

# ELM: A Low-Latency and Scalable Memory Encryption Scheme

Akiko Inoue<sup>1</sup>, Kazuhiko Minematsu<sup>2</sup>, Maya Oda, Rei Ueno, *Member, IEEE*,  
and Naofumi Homma, *Senior Member, IEEE*

**Abstract**—Memory encryption (ME) with authentication is becoming a key security feature of modern processors, as evident by the adoption of ME by Intel’s SGX. Recently ME is actively studied from the viewpoint of system architecture. This paper studies ME from the viewpoint of symmetric-key cryptographic designs, with a primal focus on latency. A significant progress in such a direction can be observed in the SGX Integrity Tree (SIT). Using a variant of AES-GCM, SIT achieves an excellent latency. However, it has a scalability issue. By carefully examining SIT, we develop a new ME scheme dubbed ELM. We present an AES-based instantiation of ELM, and show that ELM significantly reduces latency from SIT for large memories, and achieves the provable security and equivalent hardware-protected (on-chip) area. We also present preliminary hardware implementations to substantiate our advantages.

**Index Terms**—Memory encryption, authentication tree, latency, mode of operations, SGX.

## I. INTRODUCTION

MEMORY encryption (ME) is widely deployed in modern systems. One typical method is sector-wise encryption, such as XTS [1]. A sector-wise encryption scheme encrypts each memory sector in an independent and deterministic manner, keeping the key in an on-chip area that is physically protected. This prevents passive off-line attacks that try to extract the data from the storage devices, such as [2]. However, it is vulnerable against active online attacks and replay attacks, for the lack of authenticity. If we encrypt each sector using a nonce-based authenticated encryption (AE) and store all the nonces on-chip, it would protect against active attacks. However, this would incur a linear increase of the on-chip area. This is usually impractical because the on-chip area is much more expensive than the main (off-chip) memory.

A well-known classical solution to this problem is to use an authentication tree, also known as a Merkle Hash Tree [3]. By involving any unit memory data in the tree computation and storing the root hash value on-chip, the authenticity

against active attacks is guaranteed. Instead of a cryptographic hash function, we can use a message authentication code (MAC) to build an authentication tree. Merkle tree and its (possibly MAC-based) improvements, such as PAT [4] and Bonsai Tree [5], provide an authenticity of the whole memory with a constant on-chip memory overhead, at the cost of a logarithmic computation overhead for read and write operations. Confidentiality of the memory can be achieved by an additional encryption mechanism, *e.g.*, by TEC-tree [6]. Due to the increasing threat of active attacks, tree-based ME schemes, often with a confidentiality mechanism, are gradually being deployed in real-world memory/storage systems. One prominent example is Intel’s SGX [7], [8], which adopts a variant of PAT with a dedicated AES-based MAC and AE schemes similar to GMAC and GCM [9]. The widespread use of non-volatile memory also pushes the need for ME with authenticity.

Latency is a very important criterion for the aforementioned ME schemes. Merkle tree can reduce latency by utilizing parallelizability, but only for verification (read operation). When one wants to replace a piece of data with a new one (write operation), Merkle tree needs to serially update all hash values on the path from the leaf (data) to the root. PAT is the current state-of-the-art in this respect, as it is parallelizable for both read and write operations by means of a clever use of nonce-based MAC.

Tree-based MEs have been initially studied from the cryptography community, however, it recently receives significant attention from the system architecture community, such as [5], [10]–[12]. The primary focus of these studies is the data structure, such as the parameters/structures of integrity trees [5], [11] and counter/nonce representations [10], [12] that fits well into the target architecture, and cryptographic components are often considered as black boxes or instantiated by just picking a standard, *e.g.*, the use of GCM in [10]. A notable exception is the aforementioned ME inside SGX, which is also called SGX Integrity Tree (SIT). It develops dedicated AE and MAC schemes based on AES-GCM, with particular attention to latency in mind. SIT is quite efficient and enables a very low-latency read/write operation on the given tree structure that covers up to 96 Mbyte of memory on an x86 platform. Moreover, as an important subsystem of SGX, it is also quite widely deployed in practice.

In contrast to the numerous ME proposals from the system architecture community, those from cryptography community are rather scarce after [4], [6]. As mentioned above, SIT is a notable exception, however its on-chip data size is linear

Manuscript received 6 December 2021; revised 20 April 2022; accepted 9 June 2022. Date of publication 4 July 2022; date of current version 22 July 2022. This work was supported in part by the Japan Science and Technology Agency (JST) Core Research for Evolutional Science and Technology (CREST), Japan, under Grant JPMJCR19K5. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Nele Mentens. (*Corresponding author: Akiko Inoue.*)

Akiko Inoue is with NEC, Kawasaki 211-8666, Japan (e-mail: a\_inoue@nec.com).

Kazuhiko Minematsu is with NEC, Kawasaki 211-8666, Japan, and also with the Graduate School of Environment and Information Sciences, Yokohama National University, Kanagawa 240-8501, Japan.

Maya Oda, Rei Ueno, and Naofumi Homma are with the Research Institute of Electrical Communication, Tohoku University, Sendai 980-8577, Japan.

Digital Object Identifier 10.1109/TIFS.2022.3188146

to the unit data size. This poses a limitation on the amount of covered memory sizes and hence is not scalable. In fact, VAULT [11] and Morphable Counter [12] are two recent MEs that aim at covering much larger memory than SIT and improving the performance, mainly by the system architecture approach.

This paper proposes a new ME scheme dubbed ELM (Encryption for Large Memory), which enables a significantly low latency for large amounts of memory. It achieves on-chip and off-chip memory overheads comparable to existing schemes, which becomes possible by closely inspecting the required operations in SIT and its baseline PAT. ELM adopts several techniques known in the symmetric-key cryptography literature. Specifically, we find that *incremental MAC* (Inc-MAC) [13], [14] can be used securely and it works quite effectively for the trees with large arity, a common feature for the modern schemes targeting large memories. By using an Inc-MAC at the internal nodes, ELM significantly reduces the write latency without harming the read latency. Our MAC is a variant of the classical XOR-MAC [15] that optimizes the latency, number of primitive calls, and security.

Another key component of ELM is a new low-latency AE scheme. It is basically a variant of OCB [16]. OCB is already quite good in terms of latency and parallelizability, because it does not need an additional authentication function unlike GCM. However, the decryption latency of OCB is not sufficiently small due to its structure. By changing the structure, our AE mode has a smaller latency than the original for decryption, while retaining the other main features such as parallelizability and provable security. In particular, when it is viewed as a mode of a tweakable block cipher (TBC) [17], the latency is *optimally small* for both encryption and decryption. ELM is obtained by combining our MAC, our AE, and a variant of tree structure of SIT which we call PAT2. Each technique itself is not ultimately novel. However, we show how to combine them in an optimal manner to reduce latency and computation, which is, to the best of our knowledge, not known in the literature. Our proposal is generic in principle, and it can be instantiated by any block cipher or TBC. Moreover, we prove that the security of ELM (and PAT2) is reduced to the pseudorandomness of the block cipher we use, which is the first *formal* security treatment of ME with authenticity.

To showcase the effectiveness of ELM, we specify an instantiation of ELM using the same components as SIT, namely AES-128 and a full 64-bit field multiplier.<sup>1</sup> We compare them with (a generalized variant of) SIT for various memory sizes and tree parameters under a certain practical implementation setting. Our results show that ELM has a smaller latency than SIT for most of the cases,<sup>2</sup> and as memory size gets larger, the difference becomes significant. We also conducted preliminary ASIC implementations, and show that the total implementation size is comparable to that of SIT. In addition, we discuss the

optimization of hardware implementations for our proposal depending on the system constraints.

## II. PRELIMINARIES

### A. Notation

For a positive integer  $n$ ,  $\{0, 1\}^n$  denotes the set of  $n$ -bit strings and let  $[n] = \{1, \dots, n\}$ .  $X \stackrel{\$}{\leftarrow} \mathcal{X}$  means that the variable  $X$  is uniformly sampled over the set  $\mathcal{X}$ . For binary strings  $A$  and  $B$ ,  $A \parallel B$  denotes the concatenation of  $A$  and  $B$ . The bit length of  $A$  is denoted by  $|A|$ , and  $|A|_n := \lceil |A|/n \rceil$ . Dividing a string  $A$  into blocks of  $n$  bits is denoted by  $A[1] \parallel \dots \parallel A[m] \stackrel{n}{\leftarrow} A$ , where  $m = |A|_n$  and  $|A[i]| = n$ ,  $|A[m]| \leq n$  for  $1 \leq i \leq m-1$ . For  $t \in [|A|]$ ,  $\text{msb}_t(A)$  ( $\text{lsb}_t(A)$ ) denotes the first (last)  $t$  bits of  $A$ . A sequence of  $i$  zeros is written as  $0^i$ . For a function  $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  with the key space  $\mathcal{K}$ ,  $F(K, \cdot)$  may be written as  $F_K(\cdot)$ . Let  $\text{GF}(2^n)$  be a finite field of size  $2^n$  with characteristic 2 and extension degree  $n$ . We focus on the case  $n = 128$ . Following [18], we use  $(x^{128} + x^7 + x^2 + x + 1)$  for defining the field  $\text{GF}(2^{128})$ . Here, the primitive root  $x$  is interpreted as 2 in the decimal representation. For  $a \in \text{GF}(2^n)$ , let  $2a$  denote a multiplication by  $x$  and  $a$ , which is also called doubling [18]. Similarly, let  $3a$  denote  $2a \oplus a$ .

### B. (Tweakable) Block Cipher

Let  $\mathcal{K}$  and  $\mathcal{M}$  be the set of keys and messages, respectively. Let  $\mathcal{TW}$  be the set of tweaks, where a tweak is a public parameter. A tweakable block cipher (TBC) [17] is a function  $\tilde{E} : \mathcal{K} \times \mathcal{TW} \times \mathcal{M} \rightarrow \mathcal{M}$  s.t.  $\tilde{E}_K(Tw, \cdot)$  is a permutation on  $\mathcal{M}$  for  $\forall (K, Tw) \in \mathcal{K} \times \mathcal{TW}$ . It is also denoted by  $\tilde{E}_K^{Tw}$ ,  $\tilde{E}^{Tw}$ , or  $\tilde{E}$ . If  $\mathcal{TW}$  is singleton, it means a plain block cipher  $E_K$ . A TBC can be built on a block cipher [17], [18]. A block cipher  $E$  or a TBC  $\tilde{E}$  is said to be secure if it is computationally hard to distinguish from the ideal primitive with oracle access. Let  $\text{Perm}(n)$  denote the set of all permutations on  $\{0, 1\}^n$ . An  $n$ -bit tweakable permutation of  $t$ -bit tweak is a function  $\pi : \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  s.t.  $\pi(Tw, \cdot) \in \text{Perm}(n)$  for  $\forall Tw \in \{0, 1\}^t$ . The set of all  $n$ -bit tweakable permutations with  $t$ -bit tweak is denoted by  $\text{TPerm}(t, n)$ . Let  $\mathbf{P} \stackrel{\$}{\leftarrow} \text{Perm}(n)$  be a uniform random permutation (URP) and  $\tilde{\mathbf{P}} \stackrel{\$}{\leftarrow} \text{TPerm}(t, n)$  be a tweakable URP (TURP). Let  $\mathcal{A}$  be an adversary who (possibly adaptively) queries an oracle  $O$  and subsequently outputs a bit. We write  $\Pr[\mathcal{A}^O \rightarrow 1]$  to denote the probability that this bit is 1. The advantages against  $\tilde{E}$  are defined as  $\text{Adv}_{\tilde{E}}^{\text{URP}}(\mathcal{A}) := |\Pr[\mathcal{A}^{\tilde{E}} \rightarrow 1] - \Pr[\mathcal{A}^{\mathbf{P}} \rightarrow 1]|$ , and  $\text{Adv}_{\tilde{E}}^{\text{TURP}}(\mathcal{A}^\pm) := |\Pr[(\mathcal{A}^\pm)^{\tilde{E}, \tilde{E}^{-1}} \rightarrow 1] - \Pr[(\mathcal{A}^\pm)^{\tilde{\mathbf{P}}, \tilde{\mathbf{P}}^{-1}} \rightarrow 1]|$ , where  $\tilde{E}^{-1}$  and  $\tilde{\mathbf{P}}^{-1}$  are decryption functions of  $\tilde{E}$  and  $\tilde{\mathbf{P}}$ . When the advantage is sufficiently small,  $\tilde{E}$  is said to be secure against the underlying adversary.

### C. Message Authentication Code

Message authentication code (MAC) is a symmetric-key function for message authenticity. We consider *nonce-based* MAC, where a nonce is a non-repeating value used together with a message. For the nonce space  $\mathcal{N}$  and the tag space  $\mathcal{T}$ , a nonce-based MAC scheme  $\text{MAC}$  consists of two functions: the tagging function  $\text{MAC}.\mathcal{T} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \rightarrow \mathcal{T}$  and

<sup>1</sup>We also use a 128-bit multiplier, but with a very small input size.

<sup>2</sup>This holds true even when SIT adopts a part of our idea of using Inc-MAC. See Section V-B.

the verification function  $\text{MAC.V} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \times \mathcal{T} \rightarrow \{\top, \perp\}$ . A tag for  $(K, N, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{M}$  is derived as  $T = \text{MAC.T}_K(N, M)$ . The tuple  $(N, M, T)$  is considered to be authentic if  $\text{MAC.V}_K(N, M, T) = \top$ , and otherwise it is rejected. The security of MAC is defined as the probability that an adversary  $\mathcal{A}$  creates a successful forgery by accessing  $\text{MAC.T}_K$  and  $\text{MAC.V}_K$ . The security measure is  $\text{Adv}_{\text{MAC}}^{\text{mac}}(\mathcal{A}) := \Pr[\mathcal{A}^{\text{MAC.T}_K, \text{MAC.V}_K} \text{ forges}]$ , which means  $\mathcal{A}$  receives  $\top$  from  $\text{MAC.V}_K$  by querying  $(N', M', T')$  while  $(N', M')$  has never been queried to  $\text{MAC.T}_K$ . Here,  $\mathcal{A}$  is assumed to be nonce-respecting, that is, the nonces in the tagging queries are distinct. The nonces in the verification queries have no restriction, and  $\mathcal{A}$  can repeat or reuse a nonce that was used by a tagging query.

#### D. Authenticated Encryption

Authenticated encryption (AE) [19] is used to ensure the privacy and authenticity of input data simultaneously. A nonce-based AE scheme  $\text{AE}$  consists of two functions: the encryption function  $\text{AE.E} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \rightarrow \mathcal{M} \times \mathcal{T}$  and the decryption function  $\text{AE.D} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \times \mathcal{T} \rightarrow \mathcal{M} \cup \{\perp\}$ . A ciphertext and a tag for  $(K, N, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{M}$  are derived as  $(C, T) = \text{AE.E}_K(N, M)$ . The tuple  $(N, C, T)$  is considered to be authentic if  $\text{AE.D}_K(N, C, T)$  returns the message  $M \neq \perp$ , and otherwise it is rejected. The security of AE is evaluated by two criteria: privacy and authenticity. The privacy advantage is the probability that the adversary successfully distinguishes  $\text{AE.E}_K$  from the random oracle  $\$(*, *)$ . For any query  $(N, M)$ , if  $(C, T) = \text{AE.E}_K(N, M)$ ,  $\$(N, M)$  returns random bits of length  $|C| + |T|$ . Thus,  $\text{Adv}_{\text{AE}}^{\text{priv}}(\mathcal{A}) := |\Pr[\mathcal{A}^{\text{AE.E}_K} \rightarrow 1] - \Pr[\mathcal{A}^{\$} \rightarrow 1]|$ . The authenticity advantage is the probability that the adversary creates a successful forgery by accessing  $\text{AE.E}_K$  and  $\text{AE.D}_K$ . It is defined as  $\text{Adv}_{\text{AE}}^{\text{auth}}(\mathcal{A}) := \Pr[\mathcal{A}^{\text{AE.E}_K, \text{AE.D}_K} \text{ forges}]$ , which means the probability that  $\mathcal{A}$  receives  $M' \neq \perp$  from  $\text{AE.D}_K$  by querying  $(N', C', T')$  while  $(N', M')$  has never been queried to  $\text{AE.E}_K$ . For both advantages, we assume the adversary is nonce-respecting in encryption queries. For authenticity, however, there is no restriction on nonce in the decryption queries, that is,  $\mathcal{A}$  may repeat a nonce or reuse a nonce that was used in an encryption query.

#### E. Tree-Based Memory Encryption Scheme

We assume two regions in storage memory: on-chip and off-chip areas. The former is assumed to be secure in which the adversary cannot eavesdrop or tamper the stored data. The latter can be attacked by an adversary who may perform eavesdropping (getting information of plaintext from ciphertext), tampering (modify the ciphertext without being detected), and replay (replacing the ciphertext with an old legitimate one). As mentioned in the introduction, tampering can be detected by simply applying a MAC to each data unit and storing the nonce and tag off-chip. If we use a nonce-based AE instead, it also prevents eavesdropping. However, these means are not sufficient to protect from replay attacks since the adversary can perform a replay on the (nonce, ciphertext, tag) tuple. Since the on-chip area is generally much more expensive than the

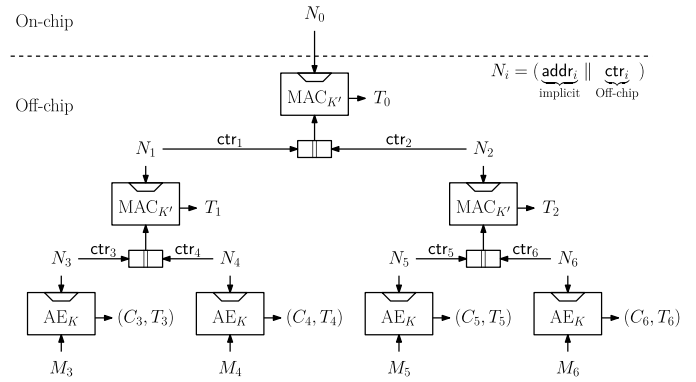


Fig. 1. An example of PAT2 with tree depth 2 and number of branches 2. A trapezoid in a box denotes the nonce input, and a box with  $\parallel$  denotes concatenation.

off-chip area, it is desirable to thwart all of these attacks with as small an on-chip area as possible. In addition, since the off-chip area is assumed to have large capacity, it is desirable to be able to perform tampering/replay detection and plaintext recovery with fewer memory accesses when only a part of the off-chip area is accessed.

As described earlier, Merkle hash tree [3] (with additional encryption layer for confidentiality) is a classical solution. It divides memory data into small chunks and associates their hashes with each leaf node of the tree. Every non-leaf node is associated with a hash of the hash values of its child nodes. By storing hash value of a root node on-chip and the others off-chip, the tampering and replay attacks can be detected. A similar scheme can be built by using MACs instead of hash functions by storing the key on-chip. Among such schemes, Parallelizable Authentication Tree (PAT) proposed by Hall and Jutla [4] (hereafter referred to as HJ05) is quite efficient for its parallelizability of both verify and update operations. It assigns a nonce to each node and stores the nonce associated with a root node in the on-chip area. PAT utilizes a MAC to compute a tag by taking the nonce assigned to its own node and nonces in its corresponding child nodes.<sup>3</sup>

In this paper, we use the term *ME tree* to refer to the tree-based ME scheme that also encrypts the leaf nodes. We introduce a generic construction of an ME tree, PAT2 (Fig. 1). It is mostly identical to PAT, but achieves confidentiality of memory by applying an AE scheme to the leaf nodes,<sup>4</sup> and it splits any nonce of PAT associated with a node into two values: an address and a local counter. The former is the memory address of the node, and the latter is a counter exclusively assigned to the node.

Fig. 1 shows an example of PAT2. Each nonce  $N_i$  assigned to node  $i$  consists of the address  $\text{addr}_i$  and the local counter  $\text{ctr}_i$ , which is initialized to 0 for all nodes. Memory data is split into four units ( $M_3$  to  $M_6$ ). After initialization, the tree keeps  $\text{ctr}_i$ ,  $T_i$  for  $i = 1, \dots, 6$ , and  $C_j$  for  $j = 3, \dots, 6$  at the off-chip area, and  $\text{ctr}_0$  at the on-chip area. When verifying

<sup>3</sup>To be more precise, [4] proposed to use a general deterministic MAC with input being prepended by a nonce, which is a typical way to convert a deterministic MAC into a nonce-based one.

<sup>4</sup>In fact, An ePrint version of HJ05 [20] specifies a combination of MAC and AE schemes for confidentiality of leaf data.



a piece of data, say  $M_3$ , we check if  $\text{AE}.\mathcal{D}_K(N_3, C_3, T_3)$  is authentic and  $\text{MAC}.\mathcal{V}_{K'}(N_i, \text{ctr}_{2i+1} \parallel \text{ctr}_{2i+2}, T_i) = \top$  for  $i = 0, 1$ . If all hold,  $M_3$  is considered to be authentic. When updating  $M_3$ , we first perform the above verification procedure, increment the corresponding local counters ( $\text{ctr}_0$ ,  $\text{ctr}_1$  and  $\text{ctr}_3$ ), and then renew  $(C_3, T_3)$ ,  $T_1$ , and  $T_0$ . Note that the steps in the verification and update procedures are independent and thus parallelizable. This is a crucial advantage of PAT/PAT2 over the classical hash tree, which only allows parallel verification. In addition, since an address is given from the outer legitimate system anyway, it does not need to be explicitly stored and it is always authentic. Therefore PAT2 enables the MAC input and off-chip overhead from the original PAT to be reduced. To the best of our knowledge, this technique was first proposed by [5]. In fact, by specifying the tree structures and the MAC and AE schemes, the resulting scheme is mostly identical to SIT. Therefore, PAT2 can be viewed as an abstraction of SIT. We consider PAT2 as our baseline scheme for its simple structure and efficiency, and present our scheme based on it in Section IV.

A number of ME trees that better handle the various criteria (except for latency) have been proposed. TEC-tree [6] provides confidentiality by encrypting data stored in all nodes. MAES [21] provides security against differential power analysis attacks. VAULT [11] and Morphable Counter [12] reduce the overhead of off-chip memory and are suitable for protecting large memory (*e.g.*, GBytes); however there is a tradeoff with the average latency because their counters are compressed.

### III. COMPONENTS OF ELM

To achieve low-latency operation, we designed dedicated MAC (PXOR-MAC) and AE (Flat-OCB).

#### A. PXOR-MAC: Incremental MAC

Fig. 2 shows the tagging function of PXOR-MAC using an  $n$ -bit block cipher  $E_K$ . The second key  $K' \xleftarrow{\$} \{0, 1\}^n$  is independent of  $K$ . The algorithm of  $\text{PXOR-MAC}.\mathcal{T}$  is obtained by replacing a TBC  $\tilde{E}$  of Alg. 1 with the following block cipher-based TBC.

$$\tilde{E}_{K,K'}^{0^n, i, j}(M) = E_K(M \oplus K' \cdot i \oplus j \cdot E_K(0^n)).$$

The verification function is trivially defined and thus omitted here. We let  $|N| = n$  and  $|T| = \tau$ . For simplicity, we exclude the case of partial blocks.

1) *Properties*: Since  $L = E_K(0^n)$  can be computed in advance, the latency of tag computation is essentially a sum of the latencies of  $n$ -bit multiplication ( $K' \cdot i$  for the block index  $i$ ) and  $E_K$ . The former can be large if  $i$  has large variations; however,  $m$  is usually not too large in practice, even when the total memory size is huge. Typically,  $m$  is upper-bounded by the number of branches, *e.g.*,  $2^7$  according to [12]. The hardware implementation is much more efficient than a full multiplier (using Gray code; see Section V). Consequently, the latency of mask computation becomes negligible, and PXOR-MAC has the optimal latency of one  $E_K$  call for tagging and verification functions thanks to the full parallelizability of the block cipher calls.




Fig. 2. PXOR-MAC.

PXOR-MAC is an Inc-MAC [13], which enables efficient tag computation when a small number of message blocks are changed. When a message block is changed together with a new nonce, the new tag is obtained by encrypting the corresponding blocks (*i.e.*, XOR of the message block and its mask) for both old and new ones, and taking an XOR of them and the old tag. We can further improve the incremental property of PXOR-MAC when used in PAT2. The MAC update function in the update procedure of PAT2 also invokes the MAC verification function, yielding some redundant  $E_K$  calls; thus, we can eliminate them. The resulting combined (verification and update) procedure, denoted by  $\text{PXOR-MAC}.\mathcal{VU}$ , is shown in Alg. 2. It takes old nonce  $N_o$ , old plaintext  $M_o$ , old tag  $T_o$ , new nonce  $N_n$ , and new plaintext  $M_n$ . It outputs new tag  $T_n$  such that  $T_n = \text{PXOR-MAC}.\mathcal{T}_{K,K'}(N_n, M_n)$  holds if  $T_o$  is authentic; otherwise, it outputs  $\perp$ . For simplicity, Alg. 2 assumes that  $M_o$  and  $M_n$  have  $m$  blocks. For example, when one message block and nonce are changed, (which is the case of ELM),  $\text{PXOR-MAC}.\mathcal{VU}$  needs only  $m + 3$   $E_K$  calls except for mask derivation, while PXOR-MAC without incremental update needs  $2m$ . Even if we use the incremental update feature of PXOR-MAC after verification, it still needs  $m + 5$  calls. This difference is not negligible as update occurs at all nodes on the path.

---

#### Algorithm 1 $\text{PXOR-MAC-T}.\mathcal{T}_{\tilde{E}}(N, M)$

---

- 1:  $M[1] \parallel \dots \parallel M[m] \xleftarrow{\$} M, T \leftarrow 0^\tau$
  - 2: **for**  $1 \leq i \leq m$  **do**
  - 3:  $T \leftarrow T \oplus \text{msb}_\tau(\tilde{E}_{K,K'}^{0^n, i, 0}(M[i]))$
  - 4:  $T \leftarrow T \oplus \text{msb}_\tau(\tilde{E}_{K,K'}^{0^n, m, 1}(N))$
  - 5: **return**  $T$
- 

---

#### Algorithm 2 $\text{PXOR-MAC}.\mathcal{VU}_{E_K}(N_o, M_o, T_o, N_n, M_n)$

---

- 1:  $L \leftarrow E_K(0^n), T' \leftarrow 0^\tau, T_n \leftarrow T_o$
  - 2:  $M_o[1] \parallel \dots \parallel M_o[m] \xleftarrow{\$} M_o, M_n[1] \parallel \dots \parallel M_n[m] \xleftarrow{\$} M_n$
  - 3: **for**  $1 \leq i \leq m$  **do**
  - 4:  $S \leftarrow \text{msb}_\tau(E_K(M_o[i] \oplus K' \cdot i)), T' \leftarrow T' \oplus S$
  - 5: **if**  $M_o[i] \neq M_n[i]$  **then**
  - 6:  $T_n \leftarrow T_n \oplus S \oplus \text{msb}_\tau(E_K(M_n[i] \oplus K' \cdot i))$
  - 7:  $S \leftarrow \text{msb}_\tau(E_K(N_o \oplus K' \cdot m \oplus L)), T' \leftarrow T' \oplus S$
  - 8:  $T_n \leftarrow T_n \oplus S \oplus \text{msb}_\tau(E_K(N_n \oplus K' \cdot m \oplus L))$
  - 9: **if**  $T' \neq T_o$  **then**
  - 10: **return**  $\perp$
  - 11: **return**  $T_n$
- 

2) *Security*: We assume the underlying block cipher of PXOR-MAC is an  $n$ -bit URP  $P$ , which is denoted by  $\text{PXOR-MAC}_P$ .

*Theorem 1:* Let  $\sigma_{\text{mac}}$  be the number of accesses to  $\mathbf{P}$  in the MAC game s.t.  $\sigma_{\text{mac}} \leq 2^{n-1}$ , and  $q_v$  be the number of queries to the verification oracle. We obtain

$$\text{Adv}_{\text{PXOR-MAC}_P}^{\text{mac}}(\mathcal{A}) \leq \frac{2q_v}{2^\tau} + \frac{4.5\sigma_{\text{mac}}^2}{2^n}.$$

Note that this advantage is an information-theoretic bound. The computational counterpart is derived from our bound. Since this is fairly standard [22], we omit it here. This theorem means that PXOR-MAC has the security bound of  $\min\{O(2^\tau), O(2^{n/2})\}$ , which is almost the same as that of MAC used in SIT. PXOR-MAC only fulfills “basic” Inc-MAC security [13], which requires that the update function of Inc-MAC always takes benign inputs (in our case this corresponds to require  $T_o = \text{PXOR-MAC}.\mathcal{T}_{K,K'}(N_o, M_o)$  holds whenever  $\mathcal{VU}$  outputs  $T_n \neq \perp$ ). We stress that, while a stronger, “tamper-proof” security notion was also defined by [13], basic security is enough for our case since the adversary cannot intervene in the update procedure.

3) *Proof of Theorem 1:* Let  $\tilde{\mathbf{P}}$  be an  $n$ -bit TURP having the same tweak space as  $\tilde{E}$ , and  $\tilde{E}$  be a TBC involving  $\mathbf{P}$  and an independent key  $K'$ , defined as  $\tilde{E}^{0^n, i, j}(M) = \mathbf{P}(M \oplus K' \cdot i \oplus j \cdot \mathbf{P}(0^n))$ . Then we obtain

$$\text{Adv}_{\text{PXOR-MAC}_P}^{\text{mac}}(\mathcal{A}) \leq \text{Adv}_{\text{PXOR-MAC}_{\tilde{P}}}^{\text{mac}}(\mathcal{A}) + \text{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{B}),$$

where  $\mathcal{B}$  is the adversary against  $\tilde{E}$ . Regarding the first term, we can prove  $\text{Adv}_{\text{PXOR-MAC}_{\tilde{P}}}^{\text{mac}}(\mathcal{A}) \leq 2q_v/2^\tau$  in the same manner as PMAC [18], and thus we omit. We then evaluate the second term. We define the offset function  $F$  as  $F_{K'}((i, j), \mathbf{P}(0^n)) = K' \cdot i \oplus j \cdot \mathbf{P}(0^n)$ , where  $i \in \{1, 2, \dots\}$ ,  $j \in \{0, 1\}$ . Then,  $\tilde{E}^{0^n, i, j}(M) = \mathbf{P}(M \oplus F_{K'}((i, j), \mathbf{P}(0^n)))$  holds for any  $(i, j, M)$ . We introduce the definition and the lemma for offset functions. They are simplified ones of Definition 4.1 and Theorem 4.1 in [23].

*Definition 1:* Let  $V \xleftarrow{\$} \{0, 1\}^n$ . An offset function  $F$  is said to be  $(\varepsilon, \gamma, \rho)$ -uniform if  $F$  satisfies the following conditions.

$$\begin{aligned} \max_{l \neq l', \delta \in \{0, 1\}^n} \Pr[F(l, V) \oplus F(l', V) = \delta] &\leq \varepsilon, \\ \max_{l, \delta \in \{0, 1\}^n} \Pr[F(l, V) = \delta] &\leq \gamma, \\ \max_{l, \delta \in \{0, 1\}^n} \Pr[F(l, V) \oplus V = \delta] &\leq \rho. \end{aligned}$$

*Lemma 1:* Suppose that  $\tilde{E}$  uses an  $(\varepsilon, \gamma, \rho)$ -uniform offset function  $F$ . We obtain the following evaluation.

$$\text{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{B}) \leq q^2 \left( 2\varepsilon + \gamma + \rho + \frac{1}{2^{n+1}} \right),$$

where  $q$  is the number of encryption queries s.t.  $q \leq 2^{n-1}$ . Since  $K'$  and  $\mathbf{P}(0^n)$  are uniformly random and independent,  $(\varepsilon, \gamma, \rho) = (1/2^n, 1/2^n, 1/2^n)$  trivially holds. Thus, we obtain  $\text{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{B}) \leq 4.5\sigma_{\text{mac}}^2/2^n$ . This concludes the proof.

## B. Flat-OCB: Low-Latency AE

To encrypt the leaf data of PAT2, an AE scheme is needed. We take OCB as the baseline AE for its efficiency: it needs  $m$  plus a few block cipher ( $E_K$ ) calls to process  $m$  blocks




Fig. 3. Flat-OCB.

(while a generic composition of Encryption and MAC needs at least  $2m$  calls). Because these  $m$  calls are parallelizable, OCB has quite a low latency. However, when looking into the dependency among “a few block cipher calls” mentioned above, we find a problem in the decryption. In more detail, the final  $E_K$  call in the OCB decryption takes the sum of the plaintext blocks after the plaintext blocks are obtained by ECB-like decryption (see *e.g.*, [16, Fig. 5]). This results in one call that cannot be computed in parallel, and adds a significant latency compared to the encryption. This is not desirable for applications requiring low latency. We present a solution to this problem. Concretely, we first propose an improved version of the TBC-based interpretation of OCB ( $\Theta\text{CB3}$  in [16], hereafter referred to as  $\Theta\text{CB}$ ), which we call Flat-OCB and show its block cipher-based instantiation, Flat-OCB. It has a smaller decryption latency than OCB, while keeping the same encryption latency.

### Algorithm 3 Flat-OCB. $\mathcal{E}_{\tilde{E}}(N, M)$

- 1:  $M[1] \parallel \dots \parallel M[m] \xleftarrow{\$} M, T \leftarrow \text{msb}_\tau(\tilde{E}_K^{N,0,0}(0^n))$
- 2: **for**  $1 \leq i \leq m-1$  **do**
- 3:  $C[i] \leftarrow \tilde{E}_K^{N,i,0}(M[i]), T \leftarrow T \oplus \text{msb}_\tau(M[i])$
- 4:  $C[m] \leftarrow \tilde{E}_K^{N,m-1,1}(M[m]), T \leftarrow T \oplus \text{msb}_\tau(M[m])$
- 5:  $C \leftarrow C[1] \parallel \dots \parallel C[m]$
- 6: **return**  $C, T$

### Algorithm 4 MASK $_{K_1, K_2, K_3, K_4}(N)$

- 1:  $N_1 \leftarrow \text{msb}_{n/2}(N), N_2 \leftarrow \text{lsb}_{n/2}(N)$
- 2: **return**  $\Delta \leftarrow (N_1 \cdot K_1 \parallel N_2 \cdot K_2) \oplus (N_2 \cdot K_3 \parallel N_1 \cdot K_4)$

Alg. 3 and Fig. 3 show the encryption of Flat-OCB. Decryption is naturally obtained, thus omitted. It is based on  $n$ -bit TBC,  $\tilde{E}$ , using  $n$ -bit nonce. We exclude the case of partial blocks for simplicity. Flat-OCB is similar to  $\Theta\text{CB}$ , however, the crucial difference is how the tag  $T$  is generated. While  $\Theta\text{CB}$  encrypts the checksum  $M[1] \oplus M[2] \oplus \dots \oplus M[m]$  using  $\tilde{E}$  to produce  $T$ , ours first encrypts  $N$  and then takes a sum with the checksum. To build a block cipher-based AE, we instantiate  $\tilde{E}$  with an  $n$ -bit block cipher  $E_K$  as

$$\tilde{E}_K^{N, i, j}(M) = E_K(M \oplus \Delta \oplus 2^i 3^j L) \oplus \Delta \oplus 2^i 3^j L, \quad (1)$$

where  $i \in \{0, 1, 2, \dots\}$ ,  $j \in \{0, 1\}$ ,  $L = E_K(0^n)$  and  $\Delta$  is derived from  $N$ . The value  $\Delta$  needs to be pairwise independent for a pair of distinct inputs, and this may be realized by various means. To keep the compatibility with SIT, we generate  $\Delta$  using  $n/2$ -bit multiplications as Alg. 4 for some even  $n$

TABLE I

COMPARISON OF AE MODES.  $m$  IS THE NUMBER OF INPUT BLOCKS, AND “MUL” IN THE LATENCY COLUMNS DENOTES  $n/2$ -bit FULL GF MULTIPLICATION. THE LAST COLUMN DENOTES THE TOTAL ON-CHIP SIZE. FOR SIMPLICITY, THE ENCRYPTION AND DECRYPTION LATENCY OF (T)BC ARE ASSUMED TO BE IDENTICAL, AND (T)BC HAS  $n$ -bit BLOCK SIZE AND  $n$ -bit KEY

Scheme	Enc latency	Dec latency	Circuit size to achieve the best latency	The total number of primitive calls	Total size of key and preprocessed data (bits)
OCB [16]	1 TBC	2 TBCs	$m + 1$ TBCs	$m + 1$ TBCs	$n$
Flat-OCB (This work)	1 TBC	1 TBC	$m + 1$ TBC	$m + 1$ TBC	$n$
OCB [16]	2 BCs	3 BC	$m + 1$ BC	$m + 2$ BC	$n$
SIT-AE [7]	1 BC + 1 Mul.	$\max\{1 \text{ BC}, 1 \text{ Mul.}\}$	$m + 1$ BC + $2m$ Mul	$m + 1$ BC + $2m$ Mul	$2n + mn$
Flat-OCB (This work)	1 BC + 1 Mul	1 BC + 1 Mul	$m + 1$ BC + 4 Mul	$m + 1$ BC + 4 Mul	$4n$

and random and independent  $n/2$ -bit keys  $(K_1, K_2, K_3, K_4)$ . We define TBC as the block cipher-based TBC defined in (1) with MASK (Alg. 4). We obtain Flat-OCB by instantiating  $\tilde{E}$  of Flat-OCB with TBC.

1) *Properties*: Table I compares Flat-OCB and Flat-OCB with other AEs. SIT-AE is a GCM-like AE inside SIT [8]. Notably, as a mode of TBC, the latency of Flat-OCB is the lowest achievable – one TBC call – for both encryption and decryption, which is achieved for the first time. In contrast, the decryption latency of OCB is two for the aforementioned reason. For other properties, Flat-OCB has the same figures as those of OCB. Similarly, Flat-OCB has the same encryption latency as OCB, and has the smaller decryption latency than OCB. SIT-AE has even smaller decryption latency, however, it requires  $2m$  multiplier circuits to achieve this, whereas Flat-OCB requires only 4 multipliers.<sup>5</sup> Moreover, the key size of SIT-AE is linear to  $m$ , which will have a non-negligible impact. In contrast, Flat-OCB requires  $3n$ -bit key, plus the  $n$ -bit preprocessed data  $L = E_K(0^n)$  used inside TBC.

2) *Security*: We assume that the underlying block cipher is an  $n$ -bit URP  $P$ , and only present an information-theoretic bound based on  $P$ .

*Theorem 2*: The advantages of Flat-OCB and Flat-OCB are obtained as follows:

$$\begin{aligned} \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{priv}}(\mathcal{A}) &= 0, \quad \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) \leq \frac{2q_d}{2^\tau}, \\ \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{priv}}(\mathcal{A}) &\leq \frac{4.5\sigma_{\text{priv}}^2}{2^n}, \\ \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) &\leq \frac{2q_d}{2^\tau} + \frac{4.5\sigma_{\text{auth}}^2}{2^n}, \end{aligned}$$

where  $\sigma_{\text{priv}}$ ,  $\sigma_{\text{auth}}$ , and  $q_d$  are the parameters for  $\mathcal{A}$  and  $\mathcal{A}^\pm$ . The parameter  $\sigma_{\text{priv}}$  (resp.  $\sigma_{\text{auth}}$ ) is the number of accesses to  $P$  in the privacy (resp. authenticity) game s.t.  $\sigma_{\text{priv}}, \sigma_{\text{auth}} \leq 2^{n-1}$ . The parameter  $q_d$  is the number of queries to the decryption oracle.

Flat-OCB has the same advantages as those of OCB ( $\mathbf{Adv}_{\text{OCB}_P}^{\text{priv}}(\mathcal{A}) = 0$ ,  $\mathbf{Adv}_{\text{OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) \leq (2^{n-\tau}q_d)/(2^n - 1)$ ), hence there is no security penalty, up to the constant. The security of Flat-OCB is comparable to those of OCB and SIT-AE. Assuming  $n = 128$  and  $\tau = 64$ , Flat-OCB has 64-bit data security that is comparable to OCB and SIT-AE.

<sup>5</sup>The proper number of multipliers for hardware implementation depends on the system constraint/architecture. See also Section V.

3) *Proof of Theorem 2*: Let  $\tilde{P}$  be an  $n$ -bit TURP having the same tweak space as TBC. Let  $\tilde{E}$  be the TBC mode obtained by replacing  $E_K$  in TBC with  $P$ . Then we obtain the following inequations.

$$\begin{aligned} \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{priv}}(\mathcal{A}) &\leq \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{priv}}(\mathcal{A}) + \mathbf{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{B}), \\ \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) &\leq \mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) + \mathbf{Adv}_{\tilde{E}}^{\text{tsprp}}(\mathcal{B}^\pm), \end{aligned}$$

where  $\mathcal{B}$  and  $\mathcal{B}^\pm$  are adversaries against  $\tilde{E}$ . To derive the advantages of  $\mathcal{B}$  and  $\mathcal{B}^\pm$ , we use the tsprp version of the methodology in [23], which is almost the same as the tprp one (*i.e.*, Def. 1 and Lem. 1), thus we omit the details. We define the offset function  $F$  of  $\tilde{E}$  as  $F((N, i, j), P(0^n)) = \text{MASK}(N) \oplus 2^i \cdot 3^j P(0^n)$ . Since keys of MASK and  $P(0^n)$  are uniformly random and independent, and MASK is pairwise independent,  $(\varepsilon, \gamma, \rho) = (1/2^n, 1/2^n, 1/2^n)$  trivially holds. Thus,  $\mathbf{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{B}) \leq 4.5\sigma_{\text{priv}}^2/2^n$ , and  $\mathbf{Adv}_{\tilde{E}}^{\text{tsprp}}(\mathcal{B}^\pm) \leq 4.5\sigma_{\text{auth}}^2/2^n$  hold.

We then evaluate the security bounds of Flat-OCB. For the privacy,  $\mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{priv}}(\mathcal{A}) = 0$  holds since every TURP call in the game takes different tweaks. For the authenticity, We start with the case  $q_d = 1$ . Let  $\{(N_1, M_1, C_1, T_1), \dots, (N_{q_e}, M_{q_e}, C_{q_e}, T_{q_e})\}$  be the transcript in encryption queries, and  $(N', C', T')$  be the decryption query. If  $\forall i \in [q_e], N' \neq N_i$ , the probability  $\mathcal{A}^\pm$  forges is at most  $1/2^\tau$ . We then evaluate the case  $\exists \alpha \in [q_e], N' = N_\alpha$ . If  $|C_\alpha|_n \neq |C'|_n := m'$ , the TURP which decrypts  $C'[m']$  takes a different tweak from all tweaks invoked in the encryption queries. Thus, the probability  $\mathcal{A}^\pm$  forges is at most  $1/2^\tau$ . If  $|C_\alpha|_n = |C'|_n = m'$ , all tweaks invoked in the decryption query are the same as those in the  $\alpha$ -th encryption query. However,  $C'[i] \neq C_\alpha[i]$  holds for  $\exists i \in [m']$ , and thus, the probability  $\mathcal{A}^\pm$  successfully guesses  $\text{msb}_\tau(M^*[i])$ , where  $M^*[i]$  is the decrypted value of  $C'[i]$ , is at most  $2/2^\tau$ . From the above, we obtain the advantage for the case  $q_d = 1$ :  $\mathbf{Adv}_{\text{Flat-OCB}_P}^{\text{auth}}(\mathcal{A}^\pm) \leq 2/2^\tau$ . We apply the standard conversion from single to multiple decryption queries [24] and obtain the bound  $q_d(2/2^\tau)$  for  $q_d \geq 1$ . This concludes the proof of Flat-OCB.

#### IV. ELM

ELM is based on PAT2, where the internal MAC and AE are instantiated by PXOR-MAC and Flat-OCB.

##### A. Notations for the Tree

Let  $b \geq 2$  be the number of branches, and let  $d$  be the tree depth, where the root has depth 0 and a leaf node has



depth  $d$ . ELM uses a balanced tree, so we have  $b^d$  leaves. The entire memory (plaintext) is divided into  $\ell$ -bit chunks and  $i$ -th leaf node is denoted by  $\text{leaf}(i)$  for  $i \in [b^d]$ . The whole input memory is  $M = M[1] \parallel \dots \parallel M[b^d]$  s.t.  $|M[i]| = \ell$  for  $i \in [b^d]$  and  $M[i]$  is associated with  $\text{leaf}(i)$ . The ciphertext chunk corresponding to  $M[i]$  is denoted by  $C[i]$ , which is stored in  $\text{leaf}(i)$ . For node  $u$ , the memory address, the counter, and the tag are denoted by  $\text{add}(u)$ ,  $\text{ctr}(u)$ , and  $\text{tag}(u)$ , respectively, where  $|\text{add}(u)| = \alpha$ ,  $|\text{ctr}(u)| = \beta$ , and  $|\text{tag}(u)| = \tau$ . The entire data to construct the tree stored in the on-chip and off-chip areas is denoted by  $\sigma$ , which includes  $C[i]$  for  $i \in [b^d]$ ,  $\text{ctr}(u)$ , and  $\text{tag}(u)$  for all nodes  $u$ . Note that we exclude the key and the preprocessed data from  $\sigma$ . As we adopt PAT2, we exclude the node addresses from  $\sigma$  and assume they are given by the system when needed. For the root node  $u_r$ , we store  $\text{ctr}(u_r)$  on-chip. The leftover data of  $\sigma$  is stored off-chip. We may also use  $\sigma$  to refer to the tree construction itself. We also write a node  $u$ , leaf node, plaintext chunk, and ciphertext chunk of  $\sigma$  as  $u^\sigma$ ,  $\text{leaf}(i)^\sigma$ ,  $M^\sigma[i]$ , and  $C^\sigma[i]$ , respectively. If no confusion is possible, we omit their superscript  $\sigma$ . For any non-leaf node  $u^\sigma$  and  $i \in [b]$ ,  $\text{ch}_i(u^\sigma)$  denotes its  $i$ -th child node.

---

**Algorithm 5** IT: Initialization of a Tree Construction  $\sigma$ 


---

**Input**  $M = M[1] \parallel \dots \parallel M[b^d]$  s.t.  $|M[i]| = \ell$  for  $i \in [b^d]$

**Output**  $\sigma$

```

1:  $\sigma \leftarrow 0^{\binom{b^d-1}{b-1} \times (\beta+\tau) + b^d \times (\ell+\beta+\tau)}$ 
2: for all nodes  $u$  do
3:    $\text{ctr}(u^\sigma) \leftarrow 0^{\beta-1} 1$ 
4:   for  $1 \leq i \leq b^d$  do
5:      $(C[i], \text{tag}(\text{leaf}(i)^\sigma))$ 
        $\leftarrow \text{Flat-OCB}.\mathcal{E}(\text{add}(\text{leaf}(i)^\sigma) \parallel \text{ctr}(\text{leaf}(i)^\sigma), M[i])$ 
6:    $M' \leftarrow \text{ctr}(\text{ch}_1(u^\sigma)) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u^\sigma))$ 
7:    $\text{tag}(u^\sigma) \leftarrow \text{PXOR-MAC}.\mathcal{T}(\text{add}(u^\sigma) \parallel \text{ctr}(u^\sigma), M')$ 
8:   for all intermediate nodes  $u$  do
9:      $\text{tag}(u^\sigma) \leftarrow \text{PXOR-MAC}.\mathcal{T}(\text{add}(u^\sigma) \parallel \text{ctr}(u^\sigma), M')$ 
10:  return  $\sigma$ 

```

---



---

**Algorithm 6** VT: Checking the Validity of  $\text{leaf}(idx)$ .

---

**Input**  $idx, \sigma$

**Output**  $\top$  or  $\perp$

```

1:  $(u_0^\sigma, \dots, u_d^\sigma) \leftarrow$  path of nodes from root to specified leaf
   (i.e.,  $u_0^\sigma$  is the root node  $u_r^\sigma$ , and  $u_d^\sigma$  is equal to  $\text{leaf}(idx)$ .)
2: for  $0 \leq i \leq d-1$  do
3:   if  $\text{PXOR-MAC}.\mathcal{V}(\text{add}(u_i^\sigma) \parallel \text{ctr}(u_i^\sigma),$ 
      $\text{ctr}(\text{ch}_1(u_i^\sigma)) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_i^\sigma)), \text{tag}(u_i^\sigma)) = \perp$  then
4:     return  $\perp$ 
5: if  $\text{Flat-OCB}.\mathcal{D}(\text{add}(u_d^\sigma) \parallel \text{ctr}(u_d^\sigma), C[idx], \text{tag}(u_d^\sigma)) = \perp$  then
6:   return  $\perp$ 
7: return  $\top$ 

```

---

### B. Specifications of ELM

ELM consists of three algorithms: IT, VT, and UT defined in Algs. 5, 6, and 7. ELM's key consists of the AE and MAC keys, which are chosen uniformly at random and independent. IT initializes the tree. It takes an input plaintext  $M$ , and outputs a tree  $\sigma$ . Here,  $\sigma$  consists of the local counters

---

**Algorithm 7** UT: Updating the Message of  $\text{leaf}(idx)$  to  $B$ .

---

**Input**  $idx, B, \sigma$

**Output**  $\tilde{\sigma}$  or  $\perp$

```

1:  $\tilde{\sigma} \leftarrow \sigma, (u_0, \dots, u_d) \leftarrow$  path of nodes from root to specified leaf
2: for  $0 \leq i \leq d$  do
3:    $\text{ctr}(u_i^{\tilde{\sigma}}) \leftarrow \text{ctr}(u_i^\sigma) + 1$ 
4:   for  $0 \leq i \leq d-1$  do
5:      $N_o \leftarrow \text{add}(u_i^\sigma) \parallel \text{ctr}(u_i^\sigma), N_n \leftarrow \text{add}(u_i^{\tilde{\sigma}}) \parallel \text{ctr}(u_i^{\tilde{\sigma}})$ 
6:      $M_o \leftarrow \text{ctr}(\text{ch}_1(u_i^\sigma)) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_i^\sigma))$ 
7:      $M_n \leftarrow \text{ctr}(\text{ch}_1(u_i^{\tilde{\sigma}})) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_i^{\tilde{\sigma}}))$ 
8:      $\text{tag}(u_i^{\tilde{\sigma}}) \leftarrow \text{PXOR-MAC}.\mathcal{V}\mathcal{U}(N_o, M_o, \text{tag}(u_i^\sigma), N_n, M_n)$ 
9:     if  $\text{tag}(u_i^{\tilde{\sigma}}) = \perp$  then
10:      return  $\perp$ 
11:   if  $\text{Flat-OCB}.\mathcal{D}(\text{add}(u_d^\sigma) \parallel \text{ctr}(u_d^\sigma), C^\sigma[idx], \text{tag}(u_d^\sigma)) = \perp$ 
     then
12:     return  $\perp$ 
13:    $(C^{\tilde{\sigma}}[idx], \text{tag}(u_d^{\tilde{\sigma}})) \leftarrow \text{Flat-OCB}.\mathcal{E}(\text{add}(u_d^\sigma) \parallel \text{ctr}(u_d^\sigma), B)$ 
14:  return  $\tilde{\sigma}$ 

```

---

being initialized to zero, the tags for the non-leaf nodes, and the (ciphertext, tag) pairs for the leaf nodes. VT checks the validity of a specified leaf node. It is associated with a read operation. VT takes an index of a leaf node  $idx \in [b^d]$  and a tree  $\sigma$  as input. It outputs  $\top$  if all the verifications of PXOR-MAC and the decryption of Flat-OCB are successful, and otherwise  $\perp$ . UT renews plaintext chunk assigned to the specified leaf node, which is associated with a write operation. UT takes an index of leaf node  $idx$ , an update value (new plaintext)  $B$  s.t.  $|B| = \ell$ , and a tree  $\sigma$  as input. It returns a renewed tree  $\tilde{\sigma}$  if the verification is successful, otherwise  $\perp$ . Note that it is essential for UT to check the validity of the data associated in node path in order to prevent a replay attack. If the verification in UT is bypassed, the adversary can roll back the value of  $\text{ctr}(\cdot)$  and mount a replay attack. For the verification and update of intermediate nodes in UT, we use  $\text{PXOR-MAC}.\mathcal{V}\mathcal{U}$  to eliminate redundant computations. Similarly, Flat-OCB in lines 7–7 of Alg. 7 can eliminate some redundant field multiplications in deriving  $\Delta$  by caching.

### C. Efficiency of ELM

ELM is designed to achieve low latency by utilizing the incremental property of MAC and full parallelizability of the cryptographic components and the tree structure. In particular, the incremental property greatly reduces the latency of UT. An Inc-MAC with basic security works (See Sec. III-A.1). Suppose  $\alpha = \beta = n/2$  and some even  $b$ . One VT call needs  $(1+2/b)d E_K$  calls for intermediate and root nodes. One UT call needs  $(3+2/b)d E_K$  calls for intermediate and root nodes, while UT with a non-incremental MAC needs at least twice as many  $E_K$  calls as VT does. In addition, ELM is scalable for its constant on-chip size (the key and the preprocessed data, total  $7n$  bits). However, (a generalized version of) SIT needs an on-chip area linear to  $b$  and  $\beta$ .

### D. Security of ELM

To the best of our knowledge, the provable security of PAT2 have not been formally shown in the literature. Among

many ME proposals, the concrete provable security of the proposed scheme (as in the same manner to the seminal Bellare *et al.* [22]) is rarely shown, or even the formal, game-based security notion is often missing. It is true that HJ05 defines the security notion for PAT and proves the security, however, PAT has slightly different tree construction from PAT2, and more importantly, it lacks a confidentiality notion as the original PAT does not encrypt data. We remedy this situation by defining two appropriate security notions for ME trees: *privacy* and *unforgeability*. The former is for the confidentiality of plaintext and is defined analogously to the standard “privacy” notion for AE. The latter is mostly identical to the notion introduced by HJ05. We show the concrete security bounds of PAT2 under these notions. Although our analysis is not surprising, we think a formal security framework is important and will facilitate further studies on ME schemes. In particular, the security model is identical to what SIT considers.

1) *Security Notion of ME Tree*: Suppose that *Tree* is an ME tree scheme defined as a tuple of three functions: the initialization function IT, the verification function VT, and the update function UT. We assume that UT contains VT and performs VT first. Recall that  $\text{IT}(M) = \sigma$ ,  $\text{VT}(idx, \sigma) = \top$  or  $\perp$ , and  $\text{UT}(idx, B, \sigma) = \tilde{\sigma}$  or  $\perp$ . Also recall that  $\sigma$  includes data stored in the on-chip area.<sup>6</sup>

For the privacy of *Tree*, we define IT-\$ and UT-\$. They return their ciphertexts and tags to be stored in the leaf nodes as random strings whose lengths are the same as those of IT and UT, respectively. Regarding other variables, for example, the data associated with the intermediate nodes, they return the same outputs as IT and UT. The privacy security of *Tree* is defined as the probability that an adversary  $\mathcal{A}$  successfully distinguishes (IT, UT) from (IT-\$, UT-\$). It is written as

$$\text{Adv}_{\text{Tree}}^{\text{pt}}(\mathcal{A}) := |\Pr[\mathcal{A}^{\text{IT,UT}} \rightarrow 1] - \Pr[\mathcal{A}^{\text{IT-}\$, \text{UT-}\$} \rightarrow 1]|,$$

where  $\mathcal{A}$  plays the following game.

- 1)  $\mathcal{A}$  queries  $M$  to the tree initialization oracle (IT or IT-\$) and obtains  $\tilde{\sigma}_0$ .
- 2)  $\mathcal{A}$  makes  $q$  adaptive queries to the update oracle (UT or UT-\$). Let  $\{(idx_1, B_1, \sigma_1, \tilde{\sigma}_1), \dots, (idx_q, B_q, \sigma_q, \tilde{\sigma}_q)\}$  be the transcript obtained by the update queries. Here, we assume that  $\sigma_i = \tilde{\sigma}_{i-1}$  for  $i \in [q]$  so that  $\mathcal{A}$  can always obtain an updated tree, not  $\perp$ .
- 3)  $\mathcal{A}$  guesses which oracle pair she has queried ((IT, UT) or (IT-\$, UT-\$)) and accordingly outputs a bit.

For the unforgeability of *Tree*, our definition follows [4]. It is defined as the advantage of an adversary  $\mathcal{A}'$  querying IT and UT successfully distinguishes VT from  $\perp_{\top}(\cdot, \cdot)$  which always returns  $\perp$  for any inputs. The unforgeability advantage of  $\mathcal{A}'$  is defined as

$$\text{Adv}_{\text{Tree}}^{\text{uft}}(\mathcal{A}') := |\Pr[\mathcal{A}'^{\text{IT,UT,VT}} \rightarrow 1] - \Pr[\mathcal{A}'^{\text{IT,UT,}\perp_{\top}} \rightarrow 1]|,$$

where  $\mathcal{A}'$  plays the following game.

<sup>6</sup>Recall that  $\sigma$  does not include the secret key and the preprocessed data. In this paper, we do not assume the confidentiality of  $\text{Sec}(\sigma)$ , thus the adversary can look into it. It is a weaker assumption than that assuming both the confidentiality and tamper freeness.

- 1)  $\mathcal{A}'$  queries  $M$  to IT and obtains  $\tilde{\sigma}_0$ .
- 2)  $\mathcal{A}'$  makes  $q'$  adaptive queries to UT. Let  $\{(idx_1, B_1, \sigma_1, \tilde{\sigma}_1), \dots, (idx_{q'}, B_{q'}, \sigma_{q'}, \tilde{\sigma}_{q'})\}$  be the transcript obtained by update queries. As well as the privacy game, we assume that  $\sigma_i = \tilde{\sigma}_{i-1}$  for  $i \in [q']$ .
- 3)  $\mathcal{A}'$  queries  $(idx', \sigma')$  to the verification oracle (VT or  $\perp_{\top}$ ) and obtains  $\top$  or  $\perp$ . Let  $(u_0, \dots, u_d)$  be the path of nodes from the root node to  $\text{leaf}(idx')$ . To exclude a trivial win, we assume that there exists  $i \in \{0, \dots, d\}$  such that  $u_i^{\sigma'}$  stores different data from that stored in  $u_i^{\tilde{\sigma}_{q'}}$ . Moreover,  $\text{Sec}(\sigma') = \text{Sec}(\tilde{\sigma}_{q'})$  also must hold since the data in the on-chip area cannot be tampered.
- 4)  $\mathcal{A}'$  guesses which oracle pair she has queried ((IT, UT, VT) or (IT, UT,  $\perp_{\top}$ )) and accordingly outputs a bit.

We stress that  $\mathcal{A}'$  can perform a verification query s.t.  $u_i^{\sigma'}$  stores the same data as that stored in  $u_i^{\tilde{\sigma}_j}$  for  $0 \leq i \leq d$  and  $0 \leq j \leq q' - 1$ , unless the data stored in  $u_i^{\tilde{\sigma}_j}$  is the same as that stored in  $u_i^{\tilde{\sigma}_{q'}}$  for all  $i \in \{0, \dots, d\}$  as described in the third operation of the above game. This condition is essential for the unforgeability notion to capture an adversary who performs a replay attack. We remark that the security notion  $\text{Adv}^{\text{uft}}$  also captures the adversary who tampers the data before update, i.e., a forgery attack against UT, since UT contains VT.

2) *Security Bounds of PAT2 and ELM*: The algorithms of PAT2 are obtained by replacing PXOR-MAC and Flat-OCB in Algs. 5, 6, and 7 to general MAC and AE denoted by MAC and AE. Note that  $\text{PXOR-MAC}.\mathcal{V}\mathcal{U}$  in Alg. 7 is interpreted as a combined function that outputs  $T_n = \text{MAC}.\mathcal{T}(N_n, M_n)$  if  $\text{MAC}.\mathcal{V}(N_o, M_o, T_o) = \top$ , otherwise outputs  $\perp$ .

*Theorem 3*: We obtain the following bounds:

$$\text{Adv}_{\text{PAT2}}^{\text{pt}}(\mathcal{A}) \leq \text{Adv}_{\text{AE}}^{\text{priv}}(\mathcal{A}_{\text{ae}}), \quad (2)$$

$$\text{Adv}_{\text{PAT2}}^{\text{uft}}(\mathcal{A}') \leq \text{Adv}_{\text{AE}}^{\text{auth}}(\mathcal{A}_{\text{ae}}^{\pm}) + \text{Adv}_{\text{MAC}}^{\text{mac}}(\mathcal{A}_{\text{mac}}). \quad (3)$$

The parameters of  $\mathcal{A}_{\text{ae}}$ ,  $\mathcal{A}_{\text{ae}}^{\pm}$ , and  $\mathcal{A}_{\text{mac}}$  can be determined by those of  $\mathcal{A}$  and  $\mathcal{A}'$ . Theorem 3 means that the security of PAT2 can be reduced to those of MAC and AE. This is not surprising, but we cannot find such a formal treatment (in particular for the combination of MAC and AE to guarantee privacy and unforgeability) in the literature. The security bounds of ELM can be derived by combining the advantages of PAT2 (2), (3) and those of Flat-OCB and PXOR-MAC. When assuming the underlying block cipher of ELM is P, we obtain

$$\text{Adv}_{\text{ELM}}^{\text{pt}}(\mathcal{A}) \leq \frac{4.5\sigma_{\text{priv}}^2}{2^n},$$

$$\text{Adv}_{\text{ELM}}^{\text{uft}}(\mathcal{A}') \leq \left( \frac{2q_d}{2^\tau} + \frac{4.5\sigma_{\text{auth}}^2}{2^n} \right) + \left( \frac{2q_v}{2^\tau} + \frac{4.5\sigma_{\text{mac}}^2}{2^n} \right),$$

where  $\sigma_{\text{priv}}$ ,  $q_d$ ,  $\sigma_{\text{auth}}$ ,  $q_v$ , and  $\sigma_{\text{mac}}$  are parameters of  $\mathcal{A}_{\text{ae}}$ ,  $\mathcal{A}_{\text{ae}}^{\pm}$ . This means that the security of ELM can be reduced to that of the underlying block cipher. Also, the security bounds of ELM are standard birthday type ( $O(2^{n/2})$ ) as with GCM and OCB.



### E. Proof of Theorem 3

Let  $K_M$  and  $K_A$  be the keys of MAC and AE. Suppose that they are uniformly random and independent.

1) *Privacy*: We assume that  $\mathcal{A}$  is given  $K_M$ , denoted by  $\mathcal{A}(K_M)$ . Since  $\mathcal{A}(K_M)$  can compute the data associated with the root node and the intermediate nodes, we can assume that  $\mathcal{A}(K_M)$  obtains only the data associated with leaf nodes from the tree initialization oracle and the update oracle. Let  $\mathcal{A}_{ae}$  be the privacy adversary against AE. If  $\mathcal{A}(K_M)$  queries IT (resp. IT-\$),  $\mathcal{A}_{ae}$  can simulate it by querying AE. $\mathcal{E}$  (resp. \$) in the same manner as Alg. 5. Note that a query to IT/IT-\$ invokes nonce-respecting encryption queries of  $\mathcal{A}_{ae}$  since  $\text{add}(\text{leaf}(i)) \parallel \text{ctr}(\text{leaf}(i)) \neq \text{add}(\text{leaf}(j)) \parallel \text{ctr}(\text{leaf}(j))$  necessarily holds for  $1 \leq i \neq j \leq b^d$ . Regarding the update queries,  $\mathcal{A}(K_M)$  invokes AE. $\mathcal{D}$ . However, AE. $\mathcal{D}$  in the update queries always outputs  $\top$  since  $\sigma_i = \tilde{\sigma}_{i-1}$  for  $i \in [q]$ . Thus,  $\mathcal{A}_{ae}$  can always output  $\top$  regardless of inputs to simulate AE. $\mathcal{D}$  in update queries. The adversary  $\mathcal{A}_{ae}$  can simulate the leftover pure update function in the same manner as the simulation of the initialization oracles. Also, the update query of  $\mathcal{A}(K_M)$  invokes nonce-respecting encryption queries of  $\mathcal{A}_{ae}$  due to the node-unique property of  $\text{add}(\cdot)$  and the one-time property of  $\text{ctr}(\cdot)$ . Namely, the sequence of queries in the privacy game of  $\mathcal{A}(K_M)$  can be simulated by  $\mathcal{A}_{ae}$ . Thus,  $\text{Adv}_{\text{PAT}_2}^{\text{priv}}(\mathcal{A})$  can be upper-bounded by

$$\begin{aligned} & |\Pr[\mathcal{A}(K_M)^{\text{IT,UT}} \rightarrow 1] - \Pr[\mathcal{A}(K_M)^{\text{IT-}\$, \text{UT-}\$} \rightarrow 1]| \\ &= \text{Adv}_{\text{AE}}^{\text{priv}}(\mathcal{A}_{ae}), \end{aligned}$$

where  $\mathcal{A}_{ae}$  queries  $b^d + q$  times to encryption oracle because an initialization query of  $\mathcal{A}(K_M)$  invokes  $b^d$  encryption queries of  $\mathcal{A}_{ae}$  and update queries of  $\mathcal{A}(K_M)$  invoke  $q$  encryption queries of  $\mathcal{A}_{ae}$ .

2) *Unforgeability*: Let  $\mathcal{A}'_{ma}$  denote an adversary who queries MAC. $\mathcal{T}$ , AE. $\mathcal{E}$ , MAC. $\mathcal{V}$ , and AE. $\mathcal{D}$ . We show  $\mathcal{A}'_{ma}$  can simulate IT, UT, and VT following the rules of the security notions of MAC and AE. As in the case of the privacy advantage,  $\mathcal{A}'_{ma}$  can simulate IT and UT by querying MAC. $\mathcal{T}$  and AE. $\mathcal{E}$ . Note that all queries are nonce-respecting. For the simulation of IT,  $\mathcal{A}'_{ma}$  queries  $(b^d - 1)/(b - 1)$  times to MAC. $\mathcal{T}$  and queries  $b^d$  times to AE. $\mathcal{E}$ . To simulate  $q'$  invocations of UT,  $\mathcal{A}'_{ma}$  queries  $q'd$  times to MAC. $\mathcal{T}$  and queries  $q'$  times to AE. $\mathcal{E}$ . The query to VT may invoke *replay queries* to MAC. $\mathcal{V}$  and AE. $\mathcal{D}$ . A replay query means that  $\mathcal{A}'_{ma}$  queries  $(N, M, T)$  (resp.  $(N, C, T)$ ) to MAC. $\mathcal{V}$  (resp. AE. $\mathcal{D}$ ) while  $(N, M)$  (resp.  $(N, M)$ ) s.t.  $(C, T) = \text{AE}.\mathcal{E}(N, M)$  has been queried to MAC. $\mathcal{T}$  (resp. AE. $\mathcal{E}$ ). Such queries are prohibited in the security notions of MAC and AE whereas they may appear in the simulation of unforgeability game since the game captures  $\mathcal{A}'$  performing a replay attack. To avoid replay queries, we suppose  $\mathcal{A}'_{ma}$  stores input/output of MAC. $\mathcal{T}$  and AE. $\mathcal{E}$  invoked in IT and UT. Let  $\text{List}_{\text{MAC}}$  and  $\text{List}_{\text{AE}}$  be the set of such tuples. Specifically, elements of  $\text{List}_{\text{MAC}}$  and  $\text{List}_{\text{AE}}$  are defined as  $(\text{add}(u^\sigma) \parallel \text{ctr}(u^\sigma), \text{ctr}(\text{ch}_1(u^\sigma)) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u^\sigma)), \text{tag}(u^\sigma))$  and  $(\text{add}(\text{leaf}(i)^\sigma) \parallel \text{ctr}(\text{leaf}(i)^\sigma), C[i], \text{tag}(\text{leaf}(i)^\sigma))$ , where  $u$  is a node of a tree construction  $\sigma$  and  $i \in [b^d]$ . When  $\mathcal{A}'$  invokes replay queries to MAC. $\mathcal{V}$  and

AE. $\mathcal{D}$ ,  $\mathcal{A}'_{ma}$  can notice it by checking  $\text{List}_{\text{MAC}}$  and  $\text{List}_{\text{AE}}$ , and can simulate the response of MAC. $\mathcal{V}$  and AE. $\mathcal{D}$  by just returning  $\top$ . Regarding the queries except for replay queries,  $\mathcal{A}'_{ma}$  can simulate them by querying MAC. $\mathcal{V}$  and AE. $\mathcal{D}$  in the same manner as Alg. 6. For the simulation of VT,  $\mathcal{A}'_{ma}$  queries at most  $d$  times to MAC. $\mathcal{V}$  and queries at most one time to AE. $\mathcal{D}$ . From the above discussion, we obtain

$$\Pr[\mathcal{A}'^{\text{IT,UT,VT}} \rightarrow 1] = \Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V, AE.D}} \rightarrow 1], \quad (4)$$

where  $\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V, AE.D}}$  means that  $\mathcal{A}'$  queries to  $\mathcal{A}'_{ma}$  pretending to be functions IT, UT, VT.

We define new adversary  $\mathcal{A}''_{ma}$  querying MAC. $\mathcal{T}$ , AE. $\mathcal{E}$ , and  $\perp_{\text{T}}$  to simulate IT, UT, and  $\perp_{\text{T}}$ .  $\mathcal{A}''_{ma}$  can simulate IT and UT by employing MAC. $\mathcal{T}$  and AE. $\mathcal{E}$  in the same way that  $\mathcal{A}'_{ma}$  does. For  $\perp_{\text{T}}$ ,  $\mathcal{A}''_{ma}$  only needs to mediate  $\mathcal{A}'$ 's query to  $\perp_{\text{T}}$ . Thus, we obtain the following equation.

$$\Pr[\mathcal{A}'^{\text{IT,UT,}\perp_{\text{T}}} \rightarrow 1] = \Pr[\mathcal{A}'^{\mathcal{A}''_{ma} \text{ MAC.T, AE.E,}\perp_{\text{T}}} \rightarrow 1]. \quad (5)$$

From (4) and (5), we obtain

$$\begin{aligned} \text{Adv}_{\text{PAT}_2}^{\text{uft}}(\mathcal{A}') &= |\Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V, AE.D}} \rightarrow 1] \\ &\quad - \Pr[\mathcal{A}'^{\mathcal{A}''_{ma} \text{ MAC.T, AE.E,}\perp_{\text{T}}} \rightarrow 1]|. \quad (6) \end{aligned}$$

We can evaluate the upper bound of (6) by adding the following three inequation's left sides and right sides, respectively, then obtain (3).

$$\begin{aligned} & |\Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V, AE.D}} \rightarrow 1] \\ &\quad - \Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V,}\perp_{\text{AE}}} \rightarrow 1]| \\ &\leq \text{Adv}_{\text{AE}}^{\text{auth}}(\mathcal{A}_{ae}^\pm), \quad (7) \end{aligned}$$

$$\begin{aligned} & |\Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E, MAC.V,}\perp_{\text{AE}}} \rightarrow 1] \\ &\quad - \Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E,}\perp_{\text{MAC,}\perp_{\text{AE}}} \rightarrow 1]| \\ &\leq \text{Adv}_{\text{MAC}}^{\text{mac}}(\mathcal{A}_{\text{mac}}), \quad (8) \end{aligned}$$

$$\begin{aligned} & |\Pr[\mathcal{A}'^{\mathcal{A}'_{ma} \text{ MAC.T, AE.E,}\perp_{\text{MAC,}\perp_{\text{AE}}} \rightarrow 1] \\ &\quad - \Pr[\mathcal{A}'^{\mathcal{A}''_{ma} \text{ MAC.T, AE.E,}\perp_{\text{T}}} \rightarrow 1]| = 0, \quad (9) \end{aligned}$$

where  $\perp_{\text{AE}}(\cdot, \cdot, \cdot)$  (resp.  $\perp_{\text{MAC}}(\cdot, \cdot, \cdot)$ ) is the function for decryption (resp. verification) queries to AE (resp. MAC), which always returns  $\perp$  for any inputs.  $\mathcal{A}_{ae}^\pm$  is an authenticity adversary against AE performing  $b^d + q'$  encryption queries and one decryption query.  $\mathcal{A}_{\text{mac}}$  is an adversary against MAC performing  $(b^d - 1)/(b - 1) + q'd$  tagging queries and  $d$  verification queries. The remaining part of this section is devoted to the proof of (7), (8), and (9). We can prove (7) and (8) in the same manner, thus we omit the proof of (8). Suppose that  $\mathcal{A}'_{ma}$  eventually outputs a bit in her simulation of the unforgeability game and  $\mathcal{A}'$  outputs the same bit as  $\mathcal{A}'_{ma}$ . Then the left side of (7) is upper-bounded by the probability that  $\mathcal{A}'_{ma}$  querying MAC. $\mathcal{T}$ , AE. $\mathcal{E}$ , and MAC. $\mathcal{V}$  successfully distinguishes AE. $\mathcal{D}$  from  $\perp_{\text{AE}}$ . Then it also can be upper-bounded by the probability that  $\mathcal{A}'_{ma}(K_M)$  querying AE. $\mathcal{E}$  successfully distinguishes AE. $\mathcal{D}$  from  $\perp_{\text{AE}}$ . Without loss of generality, we can assume that this distinguishing probability

is equal to the probability that  $\mathcal{A}'_{\text{ma}}(K_M)$  querying  $\text{AE.E}$  and  $\text{AE.D}$  obtains something other than  $\perp$  from  $\text{AE.D}$  under the unforgeability game for ME trees. The authenticity adversary against  $\text{AE}$ ,  $\mathcal{A}'_{\text{ae}}$ , can simulate the oracles that  $\mathcal{A}'_{\text{ma}}(K_M)$  queries because the query sequence of  $\mathcal{A}'_{\text{ma}}(K_M)$  respects the rule of authenticity game for AE schemes (performing nonce-respecting queries to  $\text{AE.E}$  and not performing replay queries to  $\text{AE.D}$ ). Thus, we obtain inequation (7).

We then prove (9). Recall that  $(idx', \sigma')$  is the tree verification query of  $\mathcal{A}'$  and  $(u_0^{\sigma'}, \dots, u_d^{\sigma'})$  is the path of nodes from the root to the specified leaf. The left side of (9) can be seen as the probability that  $\mathcal{A}'$  querying  $\text{IT}$  and  $\text{UT}$  obtains  $\top$  from the tree verification oracle simulated by  $\perp_{\text{MAC}}$ ,  $\perp_{\text{AE}}$ ,  $\text{ListMAC}$ , and  $\text{ListAE}$ . This is the probability that all the data associated with  $u_0^{\sigma'}, \dots, u_d^{\sigma'}$  consist of the elements of  $\text{ListMAC}$  and  $\text{ListAE}$ . We prove the probability is equal to zero in the following claim.

*Claim 1:* Either (a) or (b) described below must hold.

- (a) There exists  $i \in \{0, \dots, d-1\}$  s.t. the tuple  $(\text{add}(u_i^{\sigma'}) \parallel \text{ctr}(u_i^{\sigma'}), \text{ctr}(\text{ch}_1(u_i^{\sigma'})) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_i^{\sigma'})), \text{tag}(u_i^{\sigma'})) \notin \text{ListMAC}$ .
- (b)  $(\text{add}(u_d^{\sigma'}) \parallel \text{ctr}(u_d^{\sigma'}), C^{\sigma'}[idx], \text{tag}(u_d^{\sigma'})) \notin \text{ListAE}$ .

This claim states that  $\mathcal{A}'_{\text{ma}}$  querying  $\text{MAC.T}$ ,  $\text{AE.E}$ ,  $\perp_{\text{MAC}}$  and  $\perp_{\text{AE}}$  has to query to  $\perp_{\text{MAC}}$  or  $\perp_{\text{AE}}$  in her simulation of the tree verification query. Thus, she always obtains  $\perp$  from  $\perp_{\text{MAC}}$  or  $\perp_{\text{AE}}$  and she returns it to  $\mathcal{A}'$ . Note that HJ05 shows almost the same claim and its proof, and thus, we omit the proof of the claim.

## V. IMPLEMENTATION AND EVALUATION

We evaluate the hardware implementation of ELM using logic synthesis. We use AES-128, thus  $n = 128$ . Counter, address and tag is set to 64 bits, so  $\alpha = \beta = \tau = n/2$ . The 64-bit tag corresponds to the security level of SIT for a fair comparison. We focus on a high-throughput and area-time efficient architecture based on an unrolled and pipelined AES datapath, similar to that of SIT in [8], whose throughput is one block encryption per clock cycle. This high-throughput architecture suits the context of memory encryption. Note that our architecture can utilize other block ciphers and architectures (e.g., round-based and byte-serial ones) in accordance with the optimization goals.

### A. Hardware Architectures

Figure 4 shows the proposed hardware architecture of PXOR-MAC. The primary inputs consist of a block index, the number of branches  $\text{len}$  ( $= b$ ), nonce  $(\text{add}(u_i^{\sigma'}) \parallel \text{ctr}(u_i^{\sigma'}))$ , two  $n$ -bit keys, and an  $n$ -bit segmented plaintext block  $(\text{ctr}(\text{ch}_{2j+1}(u_i^{\sigma'})) \parallel \text{ctr}(\text{ch}_{2j+2}(u_i^{\sigma'})))$  ( $0 \leq j \leq b/2 - 2$ ), and the primary output is given as tag. One plaintext block is fed to the hardware every clock cycle one after another and an encoding is completed with 11 clock cycles. In this architecture, the AES datapath is fully unrolled and pipelined. The pipeline registers are inserted at the boundaries of each round in order to increase the throughput. This enables the encryption of one plaintext block in one clock cycle with the frequency corresponding to the critical path of one round datapath.

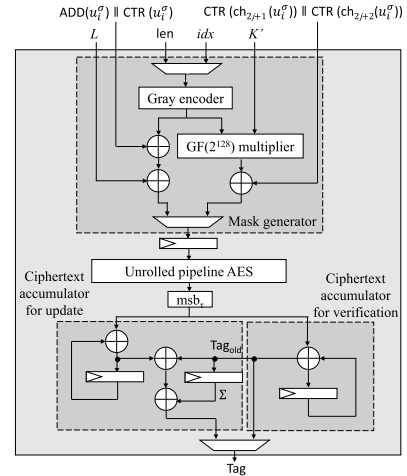


Fig. 4. Proposed MAC hardware architecture.

An up-to-date AES round datapath with a tower-field S-box presented in [25] is adopted for ELM (and SIT [8] for a comparison in this paper) in the following hardware implementation. A mask value for the input block to the AES core ( $K' \cdot i$  in Alg. 2) is generated by the multiplication of a gray code (converted from a block index) and a key  $K'$  over  $\text{GF}(2^n)$ . The conversion from a block index to a gray code is given by a combinational circuit and the generation of a mask value is implemented using a  $(n \times \log b)$ -bit GF multiplier. This multiplier generates mask values from all indices in a tree with  $b$  branches in one clock cycle. The mask value for the nonce ( $N_o \oplus K' \cdot m \oplus L$  in Alg. 2) is computed as the sum of the last mask value and  $L$  without any GF multiplication. The accumulation in the tag generator is implemented by a feedback loop consisting of a bit-parallel-XOR (i.e.,  $\text{GF}(2^{n/2})$  adder) and registers, which realizes the **for** loop at lines 2–2 in Alg. 2.

Figure 5 shows the hardware architectures of the proposed AE, where one encryption core and one decryption core are utilized. Both cores are unrolled and pipelined similarly to the above MAC hardware to ensure high throughput. The encryption and decryption cores are separately implemented (without unifying them like [25], [26]) in order to perform a decryption in the pre-verification process and an encryption in the update process simultaneously for UT. The pre-mask and post-mask generators compute the mask values for the input and output of encryption/decryption, respectively. The proposed architecture utilizes two pre-mask generators and two post-mask generators for simultaneous encryption and decryption. The field doubling and tripling for mask value generation are achieved by combinational circuit blocks denoted by  $\times 2$  and  $\times 3$ , which consist of four and 132 two-way XOR gates, respectively. The architecture can be implemented with less area than another one consisting of one pre-mask generator, one post-mask generator, and two 128-bit-wide first-in first-out (FIFO) buffers.<sup>7</sup> Plaintext accumulators obtain

<sup>7</sup>The mask value generated by the pre-mask generator should be retained for ten clock cycles for post-mask addition. This indicates that we require a  $(128 \times 10)$ -bit register to implement one FIFO if we use AES as the block cipher, which consumes a larger area and power than the four mask generators in our architecture.

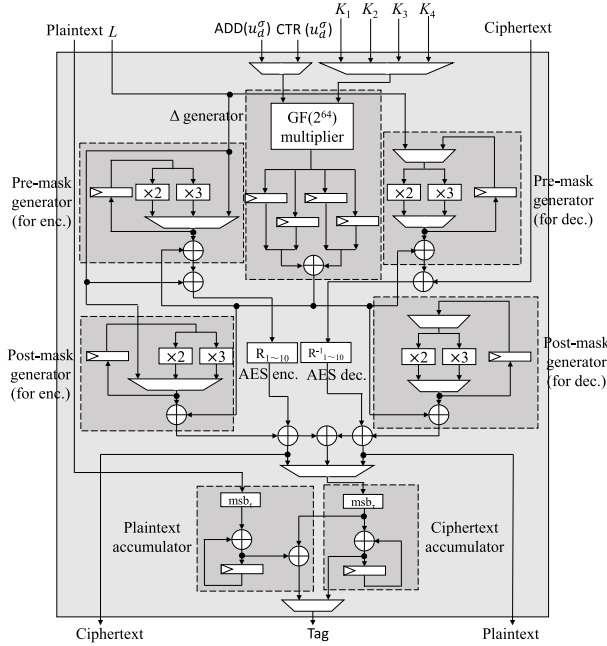


Fig. 5. Proposed AE hardware architecture.

truncated plaintext blocks to generate the tag, which is given by a feedback loop with a bit-parallel-XOR (*i.e.*,  $GF(2^{n/2})$  adder).  $\tilde{E}_K^{N,0,0}(0^n)$  is added before outputting the tag.

The mask value  $\Delta$  is generated from a nonce by the  $\Delta$  generator module following Alg. 4, using a  $((n/2) \times (n/2))$ -bit GF multiplier with four clock cycles. The generated  $\Delta$  is added to the mask values at the pre/post-mask generators.

### B. Performance Evaluation

We evaluate performance of the proposed architectures, on the basis of logic synthesis results. We assume that one MAC module is used at the top and each middle layer in the tree and one AE module is used at the lowest layer. Thus we use  $d$  MAC modules and one AE module, which fully exploits the parallelism of ELM. We investigate the best-case performance given a constraint in area and power (*i.e.*, the number of available MAC modules), and clarify the area-latency trade-offs from the evaluation results.

We use Synopsys Design Compiler I-2013.12-SP5 and Nangate 15nm Open Cell Library. The performance with  $d = 3, 5$ , and  $7$  are evaluated, following [8]. Table II lists the area obtained from the logic synthesis. We set the timing-constraint for the synthesis to the operating frequency of 4GHz, assuming modern high-end CPUs operating at 3GHz or faster. The synthesis result found no timing violation, thus the proposed architecture can be used even for modern high-end CPUs without degrading the system clock frequency. For comparison, Table II also lists the synthesis results of SIT (following [8]) implemented under the same conditions and assumption as above.

Table II shows that our implementation has a similar size as that of SIT, and the gap is closer as depth grows. The architectures for Flat-OCB require both encryption and decryption

TABLE II  
CIRCUIT AREAS SYNTHESIZED WITH 4GHz TIMING CONSTRAINT [GE]

depth $d$	SIT	ELM
3	421,083.25	519,910.50
5	601,544.50	717,801.00
7	914,155.25	915,612.50




Fig. 6. Numbers of clock cycles to update and verify tag.

cores, which resulted in a larger area than the inverse-free AE of SIT. However, our MAC hardware needs only one AES encryption core, while SIT needs a  $GF(2^{64})$  multiplier in addition to one AES encryption core.

Fig. 6 shows the numbers of clock cycles (*i.e.*, latency) of ELM and SIT. The tuple (Number of branches, Chunk size) determines the sizes of the covered memory region. Incremental SIT (Incr. SIT for short) indicates the evaluation result of SIT when UT is performed in an incremental manner. We stress that such an incremental update has not been shown in the literature including [8]; we do this for a fair comparison. Each clock cycle shown here is given by a larger one of either AE or MAC.

Fig. 6 shows that the advantage of ELM is greater as the covered region gets larger. One major reason is that ELM utilizes a  $n$ -bit block cipher for encryption, whereas SIT processes a plaintext in a  $n/2$ -bit-wise manner (*i.e.*, inner-product MAC over  $GF(2^{n/2})$ ). More precisely, since the MAC module at each layer (and AE module) should process more bits for a larger parameter, the  $n$ -bit-wise computation of PXOR-MAC in ELM enables fewer calls of the underlying pseudorandom function than the  $n/2$ -bit-wise computation of SIT, which results in a lower latency of ELM. In addition, the number of cycles in the update process of AE is reduced by using a distinct decryption core to perform the pre-verification and update processes simultaneously. Note that SIT uses an Encrypt-then-MAC style AE composing counter mode encryption and an inner-product MAC. Since SIT does not utilize any decryption function and the MAC computation becomes critical for the latency, the latency of SIT cannot be reduced in the same manner as ours. The precise numbers of cycles for various tree parameters are shown in Appendix A. The results suggest that ELM is superior to SIT when covering a larger region. As the protected memory region becomes larger, the advantage of ELM increases significantly.



TABLE III  
REQUIRED MEMORY SIZES

	SIT	ELM
on-chip	$320 + \max\{\ell, 64b\}$	952
off-chip	$113 \times \sum_{i=0}^d b^i - 56$	$112 \times \sum_{i=0}^d b^i - 56$

Table III lists the on-chip and off-chip memory sizes for each architecture. Here, both counter and tag lengths are 56 bits to be (perfectly) compliant with SIT. With respect to the on-chip storage, ELM requires four 64-bit round keys for  $\Delta$  gen., a 128-bit  $L$ , a 128-bit plaintext/ciphertext processing key  $K$ , and a 56-bit  $\text{ctr}(u^c)$ . Here, the amounts of on-chip storage is constant regardless of parameters  $b$  and  $\ell$ . In contrast, SIT needs to store the  $2 \times 128$ -bit keys and inner-product MAC for nonce processing and mask generation, depending on the size of the tree parameters. As a result, the on-chip size gets larger when  $b$  and  $\ell$  are larger because the key length of inner-product MAC increases in proportion to the input block length ( $b$ ). For example, the on-chip sizes of SIT are 768 and 8,448 bits when  $b = 8$ ,  $\ell = 512$  and  $b = 128$ ,  $\ell = 8,192$ , respectively. On the other hand, the off-chip size of ours is comparable with that of SIT. For example, when  $d = 3$  and  $b = 8$ , the off-chip sizes of SIT and ours are 66,049 and 65,464 bits, respectively.

## VI. DISCUSSION

### A. Design Optimizations

Our evaluations did not consider system constraints in order to demonstrate the scalability of our proposal. In practice, we would design the total hardware configuration depending on various system/architecture constraints, such as system clock frequency, memory size, available on-chip size, memory bandwidth, cache memory structure, and so on. The sizes of the MAC and AE modules should also be determined considering the above constraints.

A gap in latency between AE and MAC also leads to a loss of efficiency for ME tree because the entire latency is determined by a larger latency of either AE or MAC. While we evaluated the typical tree structures in Section V-B using power-of-two parameters, they should be determined such that the latencies of AE and MAC are well-balanced. In summary, when designing the ME tree and its hardware, we should first determine the optimal (*i.e.*, well-balanced) tree structure parameters for the required covered region. Then, we can mitigate the remaining gap in latency between AE and MAC on the basis of the above hardware optimization approach.

### B. Application of Split Counter

ELM can use Split Counter [10], the up-to-date method to reduce the amount of counters for ME trees. Figure 7 shows the numbers of cycles when an overflow of minor counters occurs in MAC and AE, respectively. (See Appendix A for more details.) Here, we evaluate the cases of four different numbers of branches. Since the reset counter is always the

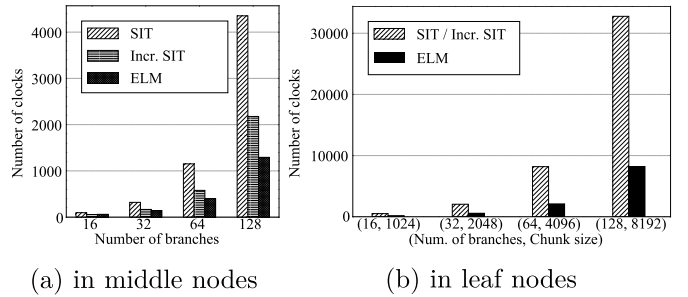


Fig. 7. Number of clocks when minor counter overflows.

same value, the encryption result of reset counter in MAC can be pre-computed for both SIT and our trees. Thus, ELM maintains superiority to SIT under the condition where an overflow occurs. We also found that the proposed AE is advantageous even with the split counter thanks to the simultaneous execution of pre-verification and update. In particular, when the number of branches increases, the proposed scheme becomes more advantageous in comparison with that without the split counter.

## VII. CONCLUSION

We have presented ELM, a new memory encryption scheme with tree-based authentication. Unlike many recent proposals from computer architecture perspective, we focus on the internal MAC and AE modes, including their interactions, to reduce the entire read/write latency. ELM combines fully parallelizable MAC and AE modes and utilizes the incremental property of the MAC mode. Our AE mode is similar to OCB, however has a better decryption latency and it can be of independent interest as a stand-alone AE mode. We provide provable security results for these components as well as the whole scheme. Since Intel SGX's scheme (SIT) is a representative work on the same direction, we instantiated ELM using the same AES and compared ELM with SIT, and presented preliminary hardware implementations. The results showed that ELM achieves significantly lower latency, while keeping the comparable implementation size of SIT. Several future directions can be considered, as follows:

### A. Other Instantiations

A low-latency (tweakable) block cipher such as PRINCE [27] or QARMA [28] will significantly improve both latency and size of ELM. Using multiple primitives of possibly different block sizes may further boost the performance.

### B. Side-Channel Attacks

Cryptographic hardware frequently needs to be resistant against side-channel attacks (SCAs). Design and evaluation of SCA-resistant hardware architecture for ELM is an interesting future topic. The use of a SCA-resistant AES implementation will thwart typical SCAs trying to recover the AES key. For example, we can use a low-latency masked round-based AES [29]. However, it will significantly reduce throughput

TABLE IV  
NUMBERS OF CLOCKS TO UPDATE AND VERIFY TAG

$b$	$\ell$	Update			Verify		Covered region [Byte]		
		SIT	Incr. SIT	ELM	SIT / Incr. SIT	ELM	$d = 3$	$d = 5$	$d = 7$
8	512	<b>20</b>	<b>20</b>	21	<b>14</b>	18	32K	2M	134M
	1,024	32	32	<b>25</b>	<b>18</b>	22	65K	4M	268M
	2,048	64	64	<b>33</b>	32	<b>30</b>	131K	8M	536M
	4,096	128	128	<b>49</b>	64	<b>46</b>	262K	16M	1G
	8,192	256	256	<b>81</b>	128	<b>78</b>	524K	33M	2G
16	512	32	<b>20</b>	22	<b>16</b>	20	262K	67M	17G
	1,024	32	32	<b>25</b>	<b>18</b>	22	524K	134M	34G
	2,048	64	64	<b>33</b>	32	<b>30</b>	1M	268M	68G
	4,096	128	128	<b>49</b>	64	<b>46</b>	2M	536M	137G
	8,192	256	256	<b>81</b>	128	<b>78</b>	4M	1G	274G
32	512	64	33	<b>30</b>	32	<b>28</b>	2M	2G	2T
	1,024	64	33	<b>30</b>	32	<b>28</b>	4M	4G	4T
	2,048	64	64	<b>33</b>	32	<b>30</b>	8M	8G	8T
	4,096	128	128	<b>49</b>	64	<b>46</b>	16M	17G	17T
	8,192	256	256	<b>81</b>	128	<b>78</b>	33M	34G	35T
64	512	128	65	<b>46</b>	64	<b>44</b>	16M	68G	281T
	1,024	128	65	<b>46</b>	64	<b>44</b>	33M	137G	562T
	2,048	128	65	<b>46</b>	64	<b>44</b>	67M	274G	1P
	4,096	128	128	<b>49</b>	64	<b>46</b>	134M	549G	2P
	8,192	256	256	<b>81</b>	128	<b>78</b>	268M	1T	4P
128	512	256	129	<b>78</b>	128	<b>76</b>	134M	2T	36P
	1,024	256	129	<b>78</b>	128	<b>76</b>	268M	4T	72P
	2,048	256	129	<b>78</b>	128	<b>76</b>	536M	8T	144P
	4,096	256	129	<b>78</b>	128	<b>76</b>	1G	17T	288P
	8,192	256	256	<b>81</b>	128	<b>78</b>	2G	35T	576P

due to a large area overhead and on-the-fly random number generation.

The usage of a (first-order) masking-friendly light-weight (tweakable) block cipher such as PRESENT [30] or Skinny [31] would be a practical alternative with less area overhead and no on-the-fly randomness.

Furthermore, it would be interesting to design leakage-resilient TBC/permutation-based AE (e.g., [32], [33]) and MAC that enable low-latency operation and are suitable to be used with ELM.

#### APPENDIX A

##### DETAILED PERFORMANCE EVALUATION

Table IV of page 15 shows the numbers of clock cycles (i.e., latency) and the size of protected memory region (namely, the covered region) of ELM and SIT in details for various tree parameters. Each clock cycle shown here is given by a larger one of either AE or MAC. For example, if our ME tree has eight branches  $b = 8$  and handles 512-bit blocks  $\ell = 512$  [bit] at the leaf (or lowest) node, ELM requires 18 and 21 clock cycles for MAC and AE, respectively; hence the clock cycle of this ME tree is given as 21 clock cycles in the table. The bold-face characters in each row highlight the scheme that achieved the lowest latency (minimal clock cycles) under the parameter condition of the row. The parameters used in Fig. 6 in Section V-B are underlined in the chunk size column in Table IV. The table shows the results of all tree structures comprehensively, where some rows hatched in gray indicate better clock cycles than those in white in terms of the

latency required for the covered regions. For example, a tree with  $b = 8$  and  $\ell = 4,096$  requires a larger latency and a smaller covered region than that with  $b = 16$  and  $\ell = 512$ , and therefore there is no reason to use the former tree rather than the latter. Such meaningless parameters are caused by the gap in latency between AE and MAC, as discussed in Section VI.

In addition, we show the numbers of clock cycles in details for the cases that the split counter is applied. The split counter is a method to reduce the amount of counters stored in an off-chip for ME trees [10]. It uses two types of counters: major and minor ones. A major counter is shared by the children nodes of the node of interest (or a parent node), and each child node is equipped with a minor counter. In other words, in an ME tree with split counter, children nodes having the same parent node share the upper bits of the same major counter. Here, we should point out that overflows of the minor counters frequently occur, since each minor counter is given with a small bit length. When such an overflow occurs, the corresponding major counter is incremented and all the minor counters associated with it are reset to zero. Accordingly, we need to update all the tags where the nonces are reset. The tag update with the split counter requires  $b$  times tag updates at the layer where a major counter is incremented (i.e., the overflow of minor counter occurs), which is non-trivial in the update process (i.e., UT). In ELM,  $b$  counters are originally used as the input for the plaintext part of tag generation by the MAC algorithm, that is,  $b \times n/2$  bits should be verified by MAC. More precisely, let  $\text{ctr}$  be the counter of the parent

TABLE V  
NUMBER OF CLOCKS WHEN MINOR COUNTER  
IN MIDDLE NODES OVERFLOWS

$b$	SIT	Incr. SIT	ELM
8	32	26	29
16	96	57	61
32	320	167	141
64	1,152	579	397
128	4,352	2,177	1,293

node and let  $\text{ctr}'_1, \text{ctr}'_2, \dots, \text{ctr}'_b$  be the counters of children nodes without the split counter, where  $b$  is the number of branches in the tree structure. In ELM, a tag  $T$  is generated as

$$T = \text{PXOR-MAC}.\mathcal{T}_{K,K'}\left(\text{add} \parallel \text{ctr}\right), \\ \left(\text{ctr}'_1 \parallel \text{ctr}'_2 \parallel \dots \parallel \text{ctr}'_b\right),$$

where  $\text{PXOR-MAC}.\mathcal{T}_{K,K'}(N, M)$  calculates a tag from a nonce  $N$  and a plaintext  $M$ . Each counter is given by  $n/2$  bits and  $b \times n/2$  bits should be encrypted.

In contrast, we consider the tag generation when utilizing the split counter. Let  $\text{Mctr}$  and  $\text{mctr}$  be the major and minor counters of a parent node, respectively. Let  $\text{Mctr}'$  and  $\text{mctr}'_1, \text{mctr}'_2, \dots, \text{mctr}'_b$  be the major counter and minor counters of children nodes, respectively. In this case, unless the overflow of minor counters occurs, a tag is generated as

$$T = \text{PXOR-MAC}.\mathcal{T}_{K,K'}\left(\text{add} \parallel \text{Mctr} \parallel \text{mctr}\right), \\ \left(\text{Mctr}' \parallel \text{mctr}'_1 \parallel \dots \parallel \text{mctr}'_b\right).$$

Here, let  $s$  and  $t$  be the bit lengths of major and minor counters, respectively ( $s + t = n/2$ ). When the split counter is used, the input is given with  $s + bt$  bits. Since  $s + bt \leq b(s + t)/2$ , the input bit length of the MAC algorithm is reduced thanks to the split counter. In addition to the tag generation algorithm (*i.e.*,  $\text{PXOR-MAC}.\mathcal{T}$ ), VT (*i.e.*,  $\text{PXOR-MAC}.\mathcal{V}$ ) and UT (*i.e.*,  $\text{PXOR-MAC}.\mathcal{VU}$ ) algorithms are performed with the split counter as well. The tags of leaf (or lowest) nodes can also be generated, verified, and updated in a similar manner even when the split counter is applied.

As a result, Table V of page 15 and Table VI of page 16 show the numbers of clock cycles when the split counter is applied and an overflow occurs in MAC and AE, respectively. When the major counter is incremented due to the overflow of a minor counter, all minor counters associated with the major counter are reset to 0.

#### APPENDIX B PROOF OF CLAIM 1

In this section, we show the proof of Claim 1. The literature [4] shows almost the same claim and its proof for the proposed ME tree without encryption of leaf nodes, however, the proof is a little bit complex. We recast it for the sake of

TABLE VI  
NUMBER OF CLOCKS WHEN MINOR COUNTER  
IN LEAF NODES OVERFLOWS

$b$	$\ell$	SIT/ Incr. SIT	ELM
8	512	136	49
	1,024	260	81
	2,048	512	147
	4,096	1,024	273
	8,192	2,048	529
16	512	136	49
	1,024	516	145
	2,048	1,024	275
	4,096	2,048	529
	8,192	4,096	1,041
32	512	520	145
	1,024	1,028	273
	2,048	2,048	531
	4,096	4,096	1,041
	8,192	8,192	2,065
64	512	1,032	273
	1,024	2,052	529
	2,048	4,096	1,043
	4,096	8,192	2,065
	8,192	16,384	4,113
128	512	2,056	529
	1,024	4,100	1,041
	2,048	8,192	2,067
	4,096	16,384	4,113
	8,192	32,768	8,209

ease to understand and the ME tree with encryption of leaf nodes.

If **(b)** occurs, the claim is simply proved. We need to see that if **(b)** does not occur, the case **(a)** must hold. First we discuss  $u_0$  (*i.e.*, the root node). Let  $(N_{u_0}, M_{u_0}, T_{u_0}) = (\text{add}(u_0^{\sigma'}) \parallel \text{ctr}(u_0^{\sigma'}), \text{ctr}(\text{ch}_1(u_0^{\sigma'})) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_0^{\sigma'})), \text{tag}(u_0^{\sigma'}))$ . If  $(N_{u_0}, M_{u_0}, T_{u_0}) \notin \text{List}_{\text{MAC}}$  holds, it means that **(a)** holds. Suppose that  $(N_{u_0}, M_{u_0}, T_{u_0}) \in \text{List}_{\text{MAC}}$ , and we obtain the following equation.

$$(N_{u_0}, M_{u_0}, T_{u_0}) \\ = (\text{add}(u_0^{\tilde{\sigma}'}) \parallel \text{ctr}(u_0^{\tilde{\sigma}'}), \\ \text{ctr}(\text{ch}_1(u_0^{\tilde{\sigma}'})) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_0^{\tilde{\sigma}'})), \text{tag}(u_0^{\tilde{\sigma}'})), \quad (10)$$

which means that the data stored in  $u_0^{\sigma'}$  is the same as that stored in  $u_0^{\tilde{\sigma}'}$ . This holds because  $N_{u_0}$  cannot be tampered by definition, and the element of  $\text{List}_{\text{MAC}}$  including  $N_{u_0}$  is uniquely determined as (10) since nonces included in  $\text{List}_{\text{MAC}}$  are distinct. Note that we also obtain  $\text{ctr}(u_1^{\sigma'}) = \text{ctr}(u_1^{\tilde{\sigma}'})$  from (10).

Next, we discuss  $u_1$ . Let  $(N_{u_1}, M_{u_1}, T_{u_1}) = (\text{add}(u_1^{\sigma'}) \parallel \text{ctr}(u_1^{\sigma'}), \text{ctr}(\text{ch}_1(u_1^{\sigma'})) \parallel \dots \parallel \text{ctr}(\text{ch}_b(u_1^{\sigma'})), \text{tag}(u_1^{\sigma'}))$ . As well as the case of  $u_0$ , we can suppose that  $(N_{u_1}, M_{u_1}, T_{u_1}) \in \text{List}_{\text{MAC}}$  since **(a)** occurs when  $(N_{u_1}, M_{u_1}, T_{u_1}) \notin \text{List}_{\text{MAC}}$  holds. In the same manner as



the case of  $u_0$ , we obtain

$$\begin{aligned} & (N_{u_1}, M_{u_1}, T_{u_1}) \\ &= (\text{add}(u_1^{\tilde{\sigma}'}) \parallel \text{ctr}(u_1^{\tilde{\sigma}'}), \\ & \quad \text{ctr}(\text{ch}_1(u_1^{\tilde{\sigma}'})) \parallel \cdots \parallel \text{ctr}(\text{ch}_b(u_1^{\tilde{\sigma}'})), \text{tag}(u_1^{\tilde{\sigma}'})), \end{aligned}$$

since  $\text{add}(u_1^{\tilde{\sigma}'})$  cannot be tampered and  $\text{ctr}(u_1^{\tilde{\sigma}'}) = \text{ctr}(u_1^{\tilde{\sigma}'})$  holds from (10).

Suppose that we repeat the same discussion as  $u_0$  and  $u_1$ . For  $2 \leq i \leq d-1$ , we define  $(N_{u_i}, M_{u_i}, T_{u_i}) = (\text{add}(u_i^{\tilde{\sigma}'}) \parallel \text{ctr}(u_i^{\tilde{\sigma}'}), \text{ctr}(\text{ch}_1(u_i^{\tilde{\sigma}'})) \parallel \cdots \parallel \text{ctr}(\text{ch}_b(u_i^{\tilde{\sigma}'})), \text{tag}(u_i^{\tilde{\sigma}'}))$ . When  $(N_{u_i}, M_{u_i}, T_{u_i}) \in \text{ListMAC}$  for  $0 \leq i \leq d-1$ , we obtain the following equation.

$$\begin{aligned} & (N_{u_i}, M_{u_i}, T_{u_i}) \\ &= (\text{add}(u_i^{\tilde{\sigma}'}) \parallel \text{ctr}(u_i^{\tilde{\sigma}'}), \\ & \quad \text{ctr}(\text{ch}_1(u_i^{\tilde{\sigma}'})) \parallel \cdots \parallel \text{ctr}(\text{ch}_b(u_i^{\tilde{\sigma}'})), \text{tag}(u_i^{\tilde{\sigma}'})). \quad (11) \end{aligned}$$

Finally, we discuss  $u_d$  (i.e., the leaf node). Let  $(N_{u_d}, C_{u_d}, T_{u_d}) = (\text{add}(u_d^{\tilde{\sigma}'}) \parallel \text{ctr}(u_d^{\tilde{\sigma}'}), C^{\tilde{\sigma}'}[idx], \text{tag}(u_d^{\tilde{\sigma}'}))$ . Recall that we assumed that **(b)** did not occur, thus  $(N_{u_d}, C_{u_d}, T_{u_d}) \in \text{ListAE}$ . Here,  $\text{add}(u_d^{\tilde{\sigma}'}) = \text{add}(u_d^{\tilde{\sigma}'})$  holds since  $\text{add}(\cdot)$  cannot be tampered, and  $\text{ctr}(u_d^{\tilde{\sigma}'}) = \text{ctr}(u_d^{\tilde{\sigma}'})$  holds due to (11) when  $i = d-1$ . Thus, we have

$$\begin{aligned} & (N_{u_d}, C_{u_d}, T_{u_d}) \\ &= (\text{add}(u_d^{\tilde{\sigma}'}) \parallel \text{ctr}(u_d^{\tilde{\sigma}'}), C^{\tilde{\sigma}'}[idx], \text{tag}(u_d^{\tilde{\sigma}'})), \quad (12) \end{aligned}$$

since nonces included in  $\text{ListAE}$  are distinct, hence the element of  $\text{ListAE}$  including  $N_{u_d}$  is uniquely determined as (12).

From (11) and (12), we proved that the data stored in  $u_i^{\tilde{\sigma}'}$  is the same as that stored in  $u_i^{\tilde{\sigma}'}$  for all  $i \in \{0, \dots, d\}$ , which is a forbidden query. Therefore, there must exist  $i \in \{0, \dots, d-1\}$  such that  $(N_{u_i}, M_{u_i}, T_{u_i}) \notin \text{ListMAC}$ . This concludes the claim.

## REFERENCES

- [1] M. Dworkin, *Recommendation for Block Cipher Modes of Operation: The XTS-AES Mode for Confidentiality on Storage Devices*, document NIST Special Publication (SP) 800-38E, 2010.
- [2] J. A. Halderman *et al.*, “Lest we remember: Cold boot attacks on encryption keys,” in *Proc. USENIX Secur. Symp.*, 2008, pp. 45–60.
- [3] R. C. Merkle, “A digital signature based on a conventional encryption function,” in *CRYPTO’87* (Lecture Notes in Computer Science), vol. 293, C. Pomerance, Ed. Berlin, Germany: Springer, Aug. 1988, pp. 369–378.
- [4] W. E. Hall and C. S. Jutla, “Parallelizable authentication trees,” in *SAC 2005* (Lecture Notes in Computer Science), vol. 3897, B. Preneel and S. Tavares, Eds. Berlin, Germany: Springer, Aug. 2006, pp. 95–109.
- [5] B. Rogers, S. Chhabra, M. Prvulovic, and Y. Solihin, “Using address independent seed encryption and bonsai Merkle trees to make secure processors OS- and performance-friendly,” in *Proc. 40th Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, Dec. 2007, pp. 183–196.
- [6] R. Elbaz, D. Champagne, R. B. Lee, L. Torres, G. Sassatelli, and P. Guillemin, “TEC-tree: A low-cost, parallelizable tree for efficient defense against memory replay attacks,” in *CHES 2007* (Lecture Notes in Computer Science), vol. 4727, P. Paillier and I. Verbauwhede, Eds. Berlin, Germany: Springer, Sep. 2007, pp. 289–302.
- [7] S. Gueron, “Memory encryption for general-purpose processors,” *IEEE Secur. Privacy*, vol. 14, no. 6, pp. 54–62, Nov. 2016.
- [8] S. Gueron, “A memory encryption engine suitable for general purpose processors,” *Cryptol. ePrint Arch., Tech. Rep. 2016/204*, 2016. [Online]. Available: <https://eprint.iacr.org/2016/204>
- [9] M. Dworkin, *Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*, document NIST Special Publication 800-38D, 2007.
- [10] C. Yan, D. Engländer, M. Prvulovic, B. Rogers, and Y. Solihin, “Improving cost, performance, and security of memory encryption and authentication,” in *Proc. 33rd Int. Symp. Comput. Archit. (ISCA)*, May 2006, pp. 179–190.
- [11] M. Taassori, A. Shafiee, and R. Balasubramonian, “VAULT: Reducing paging overheads in SGX with efficient integrity verification structures,” in *Proc. ASPLOS*. New York, NY, USA: ACM, 2018, pp. 665–678.
- [12] G. Saileshwar, P. J. Nair, P. Ramrakhiani, W. Elsasser, J. A. Joao, and M. K. Qureshi, “Morphable counters: Enabling compact integrity trees for low-overhead secure memories,” in *Proc. 51st Annu. IEEE/ACM Int. Symp. Microarchitecture (MICRO)*, Oct. 2018, pp. 416–427.
- [13] M. Bellare, O. Goldreich, and S. Goldwasser, “Incremental cryptography: The case of hashing and signing,” in *CRYPTO’94* (Lecture Notes in Computer Science), vol. 839, Y. Desmedt, Ed. Berlin, Germany: Springer, Aug. 1994, pp. 216–233.
- [14] M. Bellare and D. Micciancio, “A new paradigm for collision-free hashing: Incrementality at reduced cost,” in *EUROCRYPT’97* (Lecture Notes in Computer Science), vol. 1233, W. Fumy, Ed. Berlin, Germany: Springer, May 1997, pp. 163–192.
- [15] M. Bellare, R. Guérin, and P. Rogaway, “XOR MACs: New methods for message authentication using finite pseudorandom functions,” in *CRYPTO’95* (Lecture Notes in Computer Science), vol. 963, D. Coppersmith, Ed. Berlin, Germany: Springer, Aug. 1995, pp. 15–28.
- [16] T. Krovetz and P. Rogaway, “The software performance of authenticated-encryption modes,” in *FSE 2011* (Lecture Notes in Computer Science), vol. 6733, A. Joux, Ed. Berlin, Germany: Springer, Feb. 2011, pp. 306–327.
- [17] M. Liskov, R. L. Rivest, and D. Wagner, “Tweakable block ciphers,” in *CRYPTO 2002* (Lecture Notes in Computer Science), vol. 2442, M. Yung, Ed. Berlin, Germany: Springer, Aug. 2002, pp. 31–46.
- [18] P. Rogaway, “Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC,” in *ASIACRYPT 2004* (Lecture Notes in Computer Science), vol. 3329, P. J. Lee, Ed. Berlin, Germany: Springer, Dec. 2004, pp. 16–31.
- [19] M. Bellare and C. Namprempre, “Authenticated encryption: Relations among notions and analysis of the generic composition paradigm,” in *ASIACRYPT 2000* (Lecture Notes in Computer Science), vol. 1976, T. Okamoto, Ed. Berlin, Germany: Springer, Dec. 2000, pp. 531–545.
- [20] W. E. Hall and C. S. Jutla, “Parallelizable authentication trees,” *IACR Cryptol. ePrint Arch.*, vol. 2002, p. 190, Jan. 2002.
- [21] T. Unterluggauer, M. Werner, and S. Mangard, “MEAS: Memory encryption and authentication secure against side-channel attacks,” *J. Cryptograph. Eng.*, vol. 9, no. 2, pp. 137–158, Jun. 2019.
- [22] M. Bellare, A. Desai, E. Jorjipii, and P. Rogaway, “A concrete security treatment of symmetric encryption,” in *Proc. IEEE 38th Annu. Symp. Found. Comp. Sci.*, Oct. 1997, pp. 394–403.
- [23] K. Minematsu and T. Matsushima, “Generalization and extension of XEX\* mode,” *IEICE Trans. Fundam. Electron., Commun. Comput. Sci.*, vol. E92-A, no. 2, pp. 517–524, 2009.
- [24] M. Bellare, O. Goldreich, and A. Mityagin, “The power of verification queries in message authentication and authenticated encryption,” *Cryptol. ePrint Arch., Tech. Rep. 2004/309*, 2004. [Online]. Available: <https://eprint.iacr.org/2004/309>
- [25] R. Ueno *et al.*, “High throughput/gate AES hardware architectures based on datapath compression,” *IEEE Trans. Comput.*, vol. 69, no. 4, pp. 534–548, Apr. 2020.
- [26] R. Ueno, S. Morioka, N. Homma, and T. Aoki, “A high throughput/gate AES hardware architecture by compressing encryption and decryption datapaths—Toward efficient CBC-mode implementation,” in *CHES* (Lecture Notes in Computer Science), vol. 9813, Cham, Switzerland: Springer, 2016, pp. 538–558.
- [27] J. Borghoff *et al.*, “PRINCE—A low-latency block cipher for pervasive computing applications—Extended abstract,” in *ASIACRYPT 2012* (Lecture Notes in Computer Science), vol. 7658, X. Wang and K. Sako, Eds. Berlin, Germany: Springer, Dec. 2012, pp. 208–225.
- [28] R. Avanzi, “The QARMA block cipher family,” *IACR Trans. Symm. Cryptol.*, vol. 2017, no. 1, pp. 4–44, 2017.
- [29] P. Sasdrich, B. Bilgin, M. Hutter, and M. E. Marson, “Low-latency hardware masking with application to AES,” *IACR Trans. Cryptograph. Hardw. Embedded Syst.*, vol. 2020, pp. 300–326, Mar. 2020.

- [30] A. Bogdanov *et al.*, “PRESENT: An ultra-lightweight block cipher,” in *Proc. 9th Int. Workshop Cryptograph. Hardw. Embedded Syst. (CHES)*, in Lecture Notes in Computer Science, vol. 4727, P. Pailier and I. Verbauwhede, Eds. Vienna, Austria: Springer, Sep. 2007, pp. 450–466.
- [31] C. Beierle *et al.*, “The SKINNY family of block ciphers and its low-latency variant MANTIS,” in *CRYPTO 2016, Part II Advances in Cryptology—CRYPTO 2016* (Lecture Notes in Computer Science), vol. 9815, M. Robshaw and J. Katz, Eds. Berlin, Germany: Springer, Aug. 2016, pp. 123–153.
- [32] C. Dobraunig, M. Eichlseder, S. Mangard, F. Mendel, and T. Unterluggauer, “ISAP—Towards side-channel secure authenticated encryption,” *IACR Trans. Symmetric Cryptol.*, vol. 2017, pp. 80–105, Mar. 2017.
- [33] F. Berti, C. Guo, O. Pereira, T. Peters, and F.-X. Standaert, “TEDT: A leakage-resistant AEAD mode,” *IACR TCHES*, vol. 2020, no. 1, pp. 256–320, 2019. [Online]. Available: <https://tches.iacr.org/index.php/TCHES/article/view/8400>



**Akiko Inoue** received the B.E. and M.E. degrees from Kyushu University in 2015 and 2017, respectively. She joined NEC in 2017. She works on design and analysis of symmetric-key ciphers and related systems. She received the Best Paper Award at CRYPTO 2019.



**Kazuhiko Minematsu** received the B.E., M.E., and D.S. degrees from Waseda University in 1996, 1998, and 2008, respectively. He joined NEC in 1998. He currently works as a Research Fellow. Since 2019, he has also been a Visiting Professor with Yokohama National University. His research interests are design and analysis of symmetric-key ciphers and its application systems. He received the best paper awards at FSE 2015 and CRYPTO 2019.



**Maya Oda** received the B.E. and M.S. degrees from Tohoku University in 2019 and 2021, respectively. Her research interests include hardware security, cryptographic implementation, and memory protection.



**Rei Ueno** (Member, IEEE) is currently an Assistant Professor with the Research Institute of Electrical Communication, Tohoku University. He is also working with the JST as a Researcher for a PRESTO Project. His research interests include arithmetic circuits, cryptographic implementations, formal verification, and hardware security.



**Naofumi Homma** (Senior Member, IEEE) received the B.E., M.S., and Ph.D. degrees from Tohoku University in 1997, 1999 and 2001, respectively. He is currently a Professor with the Research Institute of Electrical Communication, Tohoku University. He is also working as a Researcher with the JST CREST. His research interests include hardware security, computer arithmetic, and EDA methodology.