .

*Remark 3.* Here we make two remarks about the definition:

1. In the one-time security definition, we do not consider the update procedure. In the next paragraph when we define continual attacks, we will handle the update procedure explicitly.
2. One-time leakage resilience of locally decodable codes can be defined in the same way as [Definition 3](#).

*Security against Continual Attacks.* In the following, we extend the security to handle continual attacks. Here we consider a third party called *updater*, who can read the underlying messages and decide how to update the codeword. Our model allows the adversary to learn the location that the updater updated the messages, so we also allow the simulator to learn this information. This is without loss of generality if the leakage class  $\mathcal{G}$  allows it, i.e. the adversary can query some  $g \in \mathcal{G}$  to figure out what location was modified. On the other hand, the updater does not tell the adversary what content was encoded of the updated messages, so the simulator needs to simulate the view without such information. We can think of the updater as an honest user interacting with the codeword (read/write). The security intuitively means that even if the adversary can launch tampering and leakage attacks when the updater is interacting with the codeword, the adversary cannot learn anything about the underlying encoded messages (or the updated messages during the interaction).

Our continual experiment consists of rounds: in each round the adversary can tamper with the codeword and get partial information. At the end of each round,

the updater will run UPDATE, and the codeword will be somewhat updated and refreshed. We note that if there is no refreshing procedure, then no coding scheme can be secure against continual leakage attack even for one-bit leakage at a time<sup>6</sup>, so this property is necessary. Our concept of “continuity” is different from that of Faust et al. [23], who considered continual attacks on the same original codeword (the tampering functions can be chosen adaptively). Our model does not allow this type of “resetting attacks.” Once a codeword has been modified to  $f(C)$ , the next tampering function will be applied on  $f(C)$ .

We remark that the one-time security can be easily extended to the continual case (using a standard hybrid argument) if the update procedure re-encodes the *whole* underlying messages (c.f. see the results in the work [42]). However, in the setting above, we emphasize on the *local property*, so this approach does not work. How to do a local update while maintaining tamper and leakage resilience makes the continual case challenging!

**Definition 7 (Continual Tampering and Leakage Experiment).** *Let  $k$  be the security parameter,  $\mathcal{F}, \mathcal{G}$  be some families of functions. Let  $(\text{ENC}, \text{DEC}, \text{UPDATE})$  be an  $(n, \hat{n}, p, q)$ -locally decodable and updatable coding scheme with respect to  $\Sigma, \hat{\Sigma}$ . Let  $\mathcal{U}$  be an updater that takes input a message  $M \in \Sigma^n$  and outputs an index  $i \in [n]$  and  $m \in \Sigma$ . Then for any blocks of messages  $M = (m_1, m_2, \dots, m_n) \in \Sigma^n$ , and any (non-uniform) adversary  $\mathcal{A}$ , any updater  $\mathcal{U}$ , define the following continual experiment  $\mathbf{CTamperLeak}_{\mathcal{A}, \mathcal{U}, M}$ :*

- The challenger first computes an initial encoding  $C^{(1)} \leftarrow \text{ENC}(M)$ .
- Then the following procedure repeats, at each round  $j$ , let  $C^{(j)}$  be the current codeword and  $M^{(j)}$  be the underlying message:
  - $\mathcal{A}$  sends either a tampering function  $f \in \mathcal{F}$  and/or a leakage function  $g \in \mathcal{G}$  to the challenger.
  - The challenger replaces the codeword with  $f(C^{(j)})$ , or sends back a leakage  $\ell^{(j)} = g(C^{(j)})$ .
  - We define  $\mathbf{m}^{(j)} \stackrel{\text{def}}{=} (\text{DEC}^{f(C^{(j)})}(1), \dots, \text{DEC}^{f(C^{(j)})}(n))$ .
  - Then the updater computes  $(i^{(j)}, m) \leftarrow \mathcal{U}(\mathbf{m}^{(j)})$  for the challenger.
  - Then the challenger runs  $\text{UPDATE}^{f(C^{(j)})}(i^{(j)}, m)$  and sends the index  $i^{(j)}$  to  $\mathcal{A}$ .
  - $\mathcal{A}$  may terminate the procedure at any point.
- Let  $t$  be the total number of rounds above. At the end, the experiment outputs

$$\left( \ell^{(1)}, \ell^{(2)}, \dots, \ell^{(t)}, \mathbf{m}^{(1)}, \dots, \mathbf{m}^{(t)}, i^{(1)}, \dots, i^{(t)} \right).$$

**Definition 8 (Non-malleability and Leakage Resilience against Continual Attacks).** *An  $(n, \hat{n}, p, q)$ -locally decodable and updatable coding scheme with respect to  $\Sigma, \hat{\Sigma}$  is continual non-malleable against  $\mathcal{F}$  and leakage resilient against  $\mathcal{G}$  if for all PPT (non-uniform) adversaries  $\mathcal{A}$ , and PPT updaters  $\mathcal{U}$ , there exists*

<sup>6</sup> If there is no refreshing procedure, then the adversary can eventually learn the whole codeword bit-by-bit by leakage. Thus he can learn the underlying message.

some PPT (non-uniform) simulator  $\mathcal{S}$  such that for any  $M = (m_1, \dots, m_n) \in \Sigma^n$ ,  $\mathbf{CTamperLeak}_{A,U,M}$  is (computationally) indistinguishable to the following ideal experiment  $\mathbf{Ideal}_{\mathcal{S},U,M}$ :

- The experiment proceeds in rounds. Let  $M^{(1)} = M$  be the initial message.
- At each round  $j$ , the experiment runs the following procedure:
  - At the beginning of each round,  $\mathcal{S}$  outputs  $(\ell^{(j)}, \mathcal{I}^{(j)}, \mathbf{w}^{(j)})$ , where  $\mathcal{I}^{(j)} \subseteq [n]$ .
  - Define

$$\mathbf{m}^{(j)} = \begin{cases} \mathbf{w}^{(j)} & \text{if } \mathcal{I}^{(j)} = [n] \\ \mathbf{m}^{(j)}|_{\mathcal{I}^{(j)}} := \perp, \mathbf{m}^{(j)}|_{\bar{\mathcal{I}}^{(j)}} := M^{(j)}|_{\bar{\mathcal{I}}^{(j)}} & \text{otherwise,} \end{cases}$$

where  $\mathbf{x}|_{\mathcal{I}}$  denotes the coordinates  $\mathbf{x}[v]$  where  $v \in \mathcal{I}$ , and the bar denotes the complement of a set.

- The updater runs  $(i^{(j)}, m) \leftarrow \mathcal{U}(\mathbf{m}^{(j)})$  and sends the index  $i^{(j)}$  to the simulator. Then the experiment updates  $M^{(j+1)}$  as follows: set  $M^{(j+1)} := M^{(j)}$  for all coordinates except  $i^{(j)}$ , and set  $M^{(j+1)}[i^{(j)}] := m$ .
- Let  $t$  be the total number of rounds above. At the end, the experiment outputs

$$\left( \ell^{(1)}, \ell^{(2)}, \dots, \ell^{(t)}, \mathbf{m}^{(1)}, \dots, \mathbf{m}^{(t)}, i^{(1)}, \dots, i^{(t)} \right).$$

### 3 Our Constructions

In this section, we present two constructions. As a warm-up, we first present a construction that is one-time secure to demonstrate the idea of achieving non-malleability, local decodability and updatability simultaneously. Then in the next section, we show how to make the construction secure against continual attacks. Due to space limit, please find in the full version of the paper [14] the security proofs.

#### 3.1 A First Attempt – One-time Security

*Construction.* Let  $\mathcal{E} = (\text{Gen}, \text{Encrypt}, \text{Decrypt})$  be a symmetric encryption scheme,  $\text{NMC} = (\text{ENC}, \text{DEC})$  be a coding scheme. Then we consider the following coding scheme:

- $\text{ENC}(M)$ : on input  $M = (m_1, m_2, \dots, m_n)$ , the algorithm first generates the encryption key  $\text{sk} \leftarrow \mathcal{E}.\text{Gen}(1^k)$ . Then it computes  $c \leftarrow \text{NMC}.\text{ENC}(\text{sk})$ ,  $e_i \leftarrow \mathcal{E}.\text{Encrypt}_{\text{sk}}(m_i, i)$  for  $i \in [n]$ . The algorithm finally outputs a codeword  $C = (c, e_1, e_2, \dots, e_n)$ .
- $\text{DEC}^C(i)$ : on input  $i \in [n]$ , the algorithm reads the first block and the  $(i+1)$ -st block of the codeword to retrieve  $(c, e_i)$ . Then it runs  $\text{sk} := \text{NMC}.\text{DEC}(c)$ . If the decoding algorithm outputs  $\perp$ , then it outputs  $\perp$  and terminates. Else, it computes  $(m_i, i^*) = \mathcal{E}.\text{Decrypt}_{\text{sk}}(e_i)$ . If  $i^* \neq i$ , or the decryption fails, the algorithm outputs  $\perp$ . If all the above verifications pass, the algorithm outputs  $m_i$ .

- $\text{UPDATE}(i, m')$ : on inputs an index  $i \in [n]$ , a block of message  $m' \in \Sigma$ , the algorithm runs  $\text{DEC}^C(i)$  to retrieve  $(c, e_i)$  and  $(\text{sk}, m_i, i)$ . If the decoding algorithm returns  $\perp$ , the algorithm writes  $\perp$  to the first block and the  $(i+1)$ -st block. Otherwise, it computes a fresh encoding  $c' \leftarrow \text{NMC.ENC}(\text{sk})$ , and a fresh ciphertext  $e'_i \leftarrow \mathcal{E}.\text{Encrypt}_{\text{sk}}(m', i)$ . Then it writes back the first block and the  $(i+1)$ -st block with  $(c', e'_i)$ .

To analyze the coding scheme, we make the following assumptions of the parameters in the underlying scheme for convenience:

1. The size of the encryption key is  $k$  (security parameter), i.e.  $|\text{sk}| = k$ .
2. Let  $\Sigma$  be a set, and the encryption scheme supports messages of length  $|\Sigma| + \log n$ . The ciphertexts are in the space  $\hat{\Sigma}$ .
3. The length of  $|\text{NMC.ENC}(\text{sk})|$  is less than  $|\hat{\Sigma}|$ .

Then clearly, the above coding scheme is an  $(n, n+1, 2, 2)$ -locally updatable and decodable code with respect to  $\Sigma, \hat{\Sigma}$ . The correctness of the scheme is obvious by inspection. The rate (ratio of the length of messages to that of codewords) of the coding scheme is  $1 - o(1)$ .

**Theorem 1.** *Assume  $\mathcal{E}$  is a symmetric authenticated encryption scheme, and NMC is a non-malleable code against the tampering function class  $\mathcal{F}$ . Then the coding scheme presented above is one-time non-malleable against the tampering class*

$$\bar{\mathcal{F}} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} f : \hat{\Sigma}^{n+1} \rightarrow \hat{\Sigma}^{n+1} \text{ and } |f| \leq \text{poly}(k), \text{ such that :} \\ f = (f_1, f_2), f_1 : \hat{\Sigma}^{n+1} \rightarrow \hat{\Sigma}, f_2 : \hat{\Sigma}^n \rightarrow \hat{\Sigma}^n, \\ \forall (x_2, \dots, x_{n+1}) \in \hat{\Sigma}^n, f_1(\cdot, x_2, \dots, x_{n+1}) \in \mathcal{F} \\ f(x_1, x_2, \dots, x_{n+1}) = (f_1(x_1, x_2, \dots, x_{n+1}), f_2(x_2, \dots, x_{n+1})). \end{array} \right\}.$$

We have presented the intuition in the introduction. Before giving the detailed proof, we make the following remark.

*Remark 4.* The function class  $\bar{\mathcal{F}}$  may look complex, yet the intuition is simple. The tampering function restricted in the first block (the underlying non-malleable code) falls into the class  $\mathcal{F}$  – this is captured by  $f_1 \in \mathcal{F}$ ; on the other hand, we just require the function restricted in the rest of the blocks to be polynomial-sized – this is captured by  $|f_2| \leq |f| \leq \text{poly}(k)$ .

For our construction, it is inherent that the function  $f_2$  cannot depend on  $x_1$  arbitrarily. Suppose this is not the case, then  $f_2$  can first decode the non-malleable code, encrypt the decoded value and write the ciphertext into  $x_2$ , which breaks non-malleability. However, if the underlying coding scheme is non-malleable and also leakage resilient to  $\mathcal{G}$ , then we can allow  $f_2$  to get additional information  $g(x_1)$  for any  $g \in \mathcal{G}$ . Moreover, the above construction is one-time leakage resilient.

We present the above simpler version for clarity of exposition, and give this remark that our construction actually achieves security against a broader class of tampering attacks.

### 3.2 Achieving Security against Continual Attacks

As discussed in the introduction, the above construction is not secure if continual tampering and leakage is allowed – the adversary can use a rewind attack to modify the underlying message to some old/related messages. We handle this challenge using a technique of Merkle tree, which preserves local properties of the above scheme. We present the construction in the following:

**Definition 9 (Merkle Tree).** Let  $h : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$  be a hash function that maps two blocks of messages to one.<sup>7</sup> A Merkle Tree  $\text{Tree}_h(M)$  takes input a message  $M = (m_1, m_2, \dots, m_n) \in \mathcal{X}^n$ . Then it applies the hash on each pair  $(m_{2i-1}, m_{2i})$ , and resulting in  $n/2$  blocks. Then again, it partitions the blocks into pairs and applies the hash on the pairs, which results in  $n/4$  blocks. This is repeated  $\log n$  times, resulting a binary tree with hash values, until one block remains. We call this value the root of Merkle Tree denoted  $\text{Root}_h(M)$ , and the internal nodes (including the root) as  $\text{Tree}_h(M)$ . Here  $M$  can be viewed as leaves.

**Theorem 2.** Assuming  $h$  is a collision resistant hash function. Then for any message  $M = (m_1, m_2, \dots, m_n) \in \mathcal{X}^n$  and any polynomial time adversary  $\mathcal{A}$ ,  $\Pr \left[ (m'_i, p_i) \leftarrow \mathcal{A}(M, h) : m'_i \neq m_i, p_i \text{ is a consistent path with } \text{Root}_h(M) \right] \leq \text{negl}(k)$ .

Moreover, given a path  $p_i$  passing the leaf  $m_i$ , and a new value  $m'_i$ , there is an algorithm that computes  $\text{Root}_h(M')$  in time  $\text{poly}(\log n, k)$ , where  $M' = (m_1, \dots, m_{i-1}, m'_i, m_{i+1}, \dots, m_n)$ .

*Construction.* Let  $\mathcal{E} = (\text{Gen}, \text{Encrypt}, \text{Decrypt})$  be a symmetric encryption scheme,  $\text{NMC} = (\text{ENC}, \text{DEC})$  be a non-malleable code,  $H$  is a family of collision resistance hash functions. Then we consider the following coding scheme:

- $\text{ENC}(M)$ : on input  $M = (m_1, m_2, \dots, m_n)$ , the algorithm first generates encryption key  $\text{sk} \leftarrow \mathcal{E}.\text{Gen}(1^k)$  and  $h \leftarrow H$ . Then it computes  $e_i \leftarrow \mathcal{E}.\text{Encrypt}_{\text{sk}}(m_i)$  for  $i \in [n]$ , and  $T = \text{Tree}_h(e_1, \dots, e_n), R = \text{Root}_h(e_1, \dots, e_n)$ . Then it computes  $c \leftarrow \text{NMC}.\text{ENC}(\text{sk}, R, h)$ , The algorithm finally outputs a codeword  $C = (c, e_1, e_2, \dots, e_n, T)$ .
- $\text{DEC}^C(i)$ : on input  $i \in [n]$ , the algorithm reads the first block, the  $(i+1)$ -st block, and a path  $p$  in the tree (from the root to the leaf  $i$ ), and it retrieve  $(c, e_i, p)$ . Then it runs  $(\text{sk}, R, h) = \text{NMC}.\text{DEC}(c)$ . If the decoding algorithm outputs  $\perp$ , or the path is not consistent with the root  $R$ , then it outputs  $\perp$  and terminates. Else, it computes  $m_i = \mathcal{E}.\text{Decrypt}_{\text{sk}}(e_i)$ . If the decryption fails, output  $\perp$ . If all the above verifications pass, the algorithm outputs  $m_i$ .
- $\text{UPDATE}(i, m')$ : on inputs an index  $i \in [n]$ , a block of message  $m' \in \Sigma$ , the algorithm runs  $\text{DEC}^C(i)$  to retrieve  $(c, e_i, p)$ . Then the algorithm can derive  $(\text{sk}, R, h) = \text{NMC}.\text{DEC}(c)$ . If the decoding algorithm returns  $\perp$ , the update writes  $\perp$  to the first block, which denotes failure. Otherwise, it computes a fresh ciphertext  $e'_i \leftarrow \mathcal{E}.\text{Encrypt}_{\text{sk}}(m')$ , a new path  $p'$  (that replaces  $e_i$  by

<sup>7</sup> Here we assume  $|\mathcal{X}|$  is greater than the security parameter.

$e'_i$ ) and a new root  $R'$ , which is consistent with the new leaf value  $e'_i$ . (Note that this can be done given only the old path  $p$  as Theorem 2.) Finally, it computes a fresh encoding  $c' \leftarrow \text{NMC.ENC}(\text{sk}, R', h)$ . Then it writes back the first block, the  $(i + 1)$ -st block, and the new path blocks with  $(c', e'_i, p')$ .

To analyze the coding scheme, we make the following assumptions of the parameters in the underlying scheme for convenience:

1. The size of the encryption key is  $k$  (security parameter), i.e.  $|\text{sk}| = k$  and the length of the output of the hash function is  $k$ .
2. Let  $\Sigma$  be a set, and the encryption scheme supports messages of length  $|\Sigma|$ . The ciphertexts are in the space  $\hat{\Sigma}$ .
3. The length of  $|\text{NMC.ENC}(\text{sk}, v)|$  is less than  $|\hat{\Sigma}|$ , where  $|v| = k$ .

Clearly, the above coding scheme is an  $(n, 2n + 1, O(\log n), O(\log n))$ -locally updatable and decodable code with respect to  $\Sigma, \hat{\Sigma}$ . The correctness of the scheme is obvious by inspection. The rate (ratio of the length of messages to that of codewords) of the coding scheme is  $1/2 - o(1)$ .

**Theorem 3.** *Assume  $\mathcal{E}$  is a semantically secure symmetric encryption scheme, and NMC is a non-malleable code against the tampering function class  $\mathcal{F}$ , and leakage resilient against the function class  $\mathcal{G}$ . Then the coding scheme presented above is non-malleable against continual attacks of the tampering class*

$$\bar{\mathcal{F}} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} f : \hat{\Sigma}^{2n+1} \rightarrow \hat{\Sigma}^{2n+1} \text{ and } |f| \leq \text{poly}(k), \text{ such that :} \\ f = (f_1, f_2), f_1 : \hat{\Sigma}^{2n+1} \rightarrow \hat{\Sigma}, f_2 : \hat{\Sigma}^{2n} \rightarrow \hat{\Sigma}^{2n}, \\ \forall (x_2, \dots, x_{2n+1}) \in \hat{\Sigma}^n, f_1(\cdot, x_2, \dots, x_{2n+1}) \in \mathcal{F}, \\ f(x_1, x_2, \dots, x_{2n+1}) = (f_1(x_1, x_2, \dots, x_{2n+1}), f_2(x_2, \dots, x_{2n+1})). \end{array} \right\},$$

and is leakage resilient against the class

$$\bar{\mathcal{G}} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} g : \hat{\Sigma}^{2n+1} \rightarrow \mathcal{Y} \text{ and } |g| \leq \text{poly}(k), \text{ such that :} \\ g = (g_1, g_2), g_1 : \hat{\Sigma}^{2n+1} \rightarrow \mathcal{Y}', g_2 : \hat{\Sigma}^{2n} \rightarrow \hat{\Sigma}^{2n}, \\ \forall (x_2, \dots, x_{2n+1}) \in \hat{\Sigma}^n, g_1(\cdot, x_2, \dots, x_{2n+1}) \in \mathcal{G}. \end{array} \right\}.$$

The intuition of this construction can be found in the introduction. Before giving the detailed proof, we make a remark.

*Remark 5.* Actually our construction is secure against a broader class of tampering functions. The  $f_2$  part can depend on  $g'(x_1)$  as long as the function  $g'(\cdot)$  together with the leakage function  $g_1(\cdot, x_2, \dots, x_{2n+1})$  belong to  $\mathcal{G}$ . That is, the tampering function  $f = (f_1, f_2, g')$  and the leakage function  $g = (g_1, g_2)$  satisfy the constraint  $g'(\cdot) \circ g_1(\cdot, x_2, \dots, x_{2n+1}) \in \mathcal{G}$  (Here we use  $\circ$  to denote concatenation). For presentation clarity, we choose to describe the simpler but slightly smaller class of functions.

### 3.3 Instantiations

In this section, we describe several constructions of non-malleable codes against different classes of tampering/leakage functions. To our knowledge, we can use



the explicit constructions (of the non-malleable codes) in the work [1, 3, 9, 22, 23, 25, 42].

First we overview different classes of tampering/leakage function allowed for these results: the constructions of [22] work for bit-wise tampering functions, and split-state functions in the random oracle model. The construction of Choi et al. [9] works for small block tampering functions. The construction of Liu and Lysyanskaya [42] achieves both tamper and leakage resilience against split-state functions in the common reference string (CRS) model. The construction of Dziembowski et al. [20] achieves information theoretic security against split-state tampering functions, but their scheme can only support encoding for bits, so it cannot be used in our construction. The subsequent construction by Aggarwal et al. [1] achieves information theoretic security against split-state tampering without CRS. We believe that their construction also achieves leakage resilience against some length bounded split-state leakage yet their paper did not claim it. The construction by Faust et al. [25] is non-malleable against small-sized tampering functions. Another construction by Faust et al. [23] achieves both tamper and leakage resilience in the split-state model with CRS. The construction of Aggarwal et al. [3] is non-malleable against permutation functions.

Then we remark that actually there are other non-explicit constructions: Cheraghchi and Guruswami [8] showed the relation non-malleable codes and non-malleable two source extractors (but constructing a non-malleable two-source extractor is still open), and in another work Cheraghchi and Guruswami [7] showed the existence of high rate non-malleable codes in the split-state model but did not give an explicit (efficient) construction.

Finally, we give a concrete example of what the resulting class looks like using the construction of Liu and Lysyanskaya [42] as building block. Recall that their construction achieves both tamper and leakage resilience for split-state functions. Our construction has the form:

$$(\text{NMC.ENC}(\text{sk}, h, T), \text{Encrypt}(m_1), \dots, \text{Encrypt}(m_n), T)$$

So the overall leakage function  $g$  restricted in the first block (i.e.  $g_1$ ) can be a (poly-sized) length-bounded split-state function;  $g$  on the other hand, can leak all the other parts. For the tampering, the overall tampering function  $f$  restricted in the first block (i.e.  $f_1$ ) can be any (poly-sized) split-state function. On the other hand  $f$  restricted in the rest (i.e.  $f_2$ ) can be just any poly-sized function. We also remark that  $f_2$  can depend on a split-state leakage on the first part, say  $g_1$ , as we discussed in the previous remark above.

## 4 Tamper and Leakage Resilient RAM

In this section, we first introduce the notations of the Random Access Machine (RAM) model of computation in presence of tampering and leakage attacks in Section 4.1. Then we define the security of tamper and leakage resilient RAM model of computation in Section 4.2, and then give a construction in Section 4.3. Due to space limit, please find in the full version of the paper [14], the building

block Oblivious RAM (ORAM), more detailed construction, and the security analysis.

#### 4.1 Random Access Machines

We consider RAM programs to be interactive stateful systems  $\langle \Pi, \text{state}, D \rangle$ , where  $\Pi$  denotes a next instruction function,  $\text{state}$  the current state stored in registers, and  $D$  the content of memory. Upon  $\text{state}$  and an input value  $d$ , the next instruction function outputs the next instruction  $I$  and an updated state  $\text{state}'$ . The initial state of the RAM machine,  $\text{state}$ , is set to  $(\text{start}, *)$ . For simplicity we often denote RAM program as  $\langle \Pi, D \rangle$ . We consider four ways of interacting with the system:

- **Execute**( $x$ ): A user can provide the system with **Execute**( $x$ ) queries, for  $x \in \{0, 1\}^u$ , where  $u$  is the input length. Upon receiving such query, the system computes  $(y, t, D') \leftarrow \langle \Pi, D \rangle(x)$ , updates the state of the system to  $D := D'$  and outputs  $(y, t)$ , where  $y$  denotes the output of the computation and  $t$  denotes the time (or number of executed instructions). By  $\text{Execute}_1(x)$  we denote the first coordinate of the output of **Execute**( $x$ ).
- **doNext**( $x$ ): A user can provide the system with **doNext**( $x$ ) queries, for  $x \in \{0, 1\}^u$ . Upon receiving such query, if  $\text{state} = (\text{start}, *)$ , set  $\text{state} := (\text{start}, x)$ , and  $d := 0^r$ ; Here  $\rho = |\text{state}|$  and  $r = |d|$ . The system does the following until termination:
  1. Compute  $(I, \text{state}') = \Pi(\text{state}, d)$ . Set  $\text{state} := \text{state}'$ .
  2. If  $I = (\text{wait})$  then set  $\text{state} := 0^\rho$ ,  $d := 0^r$  and terminate.
  3. If  $I = (\text{stop}, z)$  then set  $\text{state} := (\text{start}, *)$ ,  $d := 0^r$  and terminate with output  $z$ .
  4. If  $I = (\text{write}, v, d')$  then set  $D[v] := d'$ .
  5. If  $I = (\text{read}, v, \perp)$  then set  $d := D[v]$ .

Let  $I_1, \dots, I_\ell$  be the instructions executed by **doNext**( $x$ ). All memory addresses of executed instructions are returned to the user. Specifically, for instructions  $I_j$  of the form  $(\text{read}, v, \perp)$  or  $(\text{write}, v, d')$ ,  $v$  is returned.
- **Tamper**( $f$ ): We also consider tampering attacks against the system, modeled by **Tamper**( $f$ ) commands, for functions  $f$ . Upon receiving such command, the system sets  $D := f(D)$ .
- **Leak**( $g$ ): We finally consider leakage attacks against the system, modeled by **Leak**( $g$ ) commands, for functions  $g$ . Upon receiving such command, the value of  $g(D)$  is returned to the user.

*Remark 6.* A **doNext**( $x$ ) instruction groups together instructions performed by the CPU in a single clock cycle. Intuitively, a **(wait)** instruction indicates that a clock cycle has ended and the CPU waits for the adversary to increment the clock. In contrast, a **(stop, z)** instruction indicates that the entire execution has concluded with output  $z$ . In this case, the internal state is set back to the start state.

We require that each  $\text{doNext}(x)$  operation performs exactly  $\ell = \ell(k) = \text{poly}(k)$  instructions  $I_1, \dots, I_\ell$  where: The final instruction is of the form  $I_\ell = (\text{stop}, \cdot)$  or  $I_\ell = (\text{wait})$ . For fixed  $\ell_1 = \ell_1(k), \ell_2 = \ell_2(k)$  such that  $\ell_1 + \ell_2 = \ell - 1$ , we have that the first  $\ell_1$  instructions are of the form  $I_\ell = (\text{read}, \cdot, \perp)$  and the next  $\ell_2$  instructions are of the form  $I_\ell = (\text{write}, v, d')$ . We assume that  $\ell, \ell_1, \ell_2$  are implementation-specific and public. The limitations on space are meant to model the fact that the CPU has a limited number of registers and that no persistent state is kept by the CPU between clock cycles.

*Remark 7.* We note that  $\text{Execute}(x)$  instructions are used by the ideal world adversary—who learns only the input-output behavior of the RAM machine and the run time—as well as by the real world adversary. The real world adversary may also use the more fine-grained  $\text{doNext}(x)$  instruction. We note that given access to the  $\text{doNext}(x)$  instruction, the behavior of the  $\text{Execute}(x)$  instruction may be simulated.

*Remark 8.* We note that our model does not explicitly allow for leakage and tampering on instructions  $I$ . E.g. when an instruction  $I = (\text{write}, v, d')$  is executed, we do not directly allow tampering with the values  $v, d'$  or leakage on  $d'$  (note that  $v$  is entirely leaked to the adversary). Nevertheless, as discussed in the introduction, since we allow full leakage on the addresses, the adversary can use the tampering and leakage attacks on the memory to capture the attacks on the instructions. We defer a rigorous treatment and analysis of such attacks to future work. In this work, for simplicity of presentation we assume these instructions are not subject to direct attacks.

## 4.2 Tamper and Leakage-Resilient (TLR) RAM

A tamper and leakage resilient (TLR) RAM compiler consists of two algorithms ( $\text{CompMem}, \text{CompNext}$ ), which transform a RAM program  $\langle \Pi, D \rangle$  into another program  $\langle \widehat{\Pi}, \widehat{D} \rangle$  as follows: On input database  $D$ ,  $\text{CompMem}$  initializes the memory and internal state of the compiled machine, and generates the transformed database  $\widehat{D}$ ; On input next instruction function  $\Pi$ ,  $\text{CompNext}$  generates the next instruction function of the compiled machine.

**Definition 10.** A TLR compiler ( $\text{CompMem}, \text{CompNext}$ ) is tamper and leakage simulatable w.r.t. function families  $\mathcal{F}, \mathcal{G}$ , if for every RAM next instruction function  $\Pi$ , and for any PPT (non-uniform) adversary  $\mathcal{A}$  there exists a PPT (non-uniform) simulator  $\mathcal{S}$  such that for any initial database  $D \in \{0, 1\}^{\text{poly}(k)}$  we have

$$\text{TamperExec}(\mathcal{A}, \mathcal{F}, \mathcal{G}, \langle \text{CompNext}(\Pi), \text{CompMem}(D) \rangle) \approx \text{IdealExec}(\mathcal{S}, \langle \Pi, D \rangle)$$

where  $\text{TamperExec}$  and  $\text{IdealExec}$  are defined as follows:

- $\text{TamperExec}(\mathcal{A}, \mathcal{F}, \mathcal{G}, \langle \text{CompNext}(\Pi), \text{CompMem}(D) \rangle)$ : The adversary  $\mathcal{A}$  interacts with the system  $\langle \text{CompNext}(\Pi), \text{CompMem}(D) \rangle$  for arbitrarily many rounds of interactions where, in each round:

1. The adversary can “tamper” by executing a  $\text{Tamper}(f)$  command against the system, for some  $f \in \mathcal{F}$ .
2. The adversary can “leak” by executing a  $\text{Leak}(g)$  command against the system, and receiving  $g(D)$  in return.
3. The adversary requests a  $\text{doNext}(x)$  command to be executed by the system. Let  $I_1, \dots, I_\ell$  be the instructions executed by  $\text{doNext}(x)$ . If  $I_\ell$  is of the form  $(\text{stop}, z)$  then output  $z$  is returned to the adversary. Moreover, all memory addresses corresponding to instructions  $I_1, \dots, I_{\ell-1}$  are returned to the adversary.

The output of the game consists of the output of the adversary  $\mathcal{A}$  at the end of the interaction, along with (1) all input-output pairs  $(x_1, y_1), (x_2, y_2), \dots$ , (2) all responses to leakage queries  $\ell_1, \ell_2, \dots$  (3) all outputs of  $\text{doNext}(x_1), \text{doNext}(x_2), \dots$ .

- **IdealExec**( $\mathcal{S}, \langle \Pi, D \rangle$ ): The simulator interacts with the system  $\langle \Pi, D \rangle$  for arbitrarily many rounds of interaction where, in each round, it runs an  $\text{Execute}(x)$  query for some  $x \in \{0, 1\}^u$  and receives output  $(y, t)$ . The output of the game consists of the output of the simulator  $\mathcal{S}$  at the end of the interaction, along with all of the execute-query inputs and outputs.

For simplicity of exposition, we assume henceforth that the next instruction function  $\Pi$  to be compiled is the universal RAM next instruction function. In other words, we assume that the program to be executed is stored in the initial database  $D$ .

### 4.3 TLR-RAM Construction

Here we first give a high-level description of our construction and then state our theorem. More detailed construction and the security proof will be given in the full version of the paper [14].

*High-level Description of Construction* Let  $D$  be the initial database and let  $\text{ORAM} = (\text{oCompMem}, \text{oCompNext})$  be an ORAM compiler. Let  $\text{NMCode} = (\text{ENC}, \text{DEC}, \text{UPDATE})$  be a locally decodable and updatable code. We present the following construction  $\text{TLR-RAM} = (\text{CompMem}, \text{CompNext})$  of a tamper and leakage resilient RAM compiler. In order to make our presentation more intuitive, instead of specifying the next message function  $\text{CompNext}(\Pi)$ , we specify the pseudocode for the  $\text{doNext}(x)$  instruction of the compiled machine. We note that  $\text{CompNext}(\Pi)$  is implicitly defined by this description.

TLR-RAM takes as input an initial database  $D$  and a next instruction function  $\Pi$  and does the following:

- **CompMem**: On input security parameter  $k$  and initial database  $D$ , **CompMem** does:
  - Compute  $\tilde{D} \leftarrow \text{oCompMem}(D)$ , and output  $\hat{D} \leftarrow \text{ENC}(\tilde{D})$ .
  - Initialize the ORAM state  $\text{state}_{\text{ORAM}} := (\text{start}, *)$  and  $d_{\text{ORAM}} := 0^r$ , where  $r = |d_{\text{ORAM}}|$ .

- $\text{doNext}(x)$ : On input  $x$ , do the following until termination:
  1. If  $d_{\text{ORAM}} = \perp$  then abort.
  2. Compute  $(I, \text{state}'_{\text{ORAM}}) \leftarrow \text{oCompNext}(II)(\text{state}_{\text{ORAM}}, d_{\text{ORAM}})$  and set  $\text{state}_{\text{ORAM}} := \text{state}'_{\text{ORAM}}$ .
  3. If  $I = (\text{wait})$  then set  $\text{state}_{\text{ORAM}} := 0^\rho$  and  $d_{\text{ORAM}} := 0^r$  and terminate. Here  $\rho = |\text{state}_{\text{ORAM}}|$  and  $r = |d_{\text{ORAM}}|$ .
  4. If  $I = (\text{stop}, z)$  then set  $\text{state}_{\text{ORAM}} := (\text{start}, *)$ ,  $d := 0^r$  and terminate with output  $z$ .
  5. If  $I = (\text{write}, v, d')$  then run  $\text{UPDATE}^{\widehat{D}}(v, d')$ .
  6. If  $I = (\text{read}, v, \perp)$  then set  $d_{\text{ORAM}} := \text{DEC}^{\widehat{D}}(v)$ .

We are now ready to present the main theorem of this section:

**Theorem 4.** *Assume  $\text{ORAM} = (\text{oCompMem}, \text{oCompNext})$  is an ORAM compiler which is access-pattern hiding and assume  $\text{NMCode} = (\text{ENC}, \text{DEC}, \text{UPDATE})$  is a locally decodable and updatable code which is continual non-malleable against  $\mathcal{F}$  and leakage resilient against  $\mathcal{G}$ . Then  $\text{TLR-RAM} = (\text{CompMem}, \text{CompNext})$  presented above is tamper and leakage simulatable w.r.t. function families  $\mathcal{F}, \mathcal{G}$ .*

**Acknowledgement.** We thank Yevgeniy Dodis for helpful discussions. This research was funded in part by an NSF grant CNS-1314857, a subcontract from the DARPA PROCEED program, a Sloan Research Fellowship, and Google Faculty Research Awards. The views and conclusions contained herein are those of the authors and should not be interpreted as representing funding agencies.

## References

1. Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In *STOC*, 2014. <http://eprint.iacr.org/2013/201>. 2, 7, 17
2. Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, and Pankaj Rohatgi. The EM side-channel(s). In Burton S. Kaliski Jr., Çetin Kaya Koç, and Christof Paar, editors, *CHES 2002*, volume 2523 of *LNCS*, pages 29–45. Springer, August 2002. 2
3. Shashank Agrawal, Divya Gupta, Hemanta K. Maji, Omkant Pandey, and Manoj Prabhakaran. Explicit non-malleable codes resistant to permutations. In *Cryptology ePrint Archive, Report 2014/316*, 2014. 1, 2, 7, 17
4. Eli Biham and Adi Shamir. Differential fault analysis of secret key cryptosystems. In Burton S. Kaliski Jr., editor, *CRYPTO'97*, volume 1294 of *LNCS*, pages 513–525. Springer, August 1997. 2
5. Dan Boneh, Richard A. DeMillo, and Richard J. Lipton. On the importance of eliminating errors in cryptographic computations. *Journal of Cryptology*, 14(2):101–119, 2001. 2
6. Nishanth Chandran, Bhavana Kanukurthi, and Rafail Ostrovsky. Locally updatable and locally decodable codes. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 489–514. Springer, February 2014. 2

7. Mahdi Cheraghchi and Venkatesan Guruswami. Capacity of non-malleable codes. In Moni Naor, editor, *ITCS 2014*, pages 155–168. ACM, January 2014. [2](#), [7](#), [17](#)
8. Mahdi Cheraghchi and Venkatesan Guruswami. Non-malleable coding against bit-wise and split-state tampering. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 440–464. Springer, February 2014. [2](#), [7](#), [17](#)
9. Seung Geol Choi, Aggelos Kiayias, and Tal Malkin. BiTR: Built-in tamper resilience. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 740–758. Springer, December 2011. [2](#), [7](#), [17](#)
10. Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. Private information retrieval. *J. ACM*, 45(6):965–981, 1998. [2](#)
11. Sandro Coretti, Ueli Maurer, Bjorn Tackmann, and Daniele Venturi. From single-bit to multi-bit public-key encryption via non-malleable codes. In *Cryptology ePrint Archive, Report 2014/324*, 2014. [1](#)
12. Dana Dachman-Soled and Yael Tauman Kalai. Securing circuits against constant-rate tampering. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 533–551. Springer, August 2012. [2](#), [8](#)
13. Dana Dachman-Soled and Yael Tauman Kalai. Securing circuits and protocols against  $1/\text{poly}(k)$  tampering rate. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 540–565. Springer, February 2014. [2](#), [8](#)
14. Dana Dachman-Soled, Feng-Hao Liu, Elaine Shi, and Hong-Sheng Zhou. Locally decodable and updatable non-malleable codes and their applications. *Cryptology ePrint Archive, Report 2014/663*, 2014. <http://eprint.iacr.org/2014/663>. [13](#), [17](#), [20](#)
15. Ivan Damgård, Sebastian Faust, Pratyay Mukherjee, and Daniele Venturi. Bounded tamper resilience: How to go beyond the algebraic barrier. In Kazuo Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 140–160. Springer, December 2013. [2](#)
16. Yevgeniy Dodis and Krzysztof Pietrzak. Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 21–40. Springer, August 2010. [2](#), [8](#)
17. Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 423–440. Springer, May 2014. [2](#), [8](#)
18. Stefan Dziembowski and Sebastian Faust. Leakage-resilient cryptography from the inner-product extractor. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 702–721. Springer, December 2011. [2](#), [8](#)
19. Stefan Dziembowski and Sebastian Faust. Leakage-resilient circuits without computational assumptions. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 230–247. Springer, March 2012. [2](#), [8](#)
20. Stefan Dziembowski, Tomasz Kazana, and Maciej Obremski. Non-malleable codes from two-source extractors. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 239–257. Springer, August 2013. [2](#), [7](#), [17](#)
21. Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pages 293–302. IEEE Computer Society Press, October 2008. [2](#), [8](#)
22. Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In Andrew Chi-Chih Yao, editor, *ICS 2010*, pages 434–452. Tsinghua University Press, January 2010. [1](#), [2](#), [3](#), [7](#), [9](#), [17](#)

23. Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. Continuous non-malleable codes. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 465–488. Springer, February 2014. [2](#), [7](#), [8](#), [12](#), [17](#)
24. Sebastian Faust, Pratyay Mukherjee, Jesper Buus Nielsen, and Daniele Venturi. A tamper and leakage resilient random access machine. In *Cryptology ePrint Archive, Report 2014/338*, 2014. [8](#)
25. Sebastian Faust, Pratyay Mukherjee, Daniele Venturi, and Daniel Wichs. Efficient non-malleable codes and key-derivation for poly-size tampering circuits. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 111–128. Springer, May 2014. [2](#), [7](#), [17](#)
26. Sebastian Faust, Krzysztof Pietrzak, and Daniele Venturi. Tamper-proof circuits: How to trade leakage for tamper-resilience. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *ICALP 2011, Part I*, volume 6755 of *LNCS*, pages 391–402. Springer, July 2011. [2](#), [8](#)
27. Sebastian Faust, Tal Rabin, Leonid Reyzin, Eran Tromer, and Vinod Vaikuntanathan. Protecting circuits from leakage: the computationally-bounded and noisy cases. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 135–156. Springer, May 2010. [2](#), [8](#)
28. Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Algorithmic tamper-proof (ATP) security: Theoretical foundations for security against hardware tampering. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 258–277. Springer, February 2004. [2](#)
29. Oded Goldreich and Rafail Ostrovsky. Software protection and simulation on oblivious RAMs. *Journal of the ACM*, 43(3):431–473, 1996. [7](#)
30. Shafi Goldwasser and Guy N. Rothblum. Securing computation against continuous leakage. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 59–79. Springer, August 2010. [2](#), [8](#)
31. Shafi Goldwasser and Guy N. Rothblum. How to compute in the presence of leakage. In *53rd FOCS*, pages 31–40. IEEE Computer Society Press, October 2012. [2](#), [8](#)
32. J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Lest we remember: Cold boot attacks on encryption keys. In *USENIX Security Symposium*, pages 45–60, 2008. [2](#)
33. Yuval Ishai and Eyal Kushilevitz. On the hardness of information-theoretic multiparty computation. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 439–455. Springer, May 2004. [2](#)
34. Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and David Wagner. Private circuits II: Keeping secrets in tamperable circuits. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 308–327. Springer, May / June 2006. [2](#), [8](#)
35. Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 463–481. Springer, August 2003. [2](#), [8](#)
36. Ali Juma and Yevgeniy Vahlis. Protecting cryptographic keys against continual leakage. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 41–58. Springer, August 2010. [2](#), [8](#)
37. Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *32nd ACM STOC*, pages 80–86. ACM Press, May 2000. [2](#)

38. Aggelos Kiayias and Yiannis Tselekounis. Tamper resilient circuits: The adversary at the gates. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 161–180. Springer, December 2013. 2, 8
39. Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 104–113. Springer, August 1996. 2
40. Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 388–397. Springer, August 1999. 2
41. David Lie, Chandramohan A. Thekkath, Mark Mitchell, Patrick Lincoln, Dan Boneh, John C. Mitchell, and Mark Horowitz. Architectural support for copy and tamper resistant software. In *ASPLOS*, pages 168–177, 2000. 8
42. Feng-Hao Liu and Anna Lysyanskaya. Tamper and leakage resilience in the split-state model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 517–532. Springer, August 2012. 2, 4, 6, 7, 9, 12, 17
43. Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 278–296. Springer, February 2004. 2, 8
44. Krzysztof Pietrzak. A leakage-resilient mode of operation. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 462–482. Springer, April 2009. 2, 8
45. Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In Ehab Al-Shaer, Somesh Jha, and Angelos D. Keromytis, editors, *ACM CCS 09*, pages 199–212. ACM Press, November 2009. 2
46. Guy N. Rothblum. How to compute under  $\neg$ <sup>0</sup> leakage without secure hardware. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 552–569. Springer, August 2012. 2, 8
47. G. Edward Suh, Dwaine E. Clarke, Blaise Gassend, Marten van Dijk, and Srinivas Devadas. AEGIS: architecture for tamper-evident and tamper-resistant processing. In *Proceedings of the 17th Annual International Conference on Supercomputing, ICS 2003*, pages 160–171, 2003. 8
48. Amit Vasudevan, Jonathan M. McCune, James Newsome, Adrian Perrig, and Leendert van Doorn. CARMA: a hardware tamper-resistant isolated execution environment on commodity x86 platforms. In Heung Youl Youm and Yoojae Won, editors, *ASIACCS 12*, pages 48–49. ACM Press, May 2012. 8
49. Sergey Yekhanin. Locally decodable codes. *Foundations and Trends in Theoretical Computer Science*, 6(3):139–255, 2012. 2