

# Constructing Committing and Leakage-Resilient Authenticated Encryption

Patrick Struck<sup>1</sup> and Maximiliane Weishäupl<sup>2</sup>

<sup>1</sup> Universität Konstanz, Konstanz, Germany, [patrick.struck@uni-konstanz.de](mailto:patrick.struck@uni-konstanz.de)

<sup>2</sup> Universität Regensburg, Regensburg, Germany, [maximiliane.weishaeupl@ur.de](mailto:maximiliane.weishaeupl@ur.de)

**Abstract.** The main goal of this work is to construct authenticated encryption (AE) that is both committing and leakage-resilient. As a first approach for this we consider generic composition as a well-known method for constructing AE schemes. While the leakage resilience of generic composition schemes has already been analyzed by Barwell et al. (Asiacrypt’17), for committing security this is not the case. We fill this gap by providing a separate analysis of the generic composition paradigms with respect to committing security, giving both positive and negative results: By means of a concrete attack, we show that Encrypt-then-MAC is not committing. Furthermore, we prove that Encrypt-and-MAC is committing, given that the underlying schemes satisfy security notions we introduce for this purpose. We later prove these new notions achievable by providing schemes that satisfy them. MAC-then-Encrypt turns out to be more difficult due to the fact that the tag is not outputted alongside the ciphertext as it is done for the other two composition methods. Nevertheless, we give a detailed heuristic analysis of MAC-then-Encrypt with respect to committing security, leaving a definite result as an open task for future work. Our results, in combination with the fact that only Encrypt-then-MAC yields leakage-resilient AE schemes, show that one cannot obtain AE schemes that are both committing and leakage-resilient via generic composition. As a second approach for constructing committing and leakage-resilient AE, we develop a generic transformation that turns an arbitrary AE scheme into one that fulfills both properties—though only a slightly weakened form of leakage resilience. The transformation relies on a keyed function that is both binding, i.e., it is hard to find key-input pairs that result in the same output, leakage-resilient pseudorandom and unpredictable.

**Keywords:** Authenticated Encryption · Committing Security · Leakage Resilience

---

**\*Errata:** We noticed that the version published in ToSC 2024(1) contains an error. The transform given there does not achieve the claimed leakage resilience security notion. In this version, we give a different transform that restores the initial claim for a slightly weaker leakage resilience security notion. A more detailed errata about the mistake and fix is given in [Section 4](#). The results in [Section 3](#) are unaffected. We thank the authors of [BMOS17] for helpful feedback regarding their leakage security notions. We further thank the anonymous reviewers for the valuable feedback. Patrick Struck acknowledges support by the Hector Foundation II. Work of Maximiliane Weishäupl was funded by the German Federal Ministry of Education and Research (BMBF) under the project Quant-ID (16KISQ111).

# 1 Introduction

Authenticated encryption (AE) with associated data is the modern day version of what cryptography was historically: a method to secure the communication between two parties. Classically, secure communication is understood to entail confidentiality and authenticity. The former prevents anyone except the rightful receiver of a ciphertext to recover the message from it, while the latter ensures that the ciphertext originates indeed from the claimed sender and was not manipulated during transmission. In an AE scheme, the ciphertext is generated by encryption of the message under a context  $(K, N, A)$ , consisting of a key  $K$ , a nonce  $N$ , and associated data  $A$ . To achieve the above security goals, both the message and associated data need to be authenticated, whereas confidentiality is required only for the message.

Besides the standard security notions—confidentiality and authenticity—there are several, more advanced security notions. This encompasses, for instance, security against related-key attacks [FKOS22] which enables the adversary to obtain ciphertexts generated by keys that exhibit a certain relation, anonymous AE [CR19] where ciphertexts have to conceal the nonce to provide anonymity of the communicating parties, leakage-resilient AE [BMOS17] where the adversary can obtain extra information via side-channels, and committing security [BH22] which prevents adversaries from finding ciphertexts that decrypt under more than one key. These advanced security notions are typically analyzed isolated and not in conjunction. However, depending on the use-case, a joint consideration of multiple security notions might be sensible. As there is little analysis on the interplay between the notions, unwanted effects might occur: While there are transformations, which one applies to an AE scheme to obtain certain security properties, concatenating them might not have the desired outcome. A successive application of two or more transformations could result in a scheme that is not secure with respect to all of the security notions, but—in the worst case—only the notion corresponding to the last transformation that was applied. This stems from the fact that these transformations are usually proven to achieve the desired security property while maintaining the standard AE security, i.e., the last transformation might subvert the additional guarantees from the previous ones. On the other hand, it is possible that a transformation achieves security with respect to additional notions besides the one it was developed for. This would make further changes to the scheme redundant and might result in unnecessary overhead. In total, deepening our understanding of the relation between advanced security notions brings various benefits with it.

Two areas that have gained a lot of attention are leakage resilience [BMOS17, DJS19, DEM<sup>+</sup>20, BPPS17, BGP<sup>+</sup>19, DM21, DM19, MOSW15, BGPS21] and committing security [BH22, CR22, MLGR23, KSW23, DMVA23], though they were never considered in conjunction. We focus on developing AE schemes that fulfill both properties. Firstly, having such schemes allows deployment in different scenarios even if each of it requires only one of the properties. Secondly, we think that this concrete combination of security notions has also practical relevance: Considering key-recovery attacks on embedded devices, both leakage-resilience and committing security are necessary properties. More precisely, consider an AE scheme that is either leakage-resilient or committing but not both. Using side-channel attacks, an adversary might be able to extract the secret key if the scheme is only committing secure. If the scheme is only leakage-resilient an adversary could apply a *Partitioning Oracle Attack* [LGR21] to learn the key: the adversary crafts ciphertexts that are valid for a subset of keys. If decryption succeeds, it knows that the target key is within this set and the attack continues using smaller and smaller subsets until it finds the key.

Towards the goal of developing schemes that fulfill both notions, we start by considering generic composition as a well-known method of constructing AE schemes. Here, an AE scheme is obtained by “gluing together” a symmetric encryption (SE) scheme and a message authentication code (MAC) such that security of the AE scheme follows from security of the

two underlying components. There are three general methods: Encrypt-and-MAC (E&M), Encrypt-then-MAC (EtM), and MAC-then-Encrypt (MtE). A common misconception is that only EtM is secure and the other two approaches (E&M and MtE) are not. While initial work [BN00] indeed showed that only EtM is secure in general, later work [NRS14] provided a more fine-grained analysis. Here, associated data was taken into account as an input to the MAC and encryption schemes based on random IVs or nonces were considered. This led to the development of several secure AE schemes derived from the three classical composition methods: eight A-schemes that are composed of an IV-based encryption scheme and a vector-MAC; and three N-schemes which rely on nonce-based encryption for the underlying scheme and again a vector-MAC. Throughout this work, we will mainly focus on the latter schemes N1, N2, and N3, which follow the E&M, EtM, and MtE approach, respectively. The three N-schemes have been analyzed both in the leakage setting [BMOS17] and against related-key attacks [FKOS22]. However, they have yet to be analyzed with respect to their committing security.

## 1.1 Contribution

We first study the committing security of generic composition.<sup>1</sup> We show that N2 (and EtM in general) does not achieve committing security. The fact that N2 authenticates the ciphertext together with the circumstance that colliding ciphertexts are easy to find—due to encryption schemes being reversible by design—enables a generic attack. For N1 (following E&M), we prove that committing security can be reduced to corresponding collision resistance properties of the underlying encryption and MAC.<sup>2</sup> We call these properties wCR and CR and show that they are achievable security notions by providing instances that fulfill them—namely, the symmetric encryption scheme and MAC underlying SLAE [DJS19]. To check this, we first note that SLAE follows the FGHF' construction and hence one can further break down the notions wCR and CR to the functions  $\mathcal{F}$ ,  $\mathcal{G}$ ,  $\mathcal{H}$ , and  $\mathcal{F}'$ . We then prove that the instantiations for these functions used in SLAE fulfill the desired properties. Our results for N3 (following MtE) are not that definite as the ones for N1 and N2: The main difficulty lies in the fact that the tag gets encrypted alongside the message, whereas for N1 and N2, the tag gets appended to the ciphertext—thus any committing attack against N1 and N2 requires identical tags whereas those against N3 do not. We give an heuristic analysis, narrowing down the possible attacks and discussing a number of common attack strategies. However, a complete analysis of MtE with respect to committing security is still an open task and left for future work. Nevertheless, we can conclude that committing and leakage-resilient AE cannot be built via generic composition, as the only leakage-resilient generic composition method is Encrypt-then-MAC [BMOS17], which we prove to be not secure with respect to committing security. As a second approach for building committing and leakage-resilient AE, we turn towards generic transformations as a way to achieve the desired security properties. We observe that none of the existing transformations from the literature are suitable for our purpose. The problem is that all of them hash the key, hence the leakage resilience of a scheme might be lost after its application if the hash function leaks information. To address this problem, we develop a generic transformation that turns an arbitrary AE scheme, a hash function, and a keyed function into an AE scheme that is both leakage-resilient and committing. Note that for leakage-resilience, we consider the LAE notion from Barwell et al. [BMOS17], however, slightly weakened adopting the approach by Shrimpton and Terashima [ST13] for AE

<sup>1</sup>We focus on the N-schemes, though it is easy to see that the results also apply to the A-schemes: In case of committing security, the difference between unique nonces and random IVs becomes obsolete as the adversary can choose them at will.

<sup>2</sup>In order to get a non-trivial bound for the symmetric encryption scheme, the message length needs to be of sufficient length (e.g., 128 bits). In light of real-world committing attacks, like the *Facebook message franking attack* [DGRW18], this is a reasonable assumption.

security. At the core of the transformation is the keyed function which needs to be both pseudorandom and unpredictable under leakage as well as binding. We show existence of such a function by proving that the sponge-based function used in SLAE fulfills the binding property—pseudorandomness and unpredictability under leakage was already proven in [DJS19]. For the underlying AE scheme and hash function we rely on default AE security and collision resistance, respectively.

In total, our consideration of committing security and leakage resilience of AE schemes in conjunction, yields both negative and positive results. On the negative side, our results indicate that committing security and leakage resilience do not work well together for generic composition. On the positive side, we develop a simple generic transformation that achieves both security notions, thereby meeting our goal of building committing and leakage-resilient AE.

## 1.2 Related Work

Generic composition was initially introduced in [BN00] as a method to construct AE schemes and refined in [NRS14]. It is, especially from a theoretical point of view, an important method and was studied in various settings [FKOS22, BPP18, Ber23, BMOS17, KS20].

The necessity of committing security for AE schemes is due to real-world attacks that exploit the lack thereof [DGRW18, LGR21, ADG<sup>+</sup>22]. The first committing security notions for AE schemes were formalized in [BH22]. Later works [CR22, MLGR23] provided a more fine-grained framework for committing security notions. Several works show committing attacks against various AE schemes: GCM/AES-GCM-SIV [BH22], GCM/OCB [CR22], CCM/EAX/SIV [MLGR23], AEZ [CFG<sup>+</sup>23], and the NIST LWC finalists [KSW23]. However, there are also positive results: CAU [BH22], CTX [CR22], the sponge-based NIST LWC finalists ASCON, ISAP, and SCHWAEEMM [KSW23], SPONGEWRAP [DFG23], and a sponge-based AE scheme based on SHAKE [DMVA23] were proven to be committing.

Leakage-resilient cryptography was formalized in [DP08], following the “Only Computation Leaks Information” paradigm [MR04]. Security notions for authenticated encryption incorporating leakage were developed in [BMOS17], which also analyzed the leakage resilience of generic composition, showing inherent weaknesses for Encrypt-and-MAC and MAC-then-Encrypt. Several works developed AE schemes that are designed to be leakage-resilient: ISAP [DEM<sup>+</sup>17, DEM<sup>+</sup>20], SLAE/FGHF’ [DJS19, KS20], TEDT [BGP<sup>+</sup>19], ROMULUS-T [IKM<sup>+</sup>21], and SIVAT [BMOS17]. Furthermore, there are results concerning the leakage resilience of existing AE modes: PYJAMASK, PHOTON-BEETLE, ASCON, SPOOK, ISAP, and TEDT, are analyzed in [BBC<sup>+</sup>20], which shows different results regarding what parts need to be protected against leakage. Similarly, [VCS22] shows that out of the NIST LWC finalists, three modes (ASCON, ISAP, and ROMULUS-T) are advantageous when hardening their implementations against leakage. Lastly, [GPPS20] is worth highlighting, which provides a methodology that allows to analyze the leakage resilience of duplex sponges, e.g., ISAP, ASCON, GIBBON, and TEDTSPONGE—two of which were also shown to be committing.

Overall, there are many results with respect to either notion—for ASCON/ISAP even for both<sup>3</sup> individually—we are not aware of any work explicitly targeting both of them.

## 2 Preliminaries

**Notation.** By  $\{0, 1\}^*$ , we denote the set of bit strings with arbitrary length. The empty bit string is written as  $\varepsilon$ . For a bit string  $X$  and an integer  $y$ ,  $\lceil X \rceil_y$  (resp.  $\lfloor X \rfloor_y$ ) denote the leftmost (resp. rightmost)  $y$  bits of  $X$ . Concatenation of bit strings  $X$  and  $Y$  is written

<sup>3</sup>ROMULUS was also analyzed with respect to both leakage resilience and committing security [VCS22, KSW23], though for different variants (ROMULUS-N and ROMULUS-T).




Figure 1: The AE schemes N1 (left), N2 (middle), and N3 (right) [NRS14] in terms of an underlying symmetric encryption scheme ( $\text{Enc}, \text{Dec}$ ) and MAC ( $\text{Tag}, \text{Ver}$ ).

as  $X \parallel Y$ . We use game-based proofs [BR06] and write  $\text{G}(\mathcal{A}) \rightarrow x$ , to denote that the output of game  $\text{G}$ , when played by  $\mathcal{A}$ , is  $x$ . By  $\mathcal{A}^{\text{G}} \rightarrow x$ , we denote that  $\mathcal{A}$  outputs  $x$ , when playing game  $\text{G}$ . Throughout this work, we write  $\text{IV}$  for some public, fixed initialization vector used in the constructions.

**Primitives.** The focus of this work is on authenticated encryption (AE) scheme with associated data, which is a pair of two algorithms ( $\text{Enc}, \text{Dec}$ ). The encryption algorithm  $\text{Enc}: \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$  takes a key  $K$ , a nonce  $N$ , associated data  $A$ , and a message  $M$  as input, and outputs a ciphertext  $C$ . The decryption algorithm  $\text{Dec}: \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C} \rightarrow \mathcal{M} \cup \{\perp\}$  takes a key  $K$ , a nonce  $N$ , associated data  $A$ , and a ciphertext  $C$  as input, and outputs a message  $M$  or a special symbol  $\perp$ . The key space, nonce space, associated data space, message space, and ciphertext space are denoted by the sets  $\mathcal{K}$ ,  $\mathcal{N}$ ,  $\mathcal{A}$ ,  $\mathcal{M}$ , and  $\mathcal{C}$ , respectively. Throughout this work, we consider these sets to be bit strings of certain length, more precisely,  $\mathcal{K} = \{0, 1\}^\kappa$ ,  $\mathcal{N} = \{0, 1\}^\nu$ ,  $\mathcal{A} = \{0, 1\}^*$ ,  $\mathcal{M} = \{0, 1\}^*$ , and  $\mathcal{C} = \{0, 1\}^* \times \{0, 1\}^\tau$ . We abuse notation and write either  $(C, T)$  or  $C \parallel T$ . The length of a message is denoted by  $m$ . *Correctness* of an AE scheme means that for any  $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ , we have  $\text{Dec}(K, N, A, \text{Enc}(K, N, A, M)) = M$ . All schemes that we consider satisfy the *tidyness* property [NRS14]: it holds that  $M = \text{Dec}(K, N, A, C)$  implies that  $C = \text{Enc}(K, N, A, M)$ . A symmetric encryption scheme is defined equivalently, except that there is no associated data and decryption does not return  $\perp$ . Following [MLGR23], the triple  $(K, N, A)$  is called a *context*.

A message authentication code (MAC) consists of two algorithms ( $\text{Tag}, \text{Ver}$ ). The tagging algorithm  $\text{Tag}: \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^\tau$  takes as input a key  $K$  and an element  $X$ . It returns a tag  $T$  of size  $\{0, 1\}^\tau$ . The verification algorithm  $\text{Ver}: \mathcal{K} \times \mathcal{X} \times \{0, 1\}^\tau \rightarrow \{0, 1\}$  takes as input a key  $K$ , an element  $X$ , and a tag  $T$  and outputs 1 indicating that the input is valid, or otherwise 0. Throughout this work, we consider  $\mathcal{K} = \{0, 1\}^\kappa$  and  $\mathcal{X} = \{0, 1\}^*$ . *Correctness* of a MAC means that for any  $(K, X) \in \mathcal{K} \times \mathcal{X}$ , it holds that  $\text{Ver}(K, X, \text{Tag}(K, X)) = 1$ .

The N-schemes [NRS14] construct an AE scheme from a symmetric encryption scheme and a MAC. Figure 1 illustrates the schemes N1, N2, and N3, while their pseudocodes are provided in Figure 14.

**Sponges.** Sponges are a tool to construct primitives. They rely on an  $n$ -bit state  $S$  that is constantly updated by applying a permutation  $\pi: \{0, 1\}^n \rightarrow \{0, 1\}^n$  or a transformation  $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$  to it. A P-Sponge is centered around a permutation whereas a T-Sponge is centered around a transformation. In this work, we focus on the latter. In between two invocations of the permutation/transformation the sponge can *absorb* an input by XORing it to the first  $r$  bits of the state  $S$  or *squeeze* an output by outputting the first  $r$  bits of the state  $S$ . In the simplest setting—which is all we need in this work—the sponge will first absorb the whole input by constantly absorbing it  $r$  bits at a time and




Figure 2: The symmetric encryption scheme  $\text{SE}[\mathcal{F}, \mathcal{G}]$  (left), composed of a function  $\mathcal{F}$  and a pseudorandom generator  $\mathcal{G}$ , and the MAC  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  (right), composed of a hash function  $\mathcal{H}$  and a function  $\mathcal{F}'$ .

then squeeze the output  $r$  bits at a time. The capacity  $c = n - r$  is the part of the state that is neither modified by the input during the absorption phase nor outputted in the squeezing phase. For the sponge-based proofs, the underlying transformation is assumed to be ideal, i.e., a random transformation to which the adversary gets oracle access.

In this work, we consider the sponge-based AE scheme SLAE [DJS19], more precisely, its underlying components. SLAE follows the generic construction FGHF' [DJS19, KS20], which constructs an AE scheme (via N2) from a symmetric encryption scheme  $\text{SE}[\mathcal{F}, \mathcal{G}]$  and a MAC  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  (cf. Figure 2), which in turn are composed of functions  $\mathcal{F}$  and  $\mathcal{F}'$ , a PRG  $\mathcal{G}$ , and a hash function  $\mathcal{H}$ . The AE scheme SLAE is obtained by instantiating  $\mathcal{F}/\mathcal{F}'$ ,  $\mathcal{G}$ , and  $\mathcal{H}$  with the sponge-based primitives SLFUNC, SPRG, and SvHASH, respectively. The primitives SLFUNC, SPRG, and SvHASH are illustrated in Figure 3, while their pseudocodes are given in Figure 16. The pseudocodes of the composed encryption scheme SLENC and the MAC SLMAC are provided in Figure 15.

### 3 Committing Security and Generic Composition

In this section, we analyze the committing security of the different generic composition paradigms. In Section 3.1, we analyze the three N-schemes N1, N2, and N3, given in [NRS14]. We continue with the generic encryption scheme  $\text{SE}[\mathcal{F}, \mathcal{G}]$  and MAC  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  [DJS19, KS20] in Section 3.2, and its sponge-based instantiation in Section 3.3.

#### 3.1 Committing Security of the N-Schemes

In this section, we analyze the committing security of N1, N2, and N3. We focus on the strongest form of committing security, which we define below. It asks the adversary to output two different contexts and messages that encrypt to the same ciphertext.

**Definition 1.** Let  $\text{AE} = (\text{Enc}, \text{Dec})$  be an authenticated encryption scheme and the game CMT be defined as in Figure 4. For any adversary  $\mathcal{A}$ , its CMT advantage is defined as

$$\text{Adv}_{\text{AE}}^{\text{CMT}}(\mathcal{A}) := \Pr[\text{CMT}(\mathcal{A}) \rightarrow 1].$$

##### 3.1.1 N2 (Encrypt-then-MAC)

The following theorem shows that the N2 construction does not achieve any committing security. Due to the fact that the N2 construction authenticates the ciphertext—compared to the N1 and N3 construction which both authenticate the message—there is a generic attack exploiting the fact that colliding ciphertexts are easy to find for symmetric encryption schemes: We obtain two messages by decrypting an arbitrary ciphertext under different




Figure 3: Illustration of the sponge-based primitives SLFUNC, SPRG, and SvHASH.

#### Game CMT

---

$(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$   
**if**  $(K, N, A) = (\bar{K}, \bar{N}, \bar{A})$   
     **return** 0  
 $(C, T) \leftarrow \text{Enc}(K, N, A, M)$   
 $(\bar{C}, \bar{T}) \leftarrow \text{Enc}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$   
**return**  $((C, T) = (\bar{C}, \bar{T}))$

Figure 4: Security game CMT for authenticated encryption schemes.

keys and the same nonce, and then authenticate said ciphertext, the nonce, and arbitrary associated data using some key for the MAC.

**Theorem 1.** *Let SE be a symmetric encryption scheme and MAC be a MAC. Let further  $\text{N2}[\text{SE}, \text{MAC}]$  be the authenticated encryption scheme obtained via the N2 construction using SE and MAC. Then there exists an adversary  $\mathcal{A}$  such that*

$$\text{Adv}_{\text{N2}[\text{SE}, \text{MAC}]}^{\text{CMT}}(\mathcal{A}) = 1.$$

*Proof.* We construct a CMT adversary  $\mathcal{A}$  against  $\text{N2}[\text{SE}, \text{MAC}]$  as shown in Figure 5. It picks an encryption key  $K_e$ , a MAC key  $K_m$ , a nonce  $N$ , associated data  $A$ , and a message  $M$  at random from the respective sets and computes the ciphertext  $C \leftarrow \text{SE.Enc}(K_e, N, M)$ . After sampling a different encryption key  $\bar{K}_e$  from  $\mathcal{K} \setminus \{K_e\}$  at random,  $\mathcal{A}$  computes  $\bar{M} \leftarrow \text{SE.Dec}(\bar{K}_e, N, C)$ .<sup>4</sup> Note that  $\text{SE.Enc}(\bar{K}_e, N, \bar{M}) = C$  as we assume the symmetric

<sup>4</sup>Note that the messages  $M$  and  $\bar{M}$  might be identical which does not affect the attack as the encryption keys are guaranteed to be distinct.



---

Adversary  $\mathcal{A}$

---

$K_e, N, M \leftarrow \mathcal{K} \times \mathcal{N} \times \mathcal{M}$   
 $C \leftarrow \text{SE.Enc}(K_e, N, M)$   
 $\bar{K}_e \leftarrow \mathcal{K} \setminus \{K_e\}$   
 $\bar{M} \leftarrow \text{SE.Dec}(\bar{K}_e, N, C) \quad // \text{ by tidyness: } C = \text{SE.Enc}(\bar{K}_e, N, \bar{M})$   
 $(K_m, A) \leftarrow \mathcal{K} \times \mathcal{A}$   
**return**  $((K_e, K_m), N, A, M), ((\bar{K}_e, K_m), N, A, \bar{M})$

Figure 5: Our generic adversary breaking committing security for any instantiation of N2.

encryption scheme to be tidy. Lastly the adversary sets  $K \leftarrow (K_e, K_m)$  and  $\bar{K} \leftarrow (\bar{K}_e, K_m)$ , and outputs  $(K, N, A, M), (\bar{K}, N, A, \bar{M})$ . Then  $\mathcal{A}$  is successful as  $K_e \neq \bar{K}_e$  implies  $K \neq \bar{K}$  and

$$\begin{aligned}
& \text{N2}[\text{SE}, \text{MAC}].\text{Enc}(K, N, A, M) \\
&= \text{SE.Enc}(K_e, N, M) \parallel \text{MAC.Tag}(K_m, (N, A, \text{SE.Enc}(K_e, N, M))) \\
&= C \parallel \text{MAC.Tag}(K_m, (N, A, C)) \\
&= \text{SE.Enc}(\bar{K}_e, N, \bar{M}) \parallel \text{MAC.Tag}(K_m, (N, A, \text{SE.Enc}(\bar{K}_e, N, \bar{M}))) \\
&= \text{N2}[\text{SE}, \text{MAC}].\text{Enc}(\bar{K}, N, A, \bar{M}).
\end{aligned}$$

This finishes the proof.  $\square$

### 3.1.2 N1 (Encrypt-and-MAC)

We first define two new security notions which enable us to show committing security of N1: CR and wCR. The former requires the adversary to find a collision of a MAC, i.e., two tuples of keys and inputs that differ yet result in the same tag. The latter is similar but defined for symmetric encryption and of a more restrictive nature. Instead of asking the adversary to find any collision, it is required to find *distinct* keys and *identical* nonce-message pairs that result in the same ciphertext. While there are no tidy encryption schemes that satisfy unrestricted collision resistance—as our attack against N2 illustrates—it turns out that this restricted form is achievable as we will show later.

**Definition 2.** Let  $\text{MAC} = (\text{Tag}, \text{Ver})$  be a message authentication code and the game CR be defined as in Figure 6. For any adversary  $\mathcal{A}$ , its CR advantage is defined as

$$\text{Adv}_{\text{MAC}}^{\text{CR}}(\mathcal{A}) := \Pr[\text{CR}(\mathcal{A}) \rightarrow 1].$$

**Definition 3.** Let  $\text{SE} = (\text{Enc}, \text{Dec})$  be a symmetric encryption scheme and the game wCR be defined as in Figure 6. For any adversary  $\mathcal{A}$ , its wCR advantage is defined as

$$\text{Adv}_{\text{SE}}^{\text{wCR}}(\mathcal{A}) := \Pr[\text{wCR}(\mathcal{A}) \rightarrow 1].$$

Note that, while the tag length (in game CR) is fixed, the ciphertext length (in game wCR) is controlled by the adversary through the choice of the messages. Thus, for very short messages the game wCR can be won easily.

The theorem below shows that the N1 construction achieves the strongest form of committing security, i.e., CMT, given that the underlying encryption scheme and MAC achieve our new security notions wCR and CR.



Game CR	Game wCR
$(K, X), (\bar{K}, \bar{X}) \leftarrow \mathcal{A}()$	$(K, N, M), (\bar{K}, \bar{N}, \bar{M}) \leftarrow \mathcal{A}()$
<b>if</b> $(K, X) = (\bar{K}, \bar{X})$	<b>if</b> $K = \bar{K} \vee (N, M) \neq (\bar{N}, \bar{M})$
<b>return</b> 0	<b>return</b> 0
$T \leftarrow \text{Tag}(K, X)$	$C \leftarrow \text{Enc}(K, N, M)$
$\bar{T} \leftarrow \text{Tag}(\bar{K}, \bar{X})$	$\bar{C} \leftarrow \text{Enc}(\bar{K}, \bar{N}, \bar{M})$
<b>return</b> $(T = \bar{T})$	<b>return</b> $(C = \bar{C})$

Figure 6: Security game CR for MACs and wCR for symmetric encryption schemes.

**Theorem 2.** *Let SE be a symmetric encryption scheme and MAC be a MAC. Let further  $\text{N1}[\text{SE}, \text{MAC}]$  be the authenticated encryption scheme obtained via the N1 construction using SE and MAC. Then for any adversary  $\mathcal{A}$  there exist adversaries  $\mathcal{B}$  and  $\mathcal{C}$  such that*

$$\text{Adv}_{\text{N1}[\text{SE}, \text{MAC}]}^{\text{CMT}}(\mathcal{A}) \leq \text{Adv}_{\text{SE}}^{\text{wCR}}(\mathcal{B}) + \text{Adv}_{\text{MAC}}^{\text{CR}}(\mathcal{C}).$$

*Proof.* Let  $\mathcal{A}$  be a CMT adversary against  $\text{N1}[\text{SE}, \text{MAC}]$  and denote his output by  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}))$ , where  $K = (K_e, K_m)$  with  $K_e$  and  $K_m$  being the encryption key and MAC key, respectively. The same for  $\bar{K}$  with encryption key  $\bar{K}_e$  and MAC key  $\bar{K}_m$ . Consider  $\text{E}$  to be the event that  $K_e \neq \bar{K}_e$  and  $(N, M) = (\bar{N}, \bar{M})$ . In the following, we abbreviate  $\text{N1}[\text{SE}, \text{MAC}]$  with  $\text{N1}$  as SE and MAC are clear from the context.

We construct a wCR adversary  $\mathcal{B}$  against SE. It runs  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})) \leftarrow \mathcal{A}()$  and outputs  $((K_e, N, M), (\bar{K}_e, \bar{N}, \bar{M}))$ . If  $\mathcal{A}$  is successful, we obtain  $(K, N, A, M) \neq (\bar{K}, \bar{N}, \bar{A}, \bar{M})$  and

$$\text{N1.Enc}(K, N, A, M) = \text{N1.Enc}(\bar{K}, \bar{N}, \bar{A}, \bar{M}). \quad (1)$$

By definition of the N1 construction, the encryption of the scheme N1 is given by

$$\text{N1.Enc}(K, N, A, M) = (\text{SE.Enc}(K_e, N, M), \text{MAC.Tag}(K_m, (N, A, M))). \quad (2)$$

Then we can conclude from Eq. (1) and Eq. (2) that

$$\text{SE.Enc}(K_e, N, M) = \text{SE.Enc}(\bar{K}_e, \bar{N}, \bar{M}).$$

Next assume that, additionally to  $\mathcal{A}$  being successful, event  $\text{E}$  holds. Then adversary  $\mathcal{B}$  succeeds as  $K_e \neq \bar{K}_e$  and  $(N, M) = (\bar{N}, \bar{M})$ , but their respective encryptions under SE agree. This yields

$$\Pr[\text{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{wCR}(\mathcal{B}) \rightarrow 1].$$

Next, we construct a CR adversary  $\mathcal{C}$  against MAC. It runs  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})) \leftarrow \mathcal{A}()$  and outputs  $((K_m, N, A, M), (\bar{K}_m, \bar{N}, \bar{A}, \bar{M}))$ . If  $\mathcal{A}$  succeeds, we can conclude from Eq. (1) and Eq. (2) that

$$\text{MAC.Tag}(K_m, (N, A, M)) = \text{MAC.Tag}(\bar{K}_m, (\bar{N}, \bar{A}, \bar{M})).$$

Assume that additionally to  $\mathcal{A}$  being successful, event  $\neg \text{E}$  holds, that means either  $K_e = \bar{K}_e$  or  $(N, M) \neq (\bar{N}, \bar{M})$ . If  $(N, M) \neq (\bar{N}, \bar{M})$  holds, then  $(K_m, N, A, M), (\bar{K}_m, \bar{N}, \bar{A}, \bar{M})$  are two different inputs for MAC. The same can be deduced if  $K_e = \bar{K}_e$ , as  $\mathcal{A}$  being successful implies that  $(K, N, A, M) \neq (\bar{K}, \bar{N}, \bar{A}, \bar{M})$  and since the encryption keys agree, the difference must come from the inputs to MAC. Then adversary  $\mathcal{C}$  wins the game wCR as it outputs two different tuples that generate the same tag. Thus we obtain

$$\Pr[\neg \text{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] \leq \text{Adv}_{\text{MAC}}^{\text{CR}}(\mathcal{C}). \quad (3)$$

Combining these results we obtain

$$\begin{aligned}
\mathbf{Adv}_{\text{N1}[\text{SE}, \text{MAC}]}^{\text{CMT}}(\mathcal{A}) &= \Pr[\text{CMT}(\mathcal{A}) \rightarrow 1] \\
&= \Pr[\text{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] + \Pr[\neg \text{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] \\
&\leq \Pr[\text{wCR}(\mathcal{B}) \rightarrow 1] + \Pr[\text{CR}(\mathcal{C}) \rightarrow 1] \\
&= \mathbf{Adv}_{\text{SE}}^{\text{wCR}}(\mathcal{B}) + \mathbf{Adv}_{\text{MAC}}^{\text{CR}}(\mathcal{C}).
\end{aligned}$$

This concludes the proof.  $\square$

### 3.1.3 N3 (MAC-then-Encrypt)

The core feature that differentiates N3 from the other two N-schemes is that the tag is no longer appended to the ciphertext, but instead is encrypted alongside the message. This means—at least in theory—that a valid committing attack against N3 can have  $T \neq \bar{T}$  and  $C = \bar{C}$  at the same time, which is impossible for both N1 and N2. It turns out that this case is the problematic one for proving security of N3.

We start by considering the case that the underlying encryption scheme and MAC fulfill the security properties wCR and CR security, respectively, which we defined for the security proof of N1. While it is not clear that these security properties are fitting for N3 as well, they seem to be a good starting point—in order to prove classical security, coinciding assumptions are required for N1 and N3.

Under the assumption that the wCR and CR properties hold for SE and MAC, we can show that any CMT attack with  $T = \bar{T}$  is ruled out, leaving only attacks for which  $T \neq \bar{T}$  holds. This can be checked as follows: Assume for sake of contradiction that  $\mathcal{A}$  is an adversary that wins the game CMT and its outputs  $((K_e, K_m), N, A, M)$  and  $((\bar{K}_e, \bar{K}_m), \bar{N}, \bar{A}, \bar{M})$  fulfill  $T = \bar{T}$ . If  $K_e = \bar{K}_e$  holds, at least one of the inputs to MAC must differ by definition of CMT. This contradicts the fact that MAC is CR-secure. If  $K_e \neq \bar{K}_e$  and  $(N, M) \neq (\bar{N}, \bar{M})$  holds, we also obtain a contradiction to MAC fulfilling CR security. Lastly, the case that  $K_e \neq \bar{K}_e$  and  $(N, M) = (\bar{N}, \bar{M})$  cannot occur either, as it would contradict the wCR security of SE.

We can further restrict the scope of possible attacks: As  $T \neq \bar{T}$  must hold, SE has two different messages as input (namely  $M \parallel T \neq \bar{M} \parallel \bar{T}$ ). By correctness, this implies that  $(K_e, N) \neq (\bar{K}_e, \bar{N})$  must hold, as otherwise  $\text{SE.Dec}(K_e, N, C)$  would have to equal both  $M \parallel T$  and  $\bar{M} \parallel \bar{T}$ . This leaves us with attacks for which  $T \neq \bar{T}$ , i.e.,  $(K_m, N, A, M) \neq (\bar{K}_m, \bar{N}, \bar{A}, \bar{M})$  and  $(K_e, N) \neq (\bar{K}_e, \bar{N})$ , holds. While this is not enough to prove security of N3 in general, the restriction of possible attacks might aid security proofs for certain schemes. One example where this is the case, is N3[SENC, SLMAC], i.e., SLAE but built following N3 instead of N2: In Section 3.3, we will see that the scheme's underlying components fulfill wCR and CR, hence the above argument can be applied to N3[SENC, SLMAC]. This implies that any attack must contain different tags while yielding the same ciphertext. By construction of N3[SENC, SLMAC] the tag is computed as output of the function SLMAC, which is based on a random function. Therefore, it is highly unlikely that a pre-chosen tag can be hit by choosing the inputs fittingly. Furthermore, by construction of SENC, the key stream consists of several  $r$ -bit blocks generated by SPRG on input  $\text{SLFUNC}(K_e, N)$ . In particular, as the tag is inputted into SENC as the last message block, one needs to find  $(K, N), (\bar{K}, \bar{N})$  such that the corresponding last  $r$ -bits blocks  $(R, \bar{R})$  from the key stream satisfy  $R \oplus T = \bar{R} \oplus \bar{T}$ . However, both the computation of the key streams  $R, \bar{R}$  and the tags  $T, \bar{T}$  depend on the respective nonce and message. Thus after either the tag or the key stream is computed, the choices for the remaining one are heavily restricted and hence  $R \oplus T = \bar{R} \oplus \bar{T}$  is unlikely to occur.

For the rest of this analysis we drop the assumptions wCR and CR and play through a number of general attack strategies.

1. A naive strategy in CMT attacks is to generate the target ciphertext  $C$  from randomly sampled  $(K, N, A, M)$  and then look for a second different input tuple  $(\bar{K}, \bar{N}, \bar{A}, \bar{M})$  that encrypts to  $C$  as well. This approach is what [MLGR23] introduced as context discovery and several of the attacks against the NIST LWC finalists are of this form [KSW23]. Below, we show that context discovery attacks on the AE scheme boil down to similar attacks on the underlying encryption scheme and MAC.
  - (a) One possibility for an attack is to invert the ciphertext under  $\bar{K}_e$  and  $\bar{N}$ , which yields  $\bar{M} \parallel \bar{T}$ . This leaves the adversary with the task of finding  $\bar{K}_m$  and  $\bar{A}$  such that  $\text{MAC}(\bar{K}_m, \bar{N}, \bar{A}, \bar{M}) = \bar{T}$ . Following the context discovery notions defined for AE schemes [MLGR23], we refer to the above as a  $(K_m, A)$ -discovery attack. While there are MACs for which this is not possible (e.g. SLMAC), there are also ones that allow this attack (e.g., the MACs underlying ELEPHANT [BCDM21] and MINALPHER [STA<sup>+</sup>15] as well as the MAC CHASKEY [MMV<sup>+</sup>14]).
  - (b) Furthermore, one can try to reach the target ciphertext by first computing a second tag  $\bar{T}$  using  $(\bar{K}_m, \bar{N}, \bar{A}, \bar{M})$ . This leaves the task of finding  $\bar{K}_e$  such that  $\text{SE.Enc}(\bar{K}_e, \bar{N}, \bar{M} \parallel \bar{T}) = C$ . We call this a  $K_e$ -discovery attack. At first glance, this might look like a key recovery attack, where the adversary gets message-ciphertext pairs. Note, however, that the adversary only needs to find a key that satisfies this property for one particular message-ciphertext pair. This differentiates it from a key recovery attack and leaves the possibility for secure scheme that allow for that—though we are not aware of such an example.
2. Another strategy is trying to produce a collision, i.e., instead of fixing a target ciphertext and trying to match it with the second tuple, one varies the inputs of SE simultaneously. As message, nonce, associated data, and MAC key already need to be chosen to compute the tag as input for the encryption scheme, the collision attack boils down to finding  $K_e$  and  $\bar{K}_e$  such that  $\text{SE.Enc}(K_e, N, M \parallel T) = \text{SE.Enc}(\bar{K}_e, \bar{N}, \bar{M} \parallel \bar{T})$ .

Overall, the results for N3 are of a heterogeneous nature: On the one hand, we identify a possible proof strategy for schemes with wCR and CR security, on the other hand we discuss a number of attacks. This indicates that—at least until further analysis is carried out—an individual treatment is necessary for each scheme. However, the above results might give a first indication to whether the scheme under consideration is committing.

### 3.2 Committing Security for Symmetric Encryption and MACs

Having established the positive results for N1, the question is whether there are any schemes that satisfy the required security notions wCR and CR. In this section we analyze the generic constructions  $\text{SE}[\mathcal{F}, \mathcal{G}]$  and  $\text{MAC}[\mathcal{H}, \mathcal{F}']$ . We show that the wCR security of  $\text{SE}[\mathcal{F}, \mathcal{G}]$  reduces to its underlying components  $\mathcal{F}$  and  $\mathcal{G}$ . Likewise, CR security of  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  reduces to  $\mathcal{H}$  and  $\mathcal{F}'$ . We first give the security notions required in this section, starting with the definition of collision resistance of a hash function and a PRG.

**Definition 4.** Let  $\mathcal{H}: \{0, 1\}^* \rightarrow \{0, 1\}^w$  be a hash function with output length  $w$ . For any adversary  $\mathcal{A}$ , its CR advantage is defined as

$$\text{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{A}) := \Pr[\mathcal{H}(X_1) = \mathcal{H}(X_2) \wedge X_1 \neq X_2 \mid (X_1, X_2) \leftarrow \mathcal{A}()].$$

**Definition 5.** Let  $\mathcal{G}: \{0, 1\}^\sigma \rightarrow \{0, 1\}^*$  be a pseudorandom generator with associated seed space  $\{0, 1\}^\sigma$ . For any adversary  $\mathcal{A}$ , its CR advantage is defined as

$$\text{Adv}_{\mathcal{G}}^{\text{CR}}(\mathcal{A}) := \Pr[\mathcal{G}(z_1) = \mathcal{G}(z_2) \wedge z_1 \neq z_2 \mid (z_1, z_2) \leftarrow \mathcal{A}()].$$

Game pbind	Game bind	Game wbind
$(K, X), (\bar{K}, \bar{X}) \leftarrow \mathcal{A}()$	$(K, X), (\bar{K}, \bar{X}) \leftarrow \mathcal{A}()$	$(K, X), (\bar{K}, \bar{X}) \leftarrow \mathcal{A}()$
<b>if</b> $(K, X) = (\bar{K}, \bar{X})$	<b>if</b> $(K, X) = (\bar{K}, \bar{X})$	<b>if</b> $K = \bar{K} \vee X \neq \bar{X}$
<b>return</b> 0	<b>return</b> 0	<b>return</b> 0
$Y \leftarrow F(K, X)$	$Y \leftarrow F(K, X)$	$Y \leftarrow F(K, X)$
$\bar{Y} \leftarrow F(\bar{K}, \bar{X})$	$\bar{Y} \leftarrow F(\bar{K}, \bar{X})$	$\bar{Y} \leftarrow F(\bar{K}, \bar{X})$
<b>return</b> $([Y]_k = [\bar{Y}]_k)$	<b>return</b> $(Y = \bar{Y})$	<b>return</b> $(Y = \bar{Y})$

Figure 7: Security games pbind, bind, and wbind.

*Remark 1.* An injective pseudorandom generator trivially is collision-resistant as no collision exists. Unlike hash functions, injectivity seems likely for PRGs as they map a short seed into a larger pseudorandom bit string. However, collision resistance is not implied by a secure PRG: take any secure PRG and hardcode two seeds  $z_1 \neq z_2$  to a fixed output  $Z$ .

Next, we define the so-called binding security of a keyed function. More precisely, we give three variants of it: pbind, bind, and wbind. The first, pbind, is taken from [BH22]—note that they call this notion bind since they do not make the same distinction as we do—and requires the adversary to find two key-input pairs for which the outputs “partially” agree, i.e., on their first  $k$  bits. The second, bind, is the the same except that the whole outputs have to agree. The third, wbind, bears similarities with wCR: the adversary needs to find *distinct* key and *identical* inputs for which the outputs agree.

**Definition 6.** Let  $F: \mathcal{K} \times \{0, 1\}^x \rightarrow \{0, 1\}^y$  be a function and the games pbind, bind, wbind be defined as in Figure 7. For any adversary  $\mathcal{A}$ , its pbind, bind, and wbind advantages are defined as

$$\mathbf{Adv}_F^X(\mathcal{A}) := \Pr[X(\mathcal{A}) \rightarrow 1],$$

where  $X \in \{\text{pbind}, \text{bind}, \text{wbind}\}$ .

These security notions can be ordered hierarchically with respect to their strength, namely  $\mathbf{Adv}_F^{\text{wbind}}(\mathcal{C}) \leq \mathbf{Adv}_F^{\text{bind}}(\mathcal{D}) \leq \mathbf{Adv}_F^{\text{pbind}}(\mathcal{B})$ . Details are given in Appendix B.6.

The theorem below shows that wCR security of the encryption scheme  $\text{SE}[\mathcal{F}, \mathcal{G}]$  reduces to binding security and collision resistance of  $\mathcal{F}$  and  $\mathcal{G}$ , respectively. The details are given in Appendix B.1.

**Theorem 3.** Let  $\mathcal{F}: \{0, 1\}^\kappa \times \{0, 1\}^\nu \rightarrow \{0, 1\}^\sigma$  be a function family and  $\mathcal{G}: \{0, 1\}^\sigma \rightarrow \{0, 1\}^*$  be a pseudorandom generator. Then for any wCR adversary  $\mathcal{A}$  against  $\text{SE}[\mathcal{F}, \mathcal{G}]$ , there exists a wbind adversary  $\mathcal{B}$  against  $\mathcal{F}$  and a CR adversary  $\mathcal{C}$  against  $\mathcal{G}$  such that

$$\mathbf{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{wCR}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{F}}^{\text{wbind}}(\mathcal{B}) + \mathbf{Adv}_{\mathcal{G}}^{\text{CR}}(\mathcal{C}).$$

The binding security of  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  follows from the collision resistance of  $\mathcal{H}$  and binding security of  $\mathcal{F}'$ . This is formalized in the following theorem, which is proven in Appendix B.2.

**Theorem 4.** Let  $\mathcal{H}: \{0, 1\}^* \rightarrow \{0, 1\}^w$  be a hash function and  $\mathcal{F}': \{0, 1\}^\kappa \times \{0, 1\}^w \rightarrow \{0, 1\}^\tau$  be a function family. Then for any CR adversary  $\mathcal{A}$  against  $\text{MAC}[\mathcal{H}, \mathcal{F}']$ , there exists a CR adversary  $\mathcal{B}$  against  $\mathcal{H}$  and a bind adversary  $\mathcal{C}$  against  $\mathcal{F}'$  such that

$$\mathbf{Adv}_{\text{MAC}[\mathcal{H}, \mathcal{F}']}^{\text{CR}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{B}) + \mathbf{Adv}_{\mathcal{F}'}^{\text{bind}}(\mathcal{C}).$$

The following composition theorem shows committing security of the N1 construction, when instantiating it with  $\text{SE}[\mathcal{F}, \mathcal{G}]$  and  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  which in turn are instantiated with  $\mathcal{F}$ ,  $\mathcal{G}$ ,  $\mathcal{H}$ , and  $\mathcal{F}'$ . It follows by combining the previous results and is proven in Appendix B.3.

**Theorem 5.** Let  $\mathcal{F}: \{0,1\}^\kappa \times \{0,1\}^\nu \rightarrow \{0,1\}^\sigma$ ,  $\mathcal{F}': \{0,1\}^\kappa \times \{0,1\}^w \rightarrow \{0,1\}^\tau$  be function families,  $\mathcal{G}: \{0,1\}^\sigma \rightarrow \{0,1\}^*$  be a pseudorandom generator, and  $\mathcal{H}: \{0,1\}^* \rightarrow \{0,1\}^w$  be a hash function. Then for any CMT adversary  $\mathcal{A}$  against  $\text{N1}[\text{SE}[\mathcal{F}, \mathcal{G}], \text{MAC}[\mathcal{H}, \mathcal{F}']]$ , there exists a wbind adversary  $\mathcal{B}$  against  $\mathcal{F}$ , a CR adversary  $\mathcal{C}$  against  $\mathcal{G}$ , a CR adversary  $\mathcal{D}$  against  $\mathcal{H}$  and a bind adversary  $\mathcal{E}$  against  $\mathcal{F}'$  such that

$$\text{Adv}_{\text{N1}[\text{SE}[\mathcal{F}, \mathcal{G}], \text{MAC}[\mathcal{H}, \mathcal{F}']]}^{\text{CMT}}(\mathcal{A}) \leq \text{Adv}_{\mathcal{F}}^{\text{wbind}}(\mathcal{B}) + \text{Adv}_{\mathcal{G}}^{\text{CR}}(\mathcal{C}) + \text{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{D}) + \text{Adv}_{\mathcal{F}'}^{\text{bind}}(\mathcal{E}).$$

### 3.3 An Instantiation from Sponges

Having established [Theorem 5](#), we now give concrete bounds for the sponge-based instantiations of  $\mathcal{F}$ ,  $\mathcal{G}$ ,  $\mathcal{H}$ , and  $\mathcal{F}'$ . Recall that  $n$  is the size of the sponge state while  $c$  is its capacity.

The theorem below bounds the pbind, bind, and wbind advantages of the sponge-based function SLFUNC. Its proof can be found in [Appendix B.4](#).

**Theorem 6.** Let SLFUNC be the sponge-based function as displayed in [Figure 3](#). Then for any adversary  $\mathcal{A}$  making  $q$  queries to  $\rho$ , it holds that

$$\text{Adv}_{\text{SLFUNC}}^{\text{pbind}}(\mathcal{A}) \leq \frac{q^2 - q}{2^{c+1}} + \frac{q^2 - q}{2^{k+1}} \quad \text{and} \quad \text{Adv}_{\text{SLFUNC}}^{\text{wbind}}(\mathcal{A}) \leq \frac{q^2 - q}{2^{n+1}}.$$

In particular, for  $k = n$  we have

$$\text{Adv}_{\text{SLFUNC}}^{\text{bind}}(\mathcal{A}) \leq \frac{q^2 - q}{2^{c+1}} + \frac{q^2 - q}{2^{n+1}}.$$

Next, we give a bound on the collision-resistance of SPRG, the proof is given in [Appendix B.5](#).

**Theorem 7.** Let SPRG:  $\{0,1\}^n \rightarrow \{0,1\}^m$  be the sponge-based pseudorandom generator as displayed in [Figure 3](#). Then for any adversary  $\mathcal{A}$  making  $q$  queries to  $\rho$ , it holds that

$$\text{Adv}_{\text{SPRG}}^{\text{CR}}(\mathcal{A}) \leq \frac{q^2 - q}{2^{m+1}}.$$

Collision resistance of the sponge-based hash function SvHASH has been shown in [\[DJS19\]](#). We recall this result below.

**Theorem 8** ([\[DJS19\]](#)). Let SvHASH be the hash function as displayed in [Figure 3](#) with output length  $w$ . Then for any adversary  $\mathcal{A}$  making  $q$  queries to  $\rho$ , it holds that

$$\text{Adv}_{\text{SvHASH}}^{\text{CR}}(\mathcal{A}) \leq \frac{q(q-1)}{2^{w+1}} + \frac{q(q+2)}{2^{c-1}}.$$

Using [Theorem 5](#) together with [Theorem 6](#), [Theorem 7](#), and [Theorem 8](#), we obtain the committing security of  $\text{N1}[\text{SLENC}, \text{SLMAC}]$ .

**Theorem 9.** Let SLFUNC, SPRG, and SvHASH be the function family, PRG, and hash function described in [Figure 3](#), respectively. Let further  $\text{N1}[\text{SLENC}, \text{SLMAC}]$  be the AE scheme constructed via N1 construction from SLENC and SLMAC, which in turn are constructed from SLFUNC, SPRG, and SvHASH as described in [Figure 15](#). Then for any CMT adversary  $\mathcal{A}$  against AE, making  $q$  queries to  $\rho$ , it holds that

$$\text{Adv}_{\text{N1}[\text{SLENC}, \text{SLMAC}]}^{\text{CMT}}(\mathcal{A}) \leq \frac{q^2 - q}{2^n} + \frac{q^2 - q}{2^{m+1}} + \frac{q^2 - q}{2^{w+1}} + \frac{q^2 + 2q}{2^{c-1}} + \frac{q^2 - q}{2^{c+1}}.$$

*Remark 2.* Note that, in order to obtain a reasonable bound, the length of the messages outputted by  $\mathcal{A}$  should satisfy  $m \geq \min\{w, c\}$ . Otherwise,  $\frac{q^2 - q}{2^{m+1}}$  becomes the dominant term, which can be trivial for very small  $m$ . However, this is a non-restrictive requirement, as real-world committing attacks [\[DGRW18\]](#) usually have this property.

## 4 Committing Security and Leakage Resilience

In this section, we develop a generic transformation that turns an arbitrary AE scheme into an AE scheme that is both leakage-resilient and committing. We start with the required background on leakage-resilient security notions in [Section 4.1](#) followed by the transformation in [Section 4.2](#).

### 4.1 Leakage Security Notions

For the leakage model, we use [\[BMOS17\]](#) which is based on [\[DP08\]](#), following the “Only Computation Leaks Information” assumption [\[MR04\]](#). In this model, the adversary obtains challenge oracles that do not leak—representing the goal of the adversary—and leakage oracles that do leak—representing the power of the adversary to obtain side-channel leakage. For the leakage oracles, the adversary can choose a leakage function from some predetermined set of leakage functions. Alongside the output of the functionality that the leakage oracle represents, the adversary receives the evaluation of the leakage function. In case the scheme is composed of different components, the leakage of the composed scheme is the composition of the leakage from the individual components, i.e., for a primitive  $C$  composed of  $A$  and  $B$ , the leakage set of  $C$  is  $\mathcal{L}_C = \mathcal{L}_A \times \mathcal{L}_B$  for  $\mathcal{L}_A$  and  $\mathcal{L}_B$  the leakage sets of  $A$  and  $B$ , respectively. An assumption that we are making is that comparison of values is leak-free, for instance, when a given tag is compared with the correct, recomputed tag. This assumption is made for SLAE and its generic construction FGHF’ [\[DJS19,KS20\]](#). Methods to achieve this are presented in [\[DM21\]](#).

Our target is LAE security as defined in [\[BMOS17\]](#). In the notion, the adversary gets four oracles: Two challenge oracles which either implement the real encryption and decryption or their idealized counterparts, i.e., outputting random ciphertext (for the encryption oracle) and rejecting any query (for the decryption oracle). In addition, the adversary gets two leakage oracles, one for encryption and one for decryption, which always implement the real algorithm. The leakage oracle takes a leakage function as additional input whose output models the side-channel leakage that the adversary receives. Classical security for AE schemes is obtained by discarding the leakage oracles for the adversary. The security notion is formalized below.

**Definition 7** (LAE Security). Let  $\text{AE} = (\text{Enc}, \text{Dec})$  be an authenticated encryption scheme with associated data and the game LAE be as defined in [Figure 8](#). For any nonce-respecting adversary  $\mathcal{A}$  that never forwards or repeats queries to or from the oracles  $\text{Enc}$  and  $\text{Dec}$  and only makes encryption and decryption queries containing leakage functions in the set  $\mathcal{L}_{\text{AE}}$ , describing the leakage sets for authenticated encryption, its corresponding LAE advantage is given by

$$\text{Adv}_{\text{AE}}^{\text{LAE}}(\mathcal{A}, \mathcal{L}_{\text{AE}}) := |\Pr[\mathcal{A}^{\text{LAE}} \rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}^{\text{LAE}} \rightarrow 1 \mid b = 0]|.$$

Barwell et al. [\[BMOS17\]](#) mention that their LAE notion can be slightly weakened by following the approach for AE security due to Shrimpton and Terashima [\[ST13\]](#). Here, a ciphertext consists of the actual ciphertext  $C$  along with some, so-called, *recovery information*  $R$ . One can think of this recovery information as the tag—in case of AEAD schemes via generic composition methods. Shrimpton and Terashima [\[ST13\]](#) require only  $C$  to be indistinguishable from random but not  $R$ . We adopt this approach and refer to the resulting notion as wLAE (with security game wLAE).

In addition to LAE/wLAE security, we need LPRF security which corresponds to the standard PRF security enhanced with leakage: The adversary gets a challenge oracle, which implements either the real function or a random function. On top of that, the adversary gets a leakage oracle, which always implements the real function that also returns leakage, based on the leakage function queried by the adversary. The definition of LPRF security is given below.

Game LAE	oracle $\text{Enc}(N, A, M)$	oracle $\text{Dec}(N, A, C)$
$b \leftarrow \$ \{0, 1\}$	$C \leftarrow \text{Enc}(K, N, A, M)$	<b>if</b> $b = 0$
$K \leftarrow \$ \mathcal{K}$	<b>if</b> $b = 0$	<b>return</b> $\perp$
$b' \leftarrow \mathcal{A}^{\text{Enc}, \text{LEnc}, \text{Dec}, \text{LDec}}()$	<b>return</b> $\overline{C} \leftarrow \$ \{0, 1\}^{ C }$	$M \leftarrow \text{Dec}(K, N, A, C)$
<b>return</b> $(b' = b)$	<b>return</b> $C$	<b>return</b> $M$
	<b>oracle</b> $\text{LEnc}(N, A, M, L)$	<b>oracle</b> $\text{LDec}(N, A, C, L)$
	$A \leftarrow L(K, N, A, M)$	$A \leftarrow L(K, N, A, C)$
	$C \leftarrow \text{Enc}(K, N, A, M)$	$M \leftarrow \text{Dec}(K, N, A, C)$
	<b>return</b> $(C, A)$	<b>return</b> $(M, A)$

Figure 8: Security game LAE.

**Definition 8** (LPRF Security). Let  $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  be a function family indexed by  $\mathcal{K}$  over the domain  $\mathcal{X}$  and let the game LPRF be as defined in Figure 9. For any adversary  $\mathcal{A}$  that never forwards or repeats queries to or from the oracle  $F$  and only queries leakage functions in the set  $\mathcal{L}_F$ , describing the leakage set for the function, its corresponding LPRF advantage is given by

$$\text{Adv}_F^{\text{LPRF}}(\mathcal{A}, \mathcal{L}_F) := |\Pr[\mathcal{A}^{\text{LPRF}} \Rightarrow 1 \mid b = 1] - \Pr[\mathcal{A}^{\text{LPRF}} \Rightarrow 1 \mid b = 0]|.$$

Game LPRF	oracle $F(X)$	oracle $\text{LF}(X, L)$
$b \leftarrow \$ \{0, 1\}$	<b>if</b> $b = 0$	$A \leftarrow L(K, X)$
$K \leftarrow \$ \mathcal{K}$	<b>return</b> $Y \leftarrow \$ \mathcal{Y}$	$Y \leftarrow F(K, X)$
$b' \leftarrow \mathcal{A}^{F, \text{LF}}()$	<b>else</b>	<b>return</b> $(Y, A)$
<b>return</b> $(b' = b)$	<b>return</b> $F(K, X)$	

Figure 9: Security game LPRF.

Finally, we require LUF security as defined in [DJS19]. This notion grants the adversary a challenge oracle **Guess**, where the adversary tries to predict the output of the function. In addition,  $\mathcal{A}$  receives an oracle  $F$  which provides the adversary the output of the function for queried inputs and an oracle **Lkg** which evaluates the function on an input chosen by the adversary but returns *only* the leakage. To avoid trivial wins, the adversary does not win if it predicts the output of an input queried to  $F$ , though it can obtain the leakage for an output it later tries to predict (meaning that queries can be repeated across **Guess** and **Lkg**). Below we give the formal definition.

**Definition 9** (LUF Security). Let  $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  be a function family over the domain  $\mathcal{X}$  and indexed by  $\mathcal{K}$ , and the game LUF be as defined in Figure 10. Then, for any adversary  $\mathcal{A}$ , querying only leakage functions in the set  $\mathcal{L}_F$ , its corresponding LUF advantage is given by

$$\text{Adv}_F^{\text{LUF}}(\mathcal{A}, \mathcal{L}_F) := \Pr[\text{LUF}(\mathcal{A}) \rightarrow 1].$$

## 4.2 A Generic Transformation

Our results so far show that one can construct committing authenticated encryption via generic composition, more precisely by following the Encrypt-and-MAC paradigm.



Game LUF	oracle $\text{Guess}(X, \bar{Y})$	oracle $F(X)$	oracle $\text{Lkg}(X, L)$
$\text{win} \leftarrow 0$	$Y \leftarrow F(K, X)$	$\mathcal{S} \leftarrow_{\cup} X$	$\Lambda \leftarrow L(K, X)$
$\mathcal{S} \leftarrow \emptyset$	if $X \notin \mathcal{S} \wedge Y = \bar{Y}$	$Y \leftarrow F(K, X)$	<b>return</b> $\Lambda$
$K \leftarrow_{\$} \mathcal{K}$	$\text{win} \leftarrow 1$	<b>return</b> $Y$	
$\mathcal{A}^{\text{Guess}, F, \text{Lkg}}()$	<b>return</b> $(Y = \bar{Y})$		
<b>return</b> win			

Figure 10: Security game LUF.

However, our analysis also revealed that the Encrypt-then-MAC approach can never yield committing authenticated encryption. This presents a stark contrast to the results that are known in the realm of leakage-resilient AE schemes via generic composition: Barwell et al. [BMOS17] showed that only the Encrypt-then-MAC paradigm yields leakage-resilient AE schemes whereas the other two approaches suffer from inherent weaknesses. The key problem is that in both Encrypt-and-MAC and MAC-then-Encrypt, adversaries can obtain decryption leakage even for invalid ciphertexts, as the ciphertext needs to be decrypted in order to be validated. In contrast, Encrypt-then-MAC schemes can validate the ciphertext *before* decrypting it, thereby guaranteeing that decryption can leak only for valid ciphertexts.

In total, our results on the committing security of the generic composition methods combined with the existing ones on their leakage-resilience, expose that we cannot build AE schemes with both properties from generic composition. On our quest for committing and leakage-resilient AE schemes, we hence turn towards transformations that achieve the security notions.

To the best of our knowledge, there are two transformations that achieve CMT security: the combination of UtC and HtE from [BH22] and the CTX construction from [CR22]. Other transformations PaddingZeros, KeyHashing, and CAU-C1 [BH22] only achieve weaker forms of committing security.

Both constructions are shown to achieve CMT security while maintaining the security of the underlying AE scheme.<sup>5</sup> Here, however, security is to be understood only with respect to classical AE security. Thus, when applying the transformation to any AE scheme that achieves LAE security, it is unclear whether the resulting scheme still achieves LAE security. Note that both transformations feed the key  $K$  as input into a hash function  $H$ . If the hash function is not leakage-resilient, the adversary might be able to learn the key which would render the scheme insecure. In particular, we are not aware of any leakage-resilient hash function. While sponge constructions can achieve leakage resilience by reducing the rate to a minimum [DEM<sup>+</sup>17, DEM<sup>+</sup>20, DJS19], for hash functions such a change seems impractical.

This raises the question whether there are other transformations without this drawback. In the following we answer this question in the affirmative:

Our transform ETC (for Encrypt-then-Commit), turns an arbitrary AEAD scheme into an AEAD scheme that is both leakage-resilient—in the sense of wLAE—and achieves CMT security. This transformation is shown in Figure 11 and visualized in Figure 12. The session key  $K'$  is computed by feeding key and nonce into a keyed function  $F$ , while the commitment  $P$  stems from hashing nonce, associated data, and ciphertext, and then taking the result and the key as input to  $F$ —note that we deploy a domain separation for these two invocations of  $F$ . While this transform requires an arbitrary-length ciphertext to be hashed, the ciphertext expansion is the same as for the existing transform UTC\*. Lastly, observe that the difference between LAE and wLAE security for ETC lies in the

<sup>5</sup>Note that [BH22] also provides a transformation RtC that maintains nonce-misuse resistance.

$\text{ETC}[\text{H}, \text{F}, \text{AE}].\text{Enc}(K, N, A, M)$	$\text{ETC}[\text{H}, \text{F}, \text{AE}].\text{Dec}(K, N, A, (P^*, C))$
$K' \leftarrow \text{F}(K, N \parallel 0)$	$P \leftarrow \text{F}(K, \text{H}(N, A, C) \parallel 1)$
$C \leftarrow \text{AE}.\text{Enc}(K', N, A, M)$	<b>if</b> $P \neq P^*$
$P \leftarrow \text{F}(K, \text{H}(N, A, C) \parallel 1)$	<b>return</b> $\perp$
<b>return</b> $(P, C)$	$K' \leftarrow \text{F}(K, N \parallel 0)$
	$M \leftarrow \text{AE}.\text{Dec}(K', N, A, C)$
	<b>return</b> $M$

Figure 11: Transformation ETC.

commitment  $P$ , which is appended to the ciphertext: for LAE, it has to be unpredictable and indistinguishable from random whereas for wLAE it merely has to be unpredictable.

**Remark and Errata.** The proceedings version of this paper [SW24] contains another transformation called UTC\* which is shown in Figure 13. It turns out that this transform—while achieving CMT security—does not yield leakage resilient AE schemes as we claimed in [SW24]. More precisely, the following forgery attack is possible: Let AE be an AEAD scheme, F be a function, and H be a hash function. Further, let  $\text{UTC}^*[\text{H}, \text{F}, \text{AE}]$  be the resulting AEAD scheme after applying the UTC\* transform. Consider the set of leakage functions  $L = L_H \times L_F \times L_{\text{AE}}$ , where  $L_{\text{AE}}(K, \cdot, \cdot, \cdot) = K$ , i.e., leakage of the AEAD scheme leaks the entire key. For the UTC\*-transformed scheme, the leakage function leaks the session key  $K'$ , which is derived from the context as  $(P, K') \leftarrow \text{F}(K, \text{H}(N, A))$ . Consider an adversary  $\mathcal{A}$  that picks an arbitrary tuple  $(N, A, M)$  of nonce, associated data, and message, and queries it to its leakage encryption oracle. The response will be some ciphertext  $(P, C)$  and the leakage  $L_{\text{AE}}$  provides the adversary with the session key  $K' = \text{F}(K, \text{H}(N, A))$ . Using  $K'$ ,  $\mathcal{A}$  can pick an arbitrary message  $\bar{M}$  (this one needs to be different than the queried message  $M$ ) and compute  $\bar{C} \leftarrow \text{AE}.\text{Enc}(K', N, A, \bar{M})$ . By outputting  $(P, \bar{C})$ ,  $\mathcal{A}$  has successfully forged a ciphertext. The problem of the transform is that an adversary which obtains  $K'$  via leakage can forge arbitrary ciphertext for  $K'$ . Since  $K'$  is determined by the nonce and the associated data, this allows to forge any ciphertext for the context  $(K, N, A)$ .

The claims regarding confidentiality in the presence of leakage are unaffected. It seems plausible that requiring LAE security of the underlying AEAD scheme might be sufficient as a fix—the aforementioned attack would definitely be ruled out by that. However, the transform would then not turn an arbitrary AEAD scheme into one that is both leakage-resilient *and* committing but rather turn a leakage-resilient one into one that is committing while *maintaining* its leakage resilience. Further, we do not know how to prove it since there are certain query combinations that result in forbidden queries for the reduction. That is why we are pursuing an alternative route, which allows us to provide a transform that does *not* require leakage resilience of the underlying scheme—as was the initial goal. For this, we make two main changes: Firstly, we include the ciphertext in the computation of the commitment  $P$  and, secondly, we consider a slightly weakened form of leakage resilience wLAE as described in Section 4.1. These changes allow us to circumvent the forbidden queries and give a proof for our new transform ETC.

The theorem below establishes the security of the ETC transform.

**Theorem 10.** *Let AE be an authenticated encryption scheme,  $\text{F}: \mathcal{K} \times \{0, 1\}^x \rightarrow \{0, 1\}^k \times \mathcal{K}$  be a function, and  $\text{H}: \{0, 1\}^* \rightarrow \{0, 1\}^x$  be a hash function with associated leakage sets  $\mathcal{L}_{\text{AE}}$ ,  $\mathcal{L}_{\text{F}}^0 \cup \mathcal{L}_{\text{F}}^1$ , and  $\mathcal{L}_{\text{H}}$ , respectively.<sup>6</sup> Let further  $\text{ETC}[\text{H}, \text{F}, \text{AE}]$  be the authenticated encryption scheme resulting from H, F, and AE via the transformation shown in Figure 11 with*

<sup>6</sup>Note that  $\mathcal{L}_{\text{F}}^0$  and  $\mathcal{L}_{\text{F}}^1$  refer to the leakage for the first and second application of F in ETC, respectively.




Figure 12: Visualization of the ETC transform.

$\text{UTC}^*[\text{H}, \text{F}, \text{AE}].\text{Enc}(K, N, A, M)$	$\text{UTC}^*[\text{H}, \text{F}, \text{AE}].\text{Dec}(K, N, A, (P^*, C, T))$
$X \leftarrow \text{H}(N, A)$	$X \leftarrow \text{H}(N, A)$
$(P, K') \leftarrow \text{F}(K, X)$	$(P, K') \leftarrow \text{F}(K, X)$
$(C, T) \leftarrow \text{Enc}(K', N, A, M)$	<b>if</b> $P^* \neq P$
<b>return</b> $(P, C, T)$	<b>return</b> $\perp$
	<b>return</b> $\text{Dec}(K', N, A, (C, T))$

Figure 13: Our modified transform  $\text{UTC}^*$ .

associated leakage set  $\mathcal{L} = \mathcal{L}_H \times \mathcal{L}_F^0 \times \mathcal{L}_F^1 \times \mathcal{L}_{\text{AE}}$ . Then for any adversary  $\mathcal{A}$  there exist adversaries  $\mathcal{B}$  and  $\mathcal{C}$ , such that

$$\mathbf{Adv}_{\text{ETC}[\text{H}, \text{F}, \text{AE}]}^{\text{CMT}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{bind}}(\mathcal{B}) + \mathbf{Adv}_H^{\text{CR}}(\mathcal{C}),$$

and for any nonce-respecting adversary  $\mathcal{A}$ , making  $q$  queries to its challenge oracles, there exist adversaries  $\mathcal{A}_i$  for  $i \in \{0, \dots, 3\}$ , such that

$$\begin{aligned} \mathbf{Adv}_{\text{ETC}[\text{H}, \text{F}, \text{AE}]}^{\text{wLAE}}(\mathcal{A}, \mathcal{L}) &\leq 2\mathbf{Adv}_H^{\text{CR}}(\mathcal{A}_0) + 2\mathbf{Adv}_F^{\text{LUF}}(\mathcal{A}_1, \mathcal{L}_F^1) \\ &\quad + \mathbf{Adv}_F^{\text{LPRF}}(\mathcal{A}_2, \mathcal{L}_F^0) + q\mathbf{Adv}_{\text{AE}}^{\text{AE}}(\mathcal{A}_3). \end{aligned}$$

*Proof.* We start by showing the first part of the statement. Let  $\mathbf{G}$  be the CMT game instantiated with the scheme  $\text{ETC}[\text{H}, \text{F}, \text{AE}]$  that is composed of the hash function  $\text{H}$ , the function  $\text{F}$ , and the authenticated encryption scheme  $\text{AE}$ . Let further  $\tilde{\mathbf{G}}$  be like  $\mathbf{G}$ , except that  $\mathcal{A}$  loses if its outputs constitute a collision in the hash function  $\text{H}$ . By a game hopping argument, we have

$$\begin{aligned} \mathbf{Adv}_{\text{ETC}[\text{H}, \text{F}, \text{AE}]}^{\text{CMT}}(\mathcal{A}_0) &= \Pr[\mathbf{G}(\mathcal{A}) \rightarrow 1] \\ &\leq |\Pr[\mathbf{G}(\mathcal{A}) \rightarrow 1] - \Pr[\tilde{\mathbf{G}}(\mathcal{A}) \rightarrow 1]| + \Pr[\tilde{\mathbf{G}}(\mathcal{A}) \rightarrow 1]. \end{aligned}$$

The games  $\mathbf{G}$  and  $\tilde{\mathbf{G}}$  are identical until a collision of  $\text{H}$ , occurs. We construct  $\mathcal{C}$  that runs  $\mathcal{A}$  to obtain  $(K, N, A, M)$ ,  $(\bar{K}, \bar{N}, \bar{A}, \bar{M})$  and simply outputs  $(N, A, C)$ ,  $(\bar{N}, \bar{A}, \bar{C})$  for  $C = \text{AE.Enc}(K, N, A, M)$  and  $\bar{C} = \text{AE.Enc}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . It then holds that

$$|\Pr[\mathbf{G}(\mathcal{A}) \rightarrow 1] - \Pr[\tilde{\mathbf{G}}(\mathcal{A}) \rightarrow 1]| \leq \mathbf{Adv}_H^{\text{CR}}(\mathcal{C}).$$

Let  $\mathcal{A}$  be an adversary winning the game  $\tilde{\mathbf{G}}$ . This means that  $\mathcal{A}$  outputs  $(K, N, A, M) \neq (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ , which fulfill

$$\begin{aligned} \text{ETC}[\mathbf{H}, \mathbf{F}, \mathbf{AE}].\text{Enc}(K, N, A, M) &= (\mathbf{F}(K, \mathbf{H}(N, A, C) \parallel 1), C) \\ &= (\mathbf{F}(\bar{K}, \mathbf{H}(\bar{N}, \bar{A}, \bar{C}) \parallel 1), \bar{C}) \\ &= \text{ETC}[\mathbf{H}, \mathbf{F}, \mathbf{AE}].\text{Enc}(\bar{K}, \bar{N}, \bar{A}, \bar{M}). \end{aligned}$$

We construct an adversary  $\mathcal{B}$ , which runs  $\mathcal{A}$ —obtaining  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ —and outputs  $(K, \mathbf{H}(N, A, C))$  and  $(\bar{K}, \mathbf{H}(\bar{N}, \bar{A}, \bar{C}))$ . Using the above, we get that  $\mathcal{B}$ 's output satisfies  $\mathbf{F}(K, \mathbf{H}(N, A, C) \parallel 1) = \mathbf{F}(\bar{K}, \mathbf{H}(\bar{N}, \bar{A}, \bar{C}) \parallel 1)$ . It remains to argue that its output is valid, i.e.,  $(K, \mathbf{H}(N, A, C) \parallel 1) \neq (\bar{K}, \mathbf{H}(\bar{N}, \bar{A}, \bar{C}) \parallel 1)$ . In case that  $\mathcal{A}$ 's output satisfies  $K \neq \bar{K}$ , this holds trivially. In the case that  $\mathcal{A}$ 's output satisfies  $K = \bar{K}$  it must hold that  $(N, A) \neq (\bar{N}, \bar{A})$  and  $\mathbf{H}(N, A, C) \neq \mathbf{H}(\bar{N}, \bar{A}, \bar{C})$ , as otherwise  $\mathcal{A}$  would not be an adversary winning the game  $\tilde{\mathbf{G}}$ . Thus, in both cases, we have that  $\mathcal{B}$ 's output is valid and conclude with

$$\Pr[\tilde{\mathbf{G}}(\mathcal{A}) \rightarrow 1] = \mathbf{Adv}_{\mathbf{F}}^{\text{bind}}(\mathcal{B}).$$

Collecting the above finishes the first part of the proof.

Next, we prove the second part of the statement. The overall proof strategy follows by consecutive game hops, using the following games:

- $\mathbf{G}_0$ : game  $\mathbf{wLAE}$
- $\mathbf{G}_1$ :  $\text{Dec}$  is replaced by  $\perp$  and  $\text{LDec}$  is modified to reject all queries excepts the ones forwarded from  $\text{LEnc}$ ; this means the output will be  $\perp$  together with the leakage of verifying  $P$  via  $\mathbf{F}$
- $\mathbf{G}_2$ : the output  $K'$  of  $\mathbf{F}$  is replaced by random values—note that  $P$  is still computed via  $\mathbf{F}$  as before
- $\mathbf{G}_3$ :  $\text{Enc}$  is changed to output uniform ciphertexts  $C$
- $\mathbf{G}_4$ :  $\text{LDec}$  is changed back to no longer reject non-forwarded queries (note that this corresponds to the ideal version of  $\mathbf{LAE}$ )

To deal with hash collisions, we consider variants of these games, where the oracles reject queries (outputting  $\perp$ ) if there is a hash collision. We indicate these variants using a tilde, e.g.,  $\tilde{\mathbf{G}}_i$ . The exact sequence of game hops for the proof is thus

$$\mathbf{G}_0 \rightarrow \tilde{\mathbf{G}}_0 \rightarrow \tilde{\mathbf{G}}_1 \rightarrow \tilde{\mathbf{G}}_2 \rightarrow \tilde{\mathbf{G}}_3 \rightarrow \tilde{\mathbf{G}}_4 \rightarrow \mathbf{G}_4$$

meaning we first switch to the version where hash collisions are rejected ( $\mathbf{G}_0 \rightarrow \tilde{\mathbf{G}}_0$ ), then do the game hops according to the description above ( $\tilde{\mathbf{G}}_0 \rightarrow \tilde{\mathbf{G}}_1 \rightarrow \tilde{\mathbf{G}}_2 \rightarrow \tilde{\mathbf{G}}_3 \rightarrow \tilde{\mathbf{G}}_4$ ), and finally switch back to not reject queries leading to hash collisions ( $\tilde{\mathbf{G}}_4 \rightarrow \mathbf{G}_4$ ). In the following, we analyze the individual game hops.

**Game Hop  $\mathbf{G}_0 \rightarrow \tilde{\mathbf{G}}_0$ :**

We start by giving a bound for the first summand. The games  $\mathbf{G}_0$  and  $\tilde{\mathbf{G}}_0$  are identical until a collision of  $\mathbf{H}$  occurs. We construct  $\mathcal{A}_0$  that outputs the tuples  $(N, A, C), (\bar{N}, \bar{A}, \bar{C})$  by  $\mathcal{A}$  that yielded a hash collision. It then holds that

$$|\Pr[\mathcal{A}^{\mathbf{G}_0} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{\mathbf{G}}_0} \rightarrow 1]| \leq \mathbf{Adv}_{\mathbf{H}}^{\text{CR}}(\mathcal{A}_0).$$

Note that in the following we do not describe the check for hash collisions, in which case the adversaries simply reject a query; we thus can assume that the hash values in the next hops are always unique.

**Game Hop  $\tilde{G}_0 \rightarrow \tilde{G}_1$ :**

To bound  $|\Pr[\mathcal{A}^{\tilde{G}_0} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{G}_1} \rightarrow 1]|$ , we construct the following LUF adversary  $\mathcal{A}_1$  against  $F$ . For queries  $(N, A, M)$  to **Enc** by  $\mathcal{A}$ , the adversary  $\mathcal{A}_1$  queries  $N \parallel 0$  to  $F$ , resulting in  $K'$ . This key is used to compute the ciphertext  $C$  locally. Additionally,  $H(N, A, C) \parallel 1$  is queried to  $F$ , which returns the commitment  $P$ . Then  $(P, C)$  is returned to  $\mathcal{A}$ . For queries  $(N, A, M, (L_H, L_F^0, L_F^1, L_{AE}))$  to **LEnc**, we repeat these steps with the difference that  $N \parallel 0$  and  $H(N, A, C) \parallel 1$  are queried to  $F$  and **Lkg**. This latter gives the leakage for the two applications of the PRF  $(A_F^0, A_F^1)$ ; the leakage of the AE scheme  $(A_{AE})$  and the hash function  $(A_H)$  are computed locally by  $\mathcal{A}_1$ , which then returns  $((P, C), (A_H, A_F^0, A_F^1, A_{AE}))$ . For queries  $(N, A, (P, C))$  to **Dec**,  $\mathcal{A}_1$  queries  $(H(N, A, C) \parallel 1, P)$  to **Guess**, which outputs either 0 or 1. If the output is 0,  $\mathcal{A}_1$  returns  $\perp$  to  $\mathcal{A}$ ; if it is 1,  $\mathcal{A}_1$  queries  $N \parallel 0$  to  $F$ , which outputs  $K'$ . Using the derived key  $K'$ ,  $\mathcal{A}_1$  can compute the AE decryption locally and returns  $M$  to  $\mathcal{A}$ . For queries  $(N, A, (P, C), (L_H, L_F, L_{AE}))$  to **LDec**,  $\mathcal{A}_1$  returns  $(\perp, (\perp, \perp, A_F^1, \perp))$  if the query was not forwarded from **LEnc**, where  $A_F^1$  is obtained by querying  $H(N, A, C) \parallel 1$  to **Lkg**. Otherwise, since the query is forwarded from **LEnc**,  $\mathcal{A}$  can look up the leakage for the two evaluations of  $F$  and the hash function from the corresponding **LEnc**. It then computes the leakage of the evaluation of AE together with  $M$  locally and returns  $(M, (A_H, A_F^0, A_F^1, A_{AE}))$  to  $\mathcal{A}_1$ .

It is left to show, that  $\mathcal{A}$  being able to differentiate  $\tilde{G}_1$  and  $\tilde{G}_0$  implies that  $\mathcal{A}_1$  wins the game LUF. In order to differentiate **LDec** from  $\tilde{G}_0$  and  $\tilde{G}_1$ ,  $\mathcal{A}$  needs to make a query  $(N, A, (P, C))$ , which is *not* forwarded from **LEnc** but for which **ETC.Dec** does not return  $\perp$ —in particular this implies that  $F(K, H(N, A, C) \parallel 1) = P$  has to hold. For a query of this form, **LDec** in  $\tilde{G}_1$  returns  $(\perp, (\perp, \perp, A_F^1, \perp))$ , while **LDec** in  $\tilde{G}_0$  does not. In particular,  $\mathcal{A}_1$  can then win the game LUF by querying  $(H(N, A, C) \parallel 1, P)$  to **Guess**. Note that  $\mathcal{A}_1$  does not win the game LUF if  $H(N, A, C) \parallel 1$  was queried to  $F$  before. However, we can exclude this case: Recall that a query of  $H(N, A, C) \parallel 1$  to  $F$  could have only happened after a query of  $\mathcal{A}$  to **LEnc** or **Enc**.

Assume  $(N, A, M)$  was queried to **Enc** or **LEnc** with output  $(\bar{P}, C)$ .<sup>7</sup> This means that  $H(N, A, C) \parallel 1$  was queried to  $F$ . Note that, since we assume that  $N, A, (P, C)$  was not forwarded from **Enc/LEnc**,  $\bar{P} \neq P$  must hold. However, then  $F(K, H(N, A, C) \parallel 1) = \bar{P} \neq P$  contradicts our assumption.

Similarly, to distinguish **Dec** from  $\perp$ ,  $\mathcal{A}$  needs to make a query  $(N, A, (P, C))$ —which was not forwarded from any oracle—and for which  $F(K, H(N, A, C) \parallel 1) = P$ , as otherwise **Dec** simply returns  $\perp$ . Just as discussed above for **LDec**,  $\mathcal{A}_1$  then wins game LUF by querying  $(H(N, A, C) \parallel 1, P)$  to **Guess**.

In total we have shown

$$|\Pr[\mathcal{A}^{\tilde{G}_0} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{G}_1} \rightarrow 1]| \leq \mathbf{Adv}_F^{\text{LUF}}(\mathcal{A}_1).$$

Note that for the following hops, we do not describe how queries to **Dec** are handled, as **Dec** can easily be simulated by always returning  $\perp$ .

**Game Hop  $\tilde{G}_1 \rightarrow \tilde{G}_2$ :**

To bound  $|\Pr[\mathcal{A}^{\tilde{G}_1} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{G}_2} \rightarrow 1]|$ , we construct the following LPRF adversary  $\mathcal{A}_2$  against  $F$ . For queries  $(N, A, M)$  to **Enc** by  $\mathcal{A}$ , the adversary  $\mathcal{A}_2$  queries  $N \parallel 0$  to  $F$ , resulting in a session key  $K'$ . With this key,  $\text{AE.Enc}(K', N, A, M)$  is then computed locally and  $H(N, A, C) \parallel 1$  is queried to  $F$ , which returns the commitment  $P$ . Then,  $(P, C)$  is returned to  $\mathcal{A}$ . Regarding queries to the leakage oracles, we proceed as

<sup>7</sup>We drop the leakage functions and returned leakage here, as our argument is independent of them.

follows: For queries to **LEnc**, we use **LF** to obtain the leakage of **F** applications and compute the leakage of the other components locally. For queries to **LDec**, we return  $(\perp, (\perp, \perp, A_F^1, \perp))$ , where  $A_F^1$  is obtained from **LF** on input  $H(N, A, C) \parallel 1$ , for all queries that are not forwarded from **LEnc**, and for the ones that are, we proceed as for **LEnc**, i.e., obtain leakage for **F** from **LF** and simulate the rest locally. Whenever  $\mathcal{A}$  outputs a bit,  $\mathcal{A}_2$  outputs the same.

Note that the LPRF adversary  $\mathcal{A}_2$  makes no forbidden queries, due to  $\mathcal{A}$  being nonce-respecting and the fact that queries cannot be forwarded between **Enc** and **LDec**. Further,  $\mathcal{A}_2$  perfectly simulates  $\tilde{G}_1$  or  $\tilde{G}_2$  for  $\mathcal{A}$ —depending on the value of the secret bit in the LPRF game. Since the games only difference is whether **F** or a random function is used,  $\mathcal{A}_2$  wins the game LPRF if  $\mathcal{A}$  can distinguish  $\tilde{G}_1$  and  $\tilde{G}_2$ . Thus, we have shown

$$|\Pr[\mathcal{A}^{\tilde{G}_1} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{G}_2} \rightarrow 1]| = \mathbf{Adv}_H^{\text{LPRF}}(\mathcal{A}_2).$$

### Game Hop $\tilde{G}_2 \rightarrow \tilde{G}_3$ :

Next, we bound the game hop between  $G_2$  and  $G_3$ . Due to the fact that **ETC** uses different session keys—computed from the context—for the underlying AE scheme **AE**, we cannot simply reduce to the security of **AE**. To accommodate for that, we do a hybrid argument over the distinct session keys.<sup>8</sup> For this, we define a sequence of hybrid games  $H_0, \dots, H_q$ , for  $q$  the number of nonces that  $\mathcal{A}$  queries to **Enc**,<sup>9</sup> as follows: In  $H_i$ , the first  $i$  distinct nonces are answered as in  $\tilde{G}_3$ , while the remaining queries are handled as in  $\tilde{G}_2$ . Observe that  $H_0 = \tilde{G}_2$  and  $H_q = \tilde{G}_3$ , which yields

$$\begin{aligned} |\Pr[\mathcal{A}_1^{\tilde{G}_2} \rightarrow 1] - \Pr[\mathcal{A}_1^{\tilde{G}_3} \rightarrow 1]| &= |\Pr[\mathcal{A}_1^{H_0} \rightarrow 1] - \Pr[\mathcal{A}_1^{H_q} \rightarrow 1]| \\ &\leq \sum_{i=1}^q |\Pr[\mathcal{A}_1^{H_{i-1}} \rightarrow 1] - \Pr[\mathcal{A}_1^{H_i} \rightarrow 1]|. \end{aligned}$$

For each  $i = 1, \dots, q$ , we construct an AE adversary  $\mathcal{D}_i$  to bound the game hop between  $H_{i-1}$  and  $H_i$ . First, the adversary  $\mathcal{D}_i$  samples a key  $K$  to locally simulate the oracles **LEnc**, and **LDec** for  $\mathcal{A}$ .

For  $(N_1, \dots, N_q)$  the  $q$  distinct nonces,  $\mathcal{D}_i$  proceeds as follows: Consider the  $j$ -th query to be a query to **Enc** of the form  $(N_j, A_j, M_j)$ —this is wlog as any query to **Dec** simply returns  $\perp$ . If  $j < i$ ,  $\mathcal{D}_i$  samples  $C_j$  randomly, computes  $P_j \leftarrow F_K(H(N_j, A_j, C_j) \parallel 1)$  and sends it to  $\mathcal{A}$ . If  $j = i$ ,  $\mathcal{D}_i$  invokes its own encryption oracle on  $(N_i, A_i, M_i)$  to obtain  $C_i$ , which is sent to  $\mathcal{A}_1$  together with  $P_i \leftarrow F_K(H(N, A, C) \parallel 1)$ . If  $j > i$ ,  $K'_j$  is sampled randomly,  $\mathbf{AE.Enc}(K'_j, N_j, A_j, M_j)$  is computed locally by  $\mathcal{D}_i$  as is  $P_j \leftarrow F_K(H(N, A, C) \parallel 1)$ . Then  $(P_j, C_j)$  is returned to  $\mathcal{A}_1$ .

For each query  $(N, A, M, L = (L_H, L_F^0, L_F^1, L_{AE}))$  to **LEnc**,  $\mathcal{D}_i$  computes  $K' \leftarrow F(K, N \parallel 0)$ ,  $C \leftarrow \mathbf{AE.Enc}(K', N, A, M)$ ,  $P \leftarrow F(K, H(N, A, C) \parallel 1)$ , and the leakage locally. It sends  $(P, C)$  and  $(A_H, A_F^0, A_F^1, A_{AE})$  to  $\mathcal{A}$ . Queries to **LDec** are handled analogously with  $C$  being used instead of  $M$ ,<sup>10</sup> but only if they are forwarded from **LEnc**—otherwise  $\perp$  is returned.

We observe that  $\mathcal{D}_i$  simulates  $H_{i-1}$  or  $H_i$  depending on the secret bit from the AE game, respectively. Note further that  $\mathcal{D}_i$  only makes a single encryption query. This follows from the fact that  $\mathcal{D}_i$  makes an encryption query for the  $i$ th nonce  $N_i$  queried by  $\mathcal{A}$  which is nonce-respecting and therefore does not repeat this nonce.

<sup>8</sup>In [BH22], the same problem occurs, though they resolve this using multi-user security [BT16].

<sup>9</sup>Note that by the nonce-respecting property, the nonces are distinct.

<sup>10</sup>Note, however, that  $P^* = P$  needs to be checked.

Furthermore, we observe that any encryption/decryption query by  $\mathcal{D}_i$  stems from a query by  $\mathcal{A}$ . If any of the queries by  $\mathcal{D}_i$  would be prohibited, we would immediately get that  $\mathcal{A}$  made queries that are forbidden. Thus, we can conclude

$$\begin{aligned} |\Pr[\mathcal{A}^{\mathcal{H}_{i-1}} \rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{H}_{i-1}} \rightarrow 1]| &= |\Pr[\mathcal{D}_i^{\text{AE}} \rightarrow 1 \mid b = 1] - \Pr[\mathcal{D}_i^{\text{AE}} \rightarrow 1 \mid b = 0]| \\ &= \mathbf{Adv}_{\text{AE}}^{\text{AE}}(\mathcal{D}_i). \end{aligned}$$

We define  $\mathcal{A}_3$  to be the adversary that picks  $i$  from  $\{1, \dots, q\}$  randomly and then acts like  $\mathcal{D}_i$ . By a standard hybrid argument, we get

$$\begin{aligned} |\Pr[\mathcal{A}^{\tilde{\mathcal{G}}_2} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{\mathcal{G}}_3} \rightarrow 1]| &\leq \sum_{i=1}^q |\Pr[\mathcal{A}^{\mathcal{H}_{i-1}} \rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{H}_{i-1}} \rightarrow 1]| \\ &\leq q \mathbf{Adv}_{\text{AE}}^{\text{AE}}(\mathcal{A}_3). \end{aligned}$$

**Game Hop  $\tilde{\mathcal{G}}_3 \rightarrow \tilde{\mathcal{G}}_4$ :**

This hop reverts some of the changes done when switching from  $\tilde{\mathcal{G}}_0$  to  $\tilde{\mathcal{G}}_1$ . The same argument as above then yields that there exists an adversary  $\mathcal{A}_1$  such that

$$|\Pr[\mathcal{A}^{\tilde{\mathcal{G}}_3} \rightarrow 1] - \Pr[\mathcal{A}^{\tilde{\mathcal{G}}_4} \rightarrow 1]| \leq \mathbf{Adv}_{\text{F}}^{\text{LUF}}(\mathcal{A}_1).$$

**Game Hop  $\tilde{\mathcal{G}}_4 \rightarrow \mathcal{G}_4$ :**

Note that adversary  $\mathcal{A}$  can not distinguish  $\tilde{\mathcal{G}}_4$  and  $\mathcal{G}_4$  except if there is a collision of the hash function. Then, the encryption oracle (or the decryption oracle, respectively) of  $\tilde{\mathcal{G}}_4$  outputs  $\perp$ , whereas for  $\mathcal{G}_4$  this is not the case. Then we can construct  $\mathcal{A}_0$  that outputs the tuples  $(N, A, C)$ ,  $(\bar{N}, \bar{A}, \bar{C})$  by  $\mathcal{A}$  that yielded a hash collision. Hence, we obtain

$$|\Pr[\mathcal{A}^{\tilde{\mathcal{G}}_4} \rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{G}_4} \rightarrow 1]| \leq \mathbf{Adv}_{\text{H}}^{\text{CR}}(\mathcal{A}_0).$$

Collecting everything concludes the proof.  $\square$

*Remark 3.* Note that the proof above does not rely on the unforgeability of the underlying AEAD scheme. It seems plausible that the transform might also work when the underlying component is just a symmetric encryption scheme and unforgeability then follows from the added function  $F$ . However, the details would still need to be checked.

## 5 Conclusion and Instantiation/Implementation Aspects

Our analysis reveals that the generic composition paradigms are not suitable for constructing AE schemes that are simultaneously committing and leakage-resilient. However, we identify another way to obtain such schemes by means of a generic transformation. The latter can be applied to any secure AE scheme, i.e., it does not require any committing or leakage resilience guarantees on the scheme to begin with. In particular, this allows to apply the ETC transformation to AE schemes built from generic composition without needing the underlying symmetric encryption scheme and MAC to fulfill binding or leakage resilience properties. While we just require standard AE security for the scheme that is to be transformed, the transformation uses a function  $F$  that needs to be binding, an LPRF, and an LUF. A possible instantiation for this is the sponge-based function `SLFUNC` (cf. Figure 3) which was shown to be an LPRF and an LUF in [DJS19] and binding in Theorem 6. Note that, in order to obtain a good bound on the LPRF/LUF security, the rate  $r$  needs to be very small and is typically set to 1. Further, for any practical instantiation,



one needs to choose a concrete non-invertible permutation for  $\rho$ . Following [DJS19], a candidate can be obtained using KECCAK-P: define  $\rho(x) = \text{KECCAK-P}(x) \oplus x$ , where the additional XOR ensures non-invertibility. An alternative candidate is the tagging algorithm of the leakage-resilient MAC given in [BMOS17]<sup>11</sup> which is an LPRF but has yet to be analyzed with respect to LUF and binding security. Bellare and Hoang [BH22] provide a construction of a binding function which they call “Counter-then-XOR” (CTX). The CTX construction, which relies on a block-cipher, is also an alternative but not yet analyzed in the leakage setting.

Coming back to SLAE, which was used as an example throughout this work, we can now provide a committing and leakage-resilient variant: This is achieved by applying the ETC transformation, using SVHASH and SLFUNC, to SLAE, where  $\rho(x) = \text{KECCAK-P}(x) \oplus x$  as described above. While SLAE is leakage-resilient to start with, due to a very small rate in the underlying function SLFUNC, this is not a necessary prerequisite for applying the ETC transformation. Thus, it is even possible to use SLAE with a bigger rate in SLFUNC as input to ETC, while the two additional SLFUNC instances used in the transformation deploy  $r = 1$  to ensure the leakage-resilient guarantees.

Finally, recall that we assume the comparison of values to be leak-free. This assumption is crucial for Theorem 10, when the given  $P^*$  is compared with the recomputed  $P$ . Any implementation of ETC needs to make sure that this assumption is not violated, e.g., by hardening this step with proper counter measures like masking [CJRR99]. Otherwise, the results can be misused, as was the case in [USS<sup>+</sup>20], which used the FGHF' construction [DJS19, KS20] but did not take into account the leak-free assumption—this was pointed out in [BMPS21]. Dobraunig and Mennink [DM21] also provide methods on how to achieve this assumption, which are often readily available in the implementation anyway.

## References

- [ADG<sup>+</sup>22] Ange Albertini, Thai Duong, Shay Gueron, Stefan Kölbl, Atul Luykx, and Sophie Schmieg. How to abuse and fix authenticated encryption without key commitment. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022*, pages 3291–3308. USENIX Association, August 2022.
- [BBC<sup>+</sup>20] Davide Bellizia, Olivier Bronchain, Gaëtan Cassiers, Vincent Grosso, Chun Guo, Charles Momin, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Mode-level vs. implementation-level physical security in symmetric cryptography - A practical guide through the leakage-resistance jungle. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 369–400. Springer, Cham, August 2020.
- [BCDM21] Tim Beyne, Yu Long Chen, Christoph Dobraunig, and Bart Mennink. Elephant. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
- [Ber23] Francesco Berti. Reconsidering generic composition: The modes A10, A11 and A12 are insecure. In Leonie Simpson and Mir Ali Rezazadeh Bae, editors, *ACISP 23*, volume 13915 of *LNCS*, pages 157–176. Springer, Cham, July 2023.
- [BGP<sup>+</sup>19] Francesco Berti, Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. TEDT: a leakage-resistant AEAD mode. *IACR TCHES*, 2020(1):256–320, 2019.

---

<sup>11</sup>This MAC is inspired by the one given in [MOSW15], which, while achieving unpredictability, is not pseudorandom in the leakage setting (LPRF). This makes it unsuitable for the UTC\* transformation.

- [BGPS21] Francesco Berti, Chun Guo, Thomas Peters, and François-Xavier Standaert. Efficient leakage-resilient MACs without idealized assumptions. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part II*, volume 13091 of *LNCS*, pages 95–123. Springer, Cham, December 2021.
- [BH22] Mihir Bellare and Viet Tung Hoang. Efficient schemes for committing authenticated encryption. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 845–875. Springer, Cham, May / June 2022.
- [BMOS17] Guy Barwell, Daniel P. Martin, Elisabeth Oswald, and Martijn Stam. Authenticated encryption in the face of protocol and side channel leakage. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 693–723. Springer, Cham, December 2017.
- [BMPS21] Olivier Bronchain, Charles Momin, Thomas Peters, and François-Xavier Standaert. Improved leakage-resistant authenticated encryption based on hardware AES coprocessors. *IACR TCHES*, 2021(3):641–676, 2021.
- [BN00] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, Berlin, Heidelberg, December 2000.
- [BPP18] Francesco Berti, Olivier Pereira, and Thomas Peters. Reconsidering generic composition: The tag-then-encrypt case. In Debrup Chakraborty and Tetsu Iwata, editors, *INDOCRYPT 2018*, volume 11356 of *LNCS*, pages 70–90. Springer, Cham, December 2018.
- [BPPS17] Francesco Berti, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. On leakage-resilient authenticated encryption with decryption leakages. *IACR Trans. Symm. Cryptol.*, 2017(3):271–293, 2017.
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Berlin, Heidelberg, May / June 2006.
- [BT16] Mihir Bellare and Björn Tackmann. The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 247–276. Springer, Berlin, Heidelberg, August 2016.
- [CFG<sup>+</sup>23] Yu Long Chen, Antonio Flórez-Gutiérrez, Akiko Inoue, Ryoma Ito, Tetsu Iwata, Kazuhiko Minematsu, Nicky Mouha, Yusuke Naito, Ferdinand Sibleyras, and Yosuke Todo. Key committing security of AEZ and more. In *ToSC 2023*, 2023.
- [CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Michael J. Wiener, editor, *CRYPTO’99*, volume 1666 of *LNCS*, pages 398–412. Springer, Berlin, Heidelberg, August 1999.
- [CR19] John Chan and Phillip Rogaway. Anonymous AE. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 183–208. Springer, Cham, December 2019.

- [CR22] John Chan and Phillip Rogaway. On committing authenticated-encryption. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *ESORICS 2022, Part II*, volume 13555 of *LNCS*, pages 275–294. Springer, Cham, September 2022.
- [DEM<sup>+</sup>17] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, and Thomas Unterluggauer. ISAP – towards side-channel secure authenticated encryption. *IACR Trans. Symm. Cryptol.*, 2017(1):80–105, 2017.
- [DEM<sup>+</sup>20] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, Bart Mennink, Robert Primas, and Thomas Unterluggauer. ISAP v2.0. *IACR Trans. Symm. Cryptol.*, 2020(S1):390–416, 2020.
- [DFG23] Jean Paul Degabriele, Marc Fischlin, and Jérôme Govinden. The indistinguishability of the duplex and its practical applications. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VIII*, volume 14445 of *LNCS*, pages 237–269. Springer, Singapore, December 2023.
- [DGRW18] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast message franking: From invisible salamanders to encryption. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 155–186. Springer, Cham, August 2018.
- [DJS19] Jean Paul Degabriele, Christian Janson, and Patrick Struck. Sponges resist leakage: The case of authenticated encryption. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 209–240. Springer, Cham, December 2019.
- [DM19] Christoph Dobraunig and Bart Mennink. Leakage resilience of the duplex construction. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 225–255. Springer, Cham, December 2019.
- [DM21] Christoph Dobraunig and Bart Mennink. Leakage resilient value comparison with application to message authentication. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 377–407. Springer, Cham, October 2021.
- [DMVA23] Joan Daemen, Silvia Mella, and Gilles Van Assche. Committing authenticated encryption based on SHAKE. *IACR Cryptol. ePrint Arch.*, 2023:1494, 2023.
- [DP08] Stefan Dziembowski and Krzysztof Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pages 293–302. IEEE Computer Society Press, October 2008.
- [FKOS22] Sebastian Faust, Juliane Krämer, Maximilian Ortl, and Patrick Struck. On the related-key attack security of authenticated encryption schemes. In Clemente Galdi and Stanislaw Jarecki, editors, *SCN 22*, volume 13409 of *LNCS*, pages 362–386. Springer, Cham, September 2022.
- [GPPS20] Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Towards low-energy leakage-resistant AE from the duplex sponge. *IACR Trans. Symm. Cryptol.*, 2020(1):6–42, 2020.
- [IKM<sup>+</sup>21] Tetsu Iwata, Mustafa Khairallah, Kazuhiko Minematsu, Thomas Peyrin, and Chun Guo. Romulus. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.

- [KS20] Juliane Krämer and Patrick Struck. Leakage-resilient authenticated encryption from leakage-resilient pseudorandom functions. In Guido Marco Bertoni and Francesco Regazzoni, editors, *COSADE 2020*, volume 12244 of *LNCS*, pages 315–337. Springer, Cham, April 2020.
- [KSW23] Juliane Krämer, Patrick Struck, and Maximiliane Weishäupl. Committing AE from sponges: Security analysis of the NIST LWC finalists. Cryptology ePrint Archive, Report 2023/1525, 2023.
- [LGR21] Julia Len, Paul Grubbs, and Thomas Ristenpart. Partitioning oracle attacks. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 195–212. USENIX Association, August 2021.
- [MLGR23] Sanketh Menda, Julia Len, Paul Grubbs, and Thomas Ristenpart. Context discovery and commitment attacks - how to break CCM, EAX, SIV, and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part IV*, volume 14007 of *LNCS*, pages 379–407. Springer, Cham, April 2023.
- [MMV<sup>+</sup>14] Nicky Mouha, Bart Mennink, Anthony Van Herrewege, Dai Watanabe, Bart Preneel, and Ingrid Verbauwhede. Chaskey: An efficient MAC algorithm for 32-bit microcontrollers. In Antoine Joux and Amr M. Youssef, editors, *SAC 2014*, volume 8781 of *LNCS*, pages 306–323. Springer, Cham, August 2014.
- [MOSW15] Daniel P. Martin, Elisabeth Oswald, Martijn Stam, and Marcin Wójcik. A leakage resilient MAC. In Jens Groth, editor, *15th IMA International Conference on Cryptography and Coding*, volume 9496 of *LNCS*, pages 295–310. Springer, Cham, December 2015.
- [MR04] Silvio Micali and Leonid Reyzin. Physically observable cryptography (extended abstract). In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 278–296. Springer, Berlin, Heidelberg, February 2004.
- [NRS14] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering generic composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, Berlin, Heidelberg, May 2014.
- [ST13] Thomas Shrimpton and R. Seth Terashima. A modular framework for building variable-input-length tweakable ciphers. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 405–423. Springer, Berlin, Heidelberg, December 2013.
- [STA<sup>+</sup>15] Yu Sasaki, Yosuke Todo, Kazumaro Aoki, Yusuke Naito, Takeshi Sugawara, Yumiko Murakami, Mitsuru Matsui, and Shoichi Hirose. Minalpher v1.1. Technical report, Submission to the CAESAR Competition, 2015. Available at <https://competitions.cr.yp.to/round2/minalpherv11.pdf>.
- [SW24] Patrick Struck and Maximiliane Weishäupl. Constructing committing and leakage-resilient authenticated encryption. *IACR Trans. Symm. Cryptol.*, 2024(1):497–528, 2024.
- [USS<sup>+</sup>20] Florian Unterstein, Marc Schink, Thomas Schamberger, Lars Tebelmann, Manuel Ilg, and Johann Heyszl. Retrofitting leakage resilient authenticated encryption to microcontrollers. *IACR TCHES*, 2020(4):365–388, 2020.
- [VCS22] Corentin Verhamme, Gaëtan Cassiers, and François-Xavier Standaert. Analyzing the leakage resistance of the NIST’s lightweight crypto competition’s finalists. In *CARDIS 2022*, 2022.

## A Additional Background

The pseudocode for the N-schemes is given in Figure 14. Figure 15 provides the pseudocode of the sponge-based encryption scheme SLENC and MAC SLMAC in terms of their underlying components SLFUNC, SPRG, and SVHASH, whose pseudocodes are presented in Figure 16.

$\text{Enc}(K, N, A, M)$ in N1	$\text{Enc}(K, N, A, M)$ in N2	$\text{Enc}(K, N, A, M)$ in N3
$(K_e, K_m) \leftarrow K$	$(K_e, K_m) \leftarrow K$	$(K_e, K_m) \leftarrow K$
$C \leftarrow \text{Enc}^S(K_e, N, M)$	$C_e \leftarrow \text{Enc}^S(K_e, N, M)$	$T \leftarrow \text{Tag}(K_m, N, A, M)$
$T \leftarrow \text{Tag}(K_m, N, A, M)$	$T \leftarrow \text{Tag}(K_m, N, A, C)$	$C \leftarrow \text{Enc}^S(K_e, N, M \parallel T)$
<b>return</b> $(C, T)$	<b>return</b> $(C, T)$	<b>return</b> $C$
$\text{Dec}(K, N, A, (C, T))$ in N1	$\text{Dec}(K, N, A, (C, T))$ in N2	$\text{Dec}(K, N, A, C)$ in N3
$(K_e, K_m) \leftarrow K$	$(K_e, K_m) \leftarrow K$	$(K_e, K_m) \leftarrow K$
$M \leftarrow \text{Dec}^S(K_e, N, C)$	<b>if</b> $\text{Ver}(K_m, N, A, C, T) = 0$	$M \parallel T \leftarrow \text{Dec}^S(K_e, N, C)$
<b>if</b> $\text{Ver}(K_m, N, A, M, T) = 0$	<b>return</b> $\perp$	<b>if</b> $\text{Ver}(K_m, N, A, M, T) = 0$
<b>return</b> $\perp$	$M \leftarrow \text{Dec}^S(K_e, N, C)$	<b>return</b> $\perp$
<b>return</b> $M$	<b>return</b> $M$	<b>return</b> $M$

Figure 14: Pseudocode of N1 (left), N2 (middle), and N3 (right) in terms of a symmetric encryption scheme  $(\text{Enc}^S, \text{Dec}^S)$  and a MAC  $(\text{Tag}, \text{Ver})$ .

## B Full Proofs

### B.1 Proof of Theorem 3

*Proof.* Let  $\mathcal{A}$  be a wCR adversary against  $\text{SE}[\mathcal{F}, \mathcal{G}]$  with output  $((K, N, M), (\bar{K}, \bar{N}, \bar{M}))$ . We define  $\mathbf{E}$  to be the event that  $\mathcal{F}(K, N) = \mathcal{F}(\bar{K}, \bar{N})$ .

Firstly, we construct a wbind adversary  $\mathcal{B}$  against  $\mathcal{F}$ . It runs  $((K, N, M), (\bar{K}, \bar{N}, \bar{M})) \leftarrow \mathcal{A}()$  and outputs  $((K, N), (\bar{K}, \bar{N}))$ . If  $\mathcal{A}$  is successful, we have both  $K \neq \bar{K}$  and  $(N, M) = (\bar{N}, \bar{M})$ . Next assume that additionally to  $\mathcal{A}$  being successful, event  $\mathbf{E}$  holds, that means  $\mathcal{F}(K, N) = \mathcal{F}(\bar{K}, \bar{N})$ . Then adversary  $\mathcal{B}$  is successful in game wbind, as  $K \neq \bar{K}$  and  $N = \bar{N}$ , but they map to the same element under  $\mathcal{F}$ . Hence we have shown

$$\Pr[\mathbf{E} \wedge \text{wCR}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{wbind}(\mathcal{B}) \rightarrow 1].$$

Secondly, we construct a CR adversary  $\mathcal{C}$  against  $\mathcal{G}$ . It runs  $((K, N, M), (\bar{K}, \bar{N}, \bar{M})) \leftarrow \mathcal{A}()$  and outputs  $(\mathcal{F}(K, N), \mathcal{F}(\bar{K}, \bar{N}))$ . If  $\mathcal{A}$  is successful, then  $K \neq \bar{K}$  and  $(N, M) = (\bar{N}, \bar{M})$ , and  $\text{SE}[\mathcal{F}, \mathcal{G}].\text{Enc}(K, N, M) = \text{SE}[\mathcal{F}, \mathcal{G}].\text{Enc}(\bar{K}, \bar{N}, \bar{M})$ . By definition of  $\text{SE}[\mathcal{F}, \mathcal{G}]$  the latter is equivalent to  $\mathcal{G}(\mathcal{F}(K, N)) \oplus M = \mathcal{G}(\mathcal{F}(\bar{K}, \bar{N})) \oplus \bar{M}$ . As  $M = \bar{M}$ , this implies that  $\mathcal{G}(\mathcal{F}(K, N)) = \mathcal{G}(\mathcal{F}(\bar{K}, \bar{N}))$ . Next assume that additionally to  $\mathcal{A}$  being successful, event  $\neg \mathbf{E}$  holds, that means  $\mathcal{F}(K, N) \neq \mathcal{F}(\bar{K}, \bar{N})$ . Then adversary  $\mathcal{C}$  is successful, as he has found a collision for  $\mathcal{G}$ . Thus we obtain

$$\Pr[\neg \mathbf{E} \wedge \text{wCR}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{CR}(\mathcal{C}) \rightarrow 1].$$

<b>SENC-Enc</b> ( $K_e, N, M$ )	<b>SLMAC-Tag</b> ( $K_m, N, A, C$ )
$z \leftarrow \text{SLFUNC}(K_e, N)$	$H \leftarrow \text{SVHASH}(N, A, C)$
$C \leftarrow \text{SPRG}(z,  M ) \oplus M$	$T \leftarrow \text{SLFUNC}(K_m, H)$
<b>return</b> $C$	<b>return</b> $\lceil T \rceil_\tau$
<b>SENC-Dec</b> ( $K_e, N, C$ )	<b>SLMAC-Ver</b> ( $K_m, N, A, C, T$ )
$z \leftarrow \text{SLFUNC}(K_e, N)$	$H \leftarrow \text{SVHASH}(N, A, C)$
$M \leftarrow \text{SPRG}(z,  C ) \oplus C$	$\bar{T} \leftarrow \text{SLFUNC}(K_m, H)$
<b>return</b> $M$	<b>if</b> $\bar{T} \neq T$
	<b>return</b> $\perp$
	<b>return</b> $\top$

Figure 15: Pseudocode of the encryption scheme  $\text{SENC} = (\text{SENC-Enc}, \text{SENC-Dec})$  and the MAC  $\text{SLMAC} = (\text{SLMAC-Tag}, \text{SLMAC-Ver})$  in terms of the sponge-based primitives  $\text{SLFUNC}$ ,  $\text{SPRG}$ , and  $\text{SVHASH}$ .

In total, one can conclude

$$\begin{aligned}
\text{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{wCR}}(\mathcal{A}) &= \Pr[\text{wCR}(\mathcal{A}) \rightarrow 1] \\
&= \Pr[\text{E} \wedge \text{wCR}(\mathcal{A}) \rightarrow 1] + \Pr[\neg \text{E} \wedge \text{wCR}(\mathcal{A}) \rightarrow 1] \\
&\leq \Pr[\text{wbind}(\mathcal{B}) \rightarrow 1] + \Pr[\text{CR}(\mathcal{C}) \rightarrow 1] \\
&= \text{Adv}_{\mathcal{F}}^{\text{wbind}}(\mathcal{B}) + \text{Adv}_{\mathcal{G}}^{\text{CR}}(\mathcal{C}),
\end{aligned}$$

which finishes the proof.  $\square$

## B.2 Proof of Theorem 4

*Proof.* Let  $\mathcal{A}$  be a CR adversary against  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  with output  $((K, X), (\bar{K}, \bar{X}))$  and let  $\text{E}$  be the event that  $K = \bar{K}$  and  $\mathcal{H}(X) = \mathcal{H}(\bar{X})$ .

Firstly, we construct a CR adversary  $\mathcal{B}$  against  $\mathcal{H}$ . It runs  $((K, X), (\bar{K}, \bar{X})) \leftarrow \mathcal{A}()$  and outputs  $(X, \bar{X})$ . If  $\mathcal{A}$  is successful, we know in particular that  $(K, X) \neq (\bar{K}, \bar{X})$ . Next assume that additionally to  $\mathcal{A}$  being successful, event  $\text{E}$  holds, that means  $K = \bar{K}$  and  $\mathcal{H}(X) = \mathcal{H}(\bar{X})$ . Thus adversary  $\mathcal{B}$  is successful as  $X, \bar{X}$  are different but map to the same element under  $\mathcal{H}$ . Hence we have shown

$$\Pr[\text{E} \wedge \text{CR}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{CR}(\mathcal{B}) \rightarrow 1].$$

Secondly, we construct a bind adversary  $\mathcal{C}$  against  $\mathcal{F}'$ . It runs  $((K, X), (\bar{K}, \bar{X})) \leftarrow \mathcal{A}()$  and outputs  $((K, \mathcal{H}(X)), (\bar{K}, \mathcal{H}(\bar{X})))$ . If  $\mathcal{A}$  is successful, then  $(K, X) \neq (\bar{K}, \bar{X})$  and  $\text{MAC}[\mathcal{H}, \mathcal{F}'].\text{Tag}(K, X) = \text{MAC}[\mathcal{H}, \mathcal{F}'].\text{Tag}(\bar{K}, \bar{X})$ . By definition of  $\text{MAC}[\mathcal{H}, \mathcal{F}']$  the latter is equivalent to  $\mathcal{F}'(K, \mathcal{H}(X)) = \mathcal{F}'(\bar{K}, \mathcal{H}(\bar{X}))$ . If additionally to  $\mathcal{A}$  being successful, event  $\neg \text{E}$  holds, we know that either  $K \neq \bar{K}$  or  $\mathcal{H}(X) \neq \mathcal{H}(\bar{X})$ . Therefore adversary  $\mathcal{C}$  is successful, because  $(K, \mathcal{H}(X)) \neq (\bar{K}, \mathcal{H}(\bar{X}))$  and  $\mathcal{F}'(K, \mathcal{H}(X)) = \mathcal{F}'(\bar{K}, \mathcal{H}(\bar{X}))$ . This yields

$$\Pr[\neg \text{E} \wedge \text{CR}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{bind}(\mathcal{C}) \rightarrow 1]$$

SLFUNC( $K, X$ )	SPRG( $z, L$ )	SvHASH( $N, A, C$ )
$X_1, \dots, X_l \xleftarrow{r} X$	$l \leftarrow \lceil \frac{L}{r} \rceil$	$A_1, \dots, A_\alpha \xleftarrow{r} \text{pad}_{10^*}(A, r)$
$Y \leftarrow K \parallel \text{IV}$	$Z \leftarrow \varepsilon$	$C_1, \dots, C_\gamma \xleftarrow{r} \text{pad}_{10^*}(C, r)$
$S \leftarrow \rho(Y)$	$S \leftarrow z$	$Y \leftarrow N \parallel \text{IV}$
<b>for</b> $i = 1, \dots, l$	$Z \leftarrow Z \parallel \lceil S \rceil_r$	$S \leftarrow \rho(Y)$
$Y \leftarrow (\lceil S \rceil_r \oplus X_i) \parallel \lfloor S \rfloor_c$	<b>for</b> $i = 1, \dots, l-1$	<b>for</b> $i = 1, \dots, \alpha$
$S \leftarrow \rho(Y)$	$S \leftarrow \rho(S)$	$Y \leftarrow (\lceil S \rceil_r \oplus A_i) \parallel \lfloor S \rfloor_c$
<b>return</b> $S$	$Z \leftarrow Z \parallel \lceil S \rceil_r$	$S \leftarrow \rho(Y)$
	<b>return</b> $\lceil Z \rceil_L$	$S \leftarrow \lceil S \rceil_r \parallel (\lfloor S \rfloor_c \oplus (1 \parallel 0^{c-1}))$
		<b>for</b> $i = 1, \dots, \gamma$
		$Y \leftarrow (\lceil S \rceil_r \oplus C_i) \parallel \lfloor S \rfloor_c$
		$S \leftarrow \rho(Y)$
		$H \leftarrow \lceil S \rceil_w$
		<b>return</b> $H$

Figure 16: Pseudocode of the sponge-based function SLFUNC, pseudorandom generator SPRG, and hash function SvHASH, constructed from a random transformation  $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$ , where  $n = r + c$ .

and in total we can conclude

$$\begin{aligned}
\text{Adv}_{\text{MAC}[\mathcal{H}, \mathcal{F}']}^{\text{CR}}(\mathcal{A}) &= \Pr[\text{CR}(\mathcal{A}) \rightarrow 1] \\
&= \Pr[\text{E} \wedge \text{CR}(\mathcal{A}) \rightarrow 1] + \Pr[\neg \text{E} \wedge \text{CR}(\mathcal{A}) \rightarrow 1] \\
&\leq \Pr[\text{CR}(\mathcal{B}) \rightarrow 1] + \Pr[\text{bind}(\mathcal{C}) \rightarrow 1] \\
&= \text{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{B}) + \text{Adv}_{\mathcal{F}'}^{\text{bind}}(\mathcal{C}).
\end{aligned}$$

This finishes the proof.  $\square$

### B.3 Proof of Theorem 5

*Proof.* Consider a CMT adversary  $\mathcal{A}$  against  $\text{N1}[\text{SE}[\mathcal{F}, \mathcal{G}], \text{MAC}[\mathcal{H}, \mathcal{F}']]$ . By Theorem 2, there exist adversaries  $\mathcal{A}_1$  and  $\mathcal{A}_2$  such that

$$\text{Adv}_{\text{N1}[\text{SE}[\mathcal{F}, \mathcal{G}], \text{MAC}[\mathcal{H}, \mathcal{F}']]}^{\text{CMT}}(\mathcal{A}) \leq \text{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{wCR}}(\mathcal{A}_1) + \text{Adv}_{\text{MAC}[\mathcal{H}, \mathcal{F}']}^{\text{CR}}(\mathcal{A}_2).$$

Applying Theorem 3 for adversary  $\mathcal{A}_1$ , then yields adversaries  $\mathcal{B}$  and  $\mathcal{C}$  such that

$$\text{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{wCR}}(\mathcal{A}_1) \leq \text{Adv}_{\mathcal{F}}^{\text{wbind}}(\mathcal{B}) + \text{Adv}_{\mathcal{G}}^{\text{CR}}(\mathcal{C}).$$

Analogously, Theorem 4 for adversary  $\mathcal{A}_2$  gives adversaries  $\mathcal{D}$  and  $\mathcal{E}$  such that

$$\text{Adv}_{\text{MAC}[\mathcal{H}, \mathcal{F}']}^{\text{CR}}(\mathcal{A}_2) \leq \text{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{D}) + \text{Adv}_{\mathcal{F}'}^{\text{bind}}(\mathcal{E}).$$

Combining these results shows the claim.  $\square$

### B.4 Proof of Theorem 6

*Proof.* Let  $\mathcal{A}$  be a pbind adversary against SLFUNC with output  $((K, N), (\bar{K}, \bar{N}))$ . For the proof we will refer to the different states of the sponge as  $S_i$  and  $Y_i$  as shown in Figure 3 (upper part). We show that the probability for  $\mathcal{A}$  to win is bounded above by



the probability that for different inputs, the respective outputs of  $\rho$  collide in the first  $k$  bits (we call this a  $k$ -collision) or in the last  $c$  bits (this is referred to as an inner collision). For this we define  $\mathbf{C}$  to be the event that there are  $Y \neq \bar{Y}$  such that  $\lceil \rho(Y) \rceil_k = \lceil \rho(\bar{Y}) \rceil_k$  or  $\lfloor \rho(Y) \rfloor_c = \lfloor \rho(\bar{Y}) \rfloor_c$  for all states  $Y$  and  $\bar{Y}$  occurring during the evaluation of the tuples outputted by  $\mathcal{A}$ . Then we can split up the probability of  $\mathcal{A}$  winning the game  $\mathbf{pbind}$  as follows

$$\Pr[\mathbf{pbind}(\mathcal{A}) \rightarrow 1] = \Pr[\mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] + \Pr[\neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1]. \quad (4)$$

As a next step we show that  $\Pr[\neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] = 0$ . For this consider  $\mathbf{E}$  to be the event that  $N \neq \bar{N}$  and observe that

$$\begin{aligned} \Pr[\neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] &= \Pr[\mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] \\ &\quad + \Pr[\neg \mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1]. \end{aligned} \quad (5)$$

We proceed by showing that both summands in the above equation are zero.

First we show  $\Pr[\mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] = 0$ . We assume that  $\neg \mathbf{C}$  holds, hence we know in particular that for all  $i \in \{0, \dots, l\}$  with  $Y_i \neq \bar{Y}_i$  it holds that  $\lceil \rho(Y_i) \rceil_k \neq \lceil \rho(\bar{Y}_i) \rceil_k$  and  $\lfloor \rho(Y_i) \rfloor_c \neq \lfloor \rho(\bar{Y}_i) \rfloor_c$ . As  $\mathcal{A}$  is successful, we know that  $(K, N) \neq (\bar{K}, \bar{N})$  and  $\lceil S_{l+1} \rceil_k = \text{SLFUNC}(K, N) = \text{SLFUNC}(\bar{K}, \bar{N}) = \lceil \bar{S}_{l+1} \rceil_k$ . Furthermore  $\mathbf{E}$  holds, which implies that  $N \neq \bar{N}$ . We distinguish the following two cases:

1.  $K \neq \bar{K}$ : In this case, we have  $Y_0 = K \neq \bar{K} = \bar{Y}_0$ , hence we obtain that  $\lfloor \rho(Y_0) \rfloor_c \neq \lfloor \rho(\bar{Y}_0) \rfloor_c$  since  $\mathbf{C}$  holds. This implies that  $Y_1 \neq \bar{Y}_1$ , because  $\lfloor Y_1 \rfloor_c = \lfloor \rho(Y_0) \rfloor_c \neq \lfloor \rho(\bar{Y}_0) \rfloor_c = \lfloor \bar{Y}_1 \rfloor_c$ . Repeating this argument for all  $i \in \{1, \dots, l-1\}$  yields  $Y_l \neq \bar{Y}_l$ .
2.  $K = \bar{K}$ : In this case, we use the fact that  $N \neq \bar{N}$  implies the existence of a smallest  $m \in \{1, \dots, l\}$  with  $N_m \neq \bar{N}_m$ . Since the keys we start with are the same and  $N_i = \bar{N}_i$  holds for all  $i < m$ , we can deduce that  $S_m = \bar{S}_m$ . Then  $N_m \neq \bar{N}_m$  yields  $Y_m \neq \bar{Y}_m$ . Therefore we obtain  $\lfloor \rho(Y_m) \rfloor_c \neq \lfloor \rho(\bar{Y}_m) \rfloor_c$  since  $\mathbf{C}$  holds, which implies that  $\lfloor Y_{m+1} \rfloor_c = \lfloor \rho(Y_m) \rfloor_c \neq \lfloor \rho(\bar{Y}_m) \rfloor_c = \lfloor \bar{Y}_{m+1} \rfloor_c$ . Thus in particular it holds that  $Y_{m+1} \neq \bar{Y}_{m+1}$  and repeating this argument for all  $i \in \{m+1, \dots, l-1\}$  shows  $Y_l \neq \bar{Y}_l$ .

For both cases, we have shown that  $Y_l \neq \bar{Y}_l$ . Hence, since  $\mathbf{C}$  holds, we can deduce that  $\lceil S_{l+1} \rceil_k = \lceil \rho(Y_l) \rceil_k \neq \lceil \rho(\bar{Y}_l) \rceil_k = \lceil \bar{S}_{l+1} \rceil_k$ , which contradicts the assumption that  $\mathcal{A}$  is successful. Therefore the probability  $\Pr[\mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1]$  must be zero.

Next we show that also  $\Pr[\neg \mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] = 0$ . For this, note that event  $\mathbf{C}$  is composed of the event that there is a  $k$ -collision, denoted by  $\mathbf{C}_k$ , and the event that there is an inner collision, denoted by  $\mathbf{C}_c$ . More precisely, we have  $\mathbf{C} = \mathbf{C}_k \vee \mathbf{C}_c$ , which implies that  $\neg \mathbf{C} = \neg \mathbf{C}_k \wedge \neg \mathbf{C}_c$  holds. Using this, we obtain that

$$\begin{aligned} \Pr[\neg \mathbf{E} \wedge \neg \mathbf{C} \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] &= \Pr[\neg \mathbf{E} \wedge \neg \mathbf{C}_k \wedge \neg \mathbf{C}_c \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] \\ &\leq \Pr[\neg \mathbf{E} \wedge \neg \mathbf{C}_k \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] \end{aligned} \quad (6)$$

and hence it suffices to show  $\Pr[\neg \mathbf{E} \wedge \neg \mathbf{C}_k \wedge \mathbf{pbind}(\mathcal{A}) \rightarrow 1] = 0$ , which will be done in the following.

Since  $\mathcal{A}$  wins game  $\mathbf{bind}$ , we can deduce that  $(K, N) \neq (\bar{K}, \bar{N})$  and  $\lceil S_{l+1} \rceil_k = \text{SLFUNC}(K, N) = \text{SLFUNC}(\bar{K}, \bar{N}) = \lceil \bar{S}_{l+1} \rceil_k$ . By assumption event  $\neg \mathbf{E}$  holds, that means  $N = \bar{N}$ , which in turn implies that  $K \neq \bar{K}$ . By construction of  $\text{SLFUNC}$ , we have  $Y_0 = K \neq \bar{K} = \bar{Y}_0$  and hence  $\lceil \rho(Y_0) \rceil_k \neq \lceil \rho(\bar{Y}_0) \rceil_k$  as  $\neg \mathbf{C}_k$  holds. Then we know in particular that  $S_1 = \rho(Y_0) \neq \rho(\bar{Y}_0) = \bar{S}_1$ . By assumption, we have  $N = \bar{N}$  and thus obtain

$$Y_1 = S_1 \oplus (N_1 \parallel 0^c) \neq \bar{S}_1 \oplus (\bar{N}_1 \parallel 0^c) = \bar{Y}_1,$$

where  $N_1 = \lceil N \rceil_r$  and  $\overline{N}_1 = \lceil \overline{N} \rceil_r$ . This puts us in the same situation we started with and by repeating the above argument for all  $i \in \{1, \dots, l-1\}$ , we get  $Y_l \neq \overline{Y}_l$  and hence  $\lceil S_{l+1} \rceil_k = \lceil \rho(Y_l) \rceil_k \neq \lceil \rho(\overline{Y}_l) \rceil_k = \lceil \overline{S}_{l+1} \rceil_k$ . This contradicts the assumption that  $\mathcal{A}$  is a successful adversary hence the probability that  $\mathcal{A}$  wins the game **pbind** and simultaneously  $\neg E$  and  $\neg C_k$  hold is zero. Hence we have shown

$$\Pr[\neg E \wedge \neg C_k \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] = 0, \quad (7)$$

which proves by Eq. (6), that also  $\Pr[\neg E \wedge \neg C \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] = 0$ . Note that for this part of the proof, we split up  $C$  into  $C_k$  and  $C_c$  to emphasize that we need just the absence of  $k$ -collisions (i.e.,  $\neg C_k$ ) to show  $\Pr[\neg E \wedge \neg C \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] = 0$ . This will become relevant later in the proof when the bound for the wbind advantage is proven.

As both summands in Eq. (5) were shown to be zero, Eq. (4) simplifies to

$$\Pr[\text{pbind}(\mathcal{A}) \rightarrow 1] = \Pr[C \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] \leq \Pr[C].$$

For an adversary  $\mathcal{A}$  with  $q$  queries to the random transformation  $\rho$ , the probability to find a  $k$ -collision is

$$\sum_{j=1}^{q-1} \frac{j 2^{n-k}}{2^n} = \sum_{j=1}^{q-1} \frac{j}{2^k} = \frac{q^2 - q}{2^{k+1}}$$

and the probability to find an inner collision is

$$\sum_{j=1}^{q-1} \frac{j 2^{n-c}}{2^n} = \sum_{j=1}^{q-1} \frac{j}{2^c} = \frac{q^2 - q}{2^{c+1}}.$$

As  $C$  is the event that there is either a  $k$ -collision or an inner collision, we obtain

$$\text{Adv}_{\text{SLFUNC}}^{\text{wbind}}(\mathcal{A}) = \Pr[\text{pbind}(\mathcal{A}) \rightarrow 1] \leq \Pr[C] = \frac{q^2 - q}{2^{k+1}} + \frac{q^2 - q}{2^{c+1}},$$

which proves the first part of the theorem.

From this we can easily derive the bound for the pbind advantage, as for  $k = n$  it holds that game **pbind** equals game **bind** and thus

$$\text{Adv}_{\text{SLFUNC}}^{\text{bind}}(\mathcal{A}) = \text{Adv}_{\text{SLFUNC}}^{\text{pbind}}(\mathcal{A}) \leq \frac{q^2 - q}{2^{c+1}} + \frac{q^2 - q}{2^{n+1}}.$$

It is left to prove the bound for the advantage of a wbind adversary  $\mathcal{A}$ . Note that for such an  $\mathcal{A}$  we have  $k = n$  and the event  $E$  can never occur. The latter is the case, as  $K \neq \overline{K}$ ,  $N = \overline{N}$ , and  $S_{l+1} = \overline{S}_{l+1}$  has to hold for the output of  $\mathcal{A}$ , otherwise  $\mathcal{A}$  loses game **pbind**. Therefore it holds that

$$\begin{aligned} \Pr[\text{pbind}(\mathcal{A}) \rightarrow 1] &= \Pr[E \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] + \Pr[\neg E \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] \\ &= \Pr[\neg E \wedge \text{pbind}(\mathcal{A}) \rightarrow 1]. \end{aligned}$$

As before, we let  $C_n$  be the event that there is a  $n$ -collision, i.e., that there are two different inputs such that the respective outputs of  $\rho$  collide. Then we obtain

$$\begin{aligned} \Pr[\neg E \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] &= \Pr[\neg E \wedge \neg C_n \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] + \Pr[\neg E \wedge C_n \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] \\ &\leq \Pr[\neg E \wedge \neg C_n \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] + \Pr[C_n] \\ &= \Pr[C_n], \end{aligned}$$

where the last equality holds by Eq. (7) for  $k = n$ . Putting together the above results, and using the bound that was computed for the probability of  $n$ -collisions, yields

$$\text{Adv}_{\text{SLFUNC}}^{\text{wbind}}(\mathcal{A}) = \Pr[\text{pbind}(\mathcal{A}) \rightarrow 1] = \Pr[\neg E \wedge \text{pbind}(\mathcal{A}) \rightarrow 1] \leq \Pr[C_n] = \frac{q^2 - q}{2^{n+1}}.$$

This finishes the proof of the theorem.  $\square$

## B.5 Proof of Theorem 7

*Proof.* For sake of simplicity, we assume that  $m$  is a multiple of the rate  $r$  and we set  $l = \frac{m}{r}$ . Finding a collision for SPRG equals finding distinct states  $z_1 \neq z_2$  that agree on their outer state (first  $r$  bits) for  $l$  invocations of  $\rho$ . It holds that

$$\begin{aligned} \mathbf{Adv}_{\text{SPRG}}^{\text{CR}}(\mathcal{A}) &= \Pr \left[ [z_1]_r = [z_2]_r \wedge [\rho^1(z_1)]_r = [\rho^1(z_2)]_r \right. \\ &\quad \left. \wedge \dots \wedge [\rho^l(z_1)]_r = [\rho^l(z_2)]_r \mid z_1, z_2 \leftarrow \mathcal{A}^\rho() \right]. \end{aligned}$$

We define events  $\mathbf{E}_0, \mathbf{E}_1, \dots, \mathbf{E}_l$ , where  $\mathbf{E}_i$  equals  $[\rho^i(z_1)]_r = [\rho^i(z_2)]_r$ , for  $z_1, z_2 \leftarrow \mathcal{A}^\rho()$ . Thus

$$\begin{aligned} \mathbf{Adv}_{\text{SPRG}}^{\text{CR}}(\mathcal{A}) &= \Pr [\mathbf{E}_0 \wedge \mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_l] \\ &\leq \Pr [\mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_l] \\ &= \Pr[\mathbf{E}_1] \cdot \Pr[\mathbf{E}_2 \mid \mathbf{E}_1] \cdot \dots \cdot \Pr[\mathbf{E}_l \mid \mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_{l-1}] \\ &= \Pr[\mathbf{E}_1] \cdot \prod_{i=2}^l \Pr[\mathbf{E}_i \mid \mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_{i-1}]. \end{aligned}$$

Bounding event  $\mathbf{E}_1$  is a simple counting argument. Since  $\mathcal{A}$  makes  $q$  queries to  $\rho$ , let  $\mathbf{X}_i$  be the event that the  $i$ -th query to  $\rho$  triggers  $\mathbf{E}_1$ . It holds that

$$\Pr[\mathbf{E}_1] \leq \sum_{i=1}^q \Pr[\mathbf{X}_i] = \sum_{i=1}^q \frac{(i-1)2^c}{2^n} = \sum_{i=1}^q \frac{i-1}{2^r} = \frac{q^2 - q}{2^{r+1}}.$$

Note that only the probability for  $\mathbf{E}_1$  depends on  $q$  as  $\mathcal{A}$  can only influence the outcome of the first application of  $\rho$  with his queries. For the remaining events  $(\mathbf{E}_2, \dots, \mathbf{E}_l)$ , the fact that  $\rho$  is a random function yields that the events are independent from one another which gives for all  $i \in \{2, \dots, l\}$

$$\Pr[\mathbf{E}_i \mid \mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_{i-1}] = \Pr[\mathbf{E}_i] \leq \frac{1}{2^r}.$$

Collecting the above finally provides

$$\mathbf{Adv}_{\text{SPRG}}^{\text{CR}}(\mathcal{A}) \leq \Pr[\mathbf{E}_1] \cdot \prod_{i=2}^l \Pr[\mathbf{E}_i \mid \mathbf{E}_1 \wedge \dots \wedge \mathbf{E}_{i-1}] \leq \frac{q^2 - q}{2^{r+1}} \cdot \prod_{i=2}^l \frac{1}{2^r} = \frac{q^2 - q}{2^{m+1}}.$$

This finishes the proof.  $\square$

## B.6 Hierarchy of Binding Notions

The following theorem describes the relation between the different binding notions.

**Theorem 11.** *Let  $F: \mathcal{K} \times \{0, 1\}^x \rightarrow \{0, 1\}^y$  be a function. Then for any adversary  $\mathcal{A}$  there exists an adversary  $\mathcal{B}$  such that*

$$\mathbf{Adv}_F^{\text{bind}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{pbind}}(\mathcal{B})$$

*and for any adversary  $\mathcal{C}$  there exists an adversary  $\mathcal{D}$  such that*

$$\mathbf{Adv}_F^{\text{wbind}}(\mathcal{C}) \leq \mathbf{Adv}_F^{\text{bind}}(\mathcal{D}).$$

*Proof.* For the first part, we consider a bind adversary  $\mathcal{A}$  against  $F$  and construct a pbind adversary  $\mathcal{B}$  against  $F$ . It runs  $((K, X), (\bar{K}, \bar{X})) \leftarrow \mathcal{A}$  and outputs  $((K, X), (\bar{K}, \bar{X}))$ . If  $\mathcal{A}$

succeeds then  $(K, X) \neq (\overline{K}, \overline{X})$  and  $F(K, X) = F(\overline{K}, \overline{X})$ . The latter implies in particular that  $\lceil Y_1 \rceil_l = \lceil Y_2 \rceil_l$  for  $Y_i \leftarrow F(K_i, X_i)$ . Thus  $\mathcal{B}$  is also successful in game **pbind**.

For the second part, let  $\mathcal{C}$  be a **wbind** adversary against  $F$  and construct a **bind** adversary  $\mathcal{D}$  against  $F$ . It runs  $((K, X), (\overline{K}, \overline{X})) \leftarrow \mathcal{A}$  and outputs  $((K, X), (\overline{K}, \overline{X}))$ . If  $\mathcal{C}$  succeeds then  $K \neq \overline{K}$  and  $F(K, X) = F(\overline{K}, \overline{X})$ . Note that the first condition implies in particular that  $(K, X) \neq (\overline{K}, \overline{X})$  and hence  $\mathcal{D}$  is successful as well.  $\square$