

# Committing Authenticated Encryption

## Sponges vs. Block-Ciphers in the case of the NIST LWC Finalists

Juliane Krämer<sup>1</sup>, Patrick Struck<sup>2</sup>, and Maximiliane Weishäupl<sup>1</sup>

<sup>1</sup> Universität Regensburg, Germany  
{juliane.kraemer,maximiliane.weishaeupl}@ur.de

<sup>2</sup> Universität Konstanz, Germany  
patrick.struck@uni-konstanz.de

**Abstract.** Committing security has gained considerable attention in the field of authenticated encryption (AE). This can be traced back to a line of recent attacks, which suggests that AE schemes used in practice should not only provide confidentiality and authenticity, but also committing security. Roughly speaking, a committing AE scheme guarantees that ciphertexts will decrypt only for one key. Despite the recent research effort in this area, the finalists of the NIST lightweight cryptography standardization process have not been put under consideration yet. We close this gap by providing an analysis of these schemes with respect to their committing security. Despite the structural similarities the finalists exhibit, our results are of a quite heterogeneous nature: We break four of the schemes with effectively no costs, while for two schemes our attacks are costlier, yet still efficient. For the remaining three schemes ISAP, ASCON, and (a slightly modified version of) SCHWAEMM, we give formal security proofs. Our analysis reveals that sponges—due to their large states—are more favorable for committing security compared to block-ciphers.

---

\* Work of Juliane Krämer was supported by the German Research Foundation (DFG) – SFB 1119 – 236615297. Patrick Struck acknowledges funding by the Bavarian State Ministry of Science and the Arts in the framework of the bidt Graduate Center for Postdocs (while working at University of Regensburg) and 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).

# Table of Contents

1	Introduction.....	3
1.1	Contribution.....	4
1.2	Related Work.....	6
2	Authenticated Encryption and the NIST LWC Finalists.....	7
2.1	Notation.....	7
2.2	Definitions.....	8
2.3	NIST LWC Finalists.....	8
3	Security Analysis.....	13
3.1	Elephant.....	13
3.2	Romulus.....	15
3.3	GIFT-COFB.....	16
3.4	PHOTON-Beetle.....	18
3.5	Xoodyak.....	19
3.6	TinyJAMBU.....	20
3.7	ISAP.....	22
3.8	Ascon.....	24
3.9	Schwaemm.....	26
4	Conclusion.....	28
	References.....	29
A	Additional Preliminaries.....	34
A.1	Paddings and Security Notions.....	34
A.2	(Tweakable) Block-Ciphers.....	36
A.3	Sponges.....	36
A.4	Existing Results.....	37
B	Deferred Proofs.....	40
B.1	Proof of Theorem 3 (Elephant).....	40
B.2	Proof of Theorem 4 (Romulus).....	42
B.3	Proof of Theorem 5 (GIFT-COFB).....	47
B.4	Proof of Theorem 6 (PHOTON-Beetle).....	52
B.5	Proof of Theorem 7 (Xoodyak).....	54
B.6	Proof of Theorem 8 (TinyJAMBU).....	56
B.7	Proof of Theorem 9 (ISAP).....	58
B.8	Proof of Theorem 10 (Ascon).....	62
B.9	Proof of Theorem 11 (Schwaemm).....	69

## 1 Introduction

The most fundamental cryptographic concept is symmetric encryption, allowing two parties, Alice and Bob, which share some secret key, to securely exchange messages. The initial goal—and still a cornerstone—is confidentiality which prevents anyone but Alice and Bob from recovering the message from a ciphertext. In modern cryptography, security requirements have been enhanced to also incorporate authenticity<sup>3</sup> which ensures that no third party can produce a ciphertext that Bob would accept as one generated by Alice. On that account, *authenticated encryption* (AE), which encompasses both confidentiality and authenticity was introduced and has, since then, become the gold standard [11, 48]. While authenticated encryption has undergone some changes—from probabilistic over IV-based to nonce-based—nowadays, the research community agrees on *authenticated encryption with associated data* as the right approach. Such a scheme generates a ciphertext  $C$  by encrypting a message  $M$  under a *context*  $(K, N, A)$ , consisting of a key  $K$ , a nonce  $N$ , and associated data  $A$ . Authenticity should hold for both the associated data and the message. In contrast, confidentiality is required only for the message.

The relevance of authenticated encryption is not only reflected by the conducted research, but also by the fact that AE schemes are deployed ubiquitously, e.g., in TLS 1.3 [51]. The CAESAR competition for authenticated encryption [11] and the recent NIST lightweight cryptography (LWC) standardization process [48], both called specifically for AE schemes which are deemed secure if they provide both confidentiality and authenticity.

However, a series of recent attacks [2, 32, 43] has shown that our understanding of what a secure AE scheme is, has to change once again. The *Facebook message franking attack* [32] enabled Alice, a malicious user, to send an offensive or even illegal content to Bob. If Bob tries to report this, it will fail as Facebook will see a harmless content—prepared by Alice as part of the attack—instead. Further examples are the *subscribe with Google attack* [2] and the *partitioning oracle attack* [43]. The latter allows for more efficient key recovery: The adversary crafts a ciphertext that decrypts validly under multiple keys (for instance taken from a leaked list of candidate keys) and sends it to the recipient; If the recipient rejects the ciphertext as invalid, the adversary can rule out all keys which are valid for the sent ciphertext.

In fact, these attacks can all be traced back to the same problem: The existence of ciphertexts that decrypt validly under more than one key. This is neither prevented by confidentiality nor by authenticity, barring the need for an additional security notion. To this end, committing security [9] was defined by requiring each ciphertext to be a commitment to the key ( $\text{CMT}_K$ ) or even to the whole context (CMT). The latter notion is the strongest one and is formalized by the following security game: The adversary outputs two tuples  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ , each consisting of key, nonce, associated data,

<sup>3</sup> Otherwise attacks like the *padding oracle attack* [53], which exploits the absence of any authentication mechanism, are possible.

and message, and wins if their contexts differ, i.e.,  $(K, N, A) \neq (\overline{K}, \overline{N}, \overline{A})$ , and  $\text{AE.ENC}(K, N, A, M) = \text{AE.ENC}(\overline{K}, \overline{N}, \overline{A}, \overline{M})$  holds, for AE the scheme under consideration.<sup>4</sup>

The aforementioned attacks demonstrate that the consequences of using non-committing authenticated encryption can be severe. Considering that there are most likely more attacks, which have yet to be discovered, it is important to deal with this problem. One possibility would be to design protocols in such a way that usage of non-committing authenticated encryption does not result in attacks.<sup>5</sup> However, this approach is ill-advised as it requires separate analysis for the different protocols and simply puts the burden on the designers of a protocol. A better approach is to design authenticated encryption schemes that are committing which can then be used in different protocols without worrying about committing attacks.

To this end, AE schemes used in practice need to be analyzed with respect to committing security. This process has already begun and a number of commonly used AE schemes (GCM, SIV, CCM, EAX, OCB3) have been examined [9,45]. A majority of them was shown to not achieve committing security. Arguably among the most important AE schemes are the finalists of the NIST LWC standardization process. While these schemes have received significant analysis with respect to confidentiality and authenticity [52], we are not aware of any research with respect to their committing security. Despite the announcement of the NIST that ASCON will be standardized, we consider all finalists relevant objects to study—especially considering that none of the finalists suffer from any weaknesses regarding their claimed security levels [52].

## 1.1 Contribution

In this paper we analyze the committing security of the NIST LWC finalists that are based on (tweakable) block-ciphers or sponges.<sup>6</sup> More precisely, we focus on the authenticated encryption mode of the schemes, while the underlying primitives, i.e., (tweakable) block-ciphers or permutations, are assumed to be ideal. We follow the example of [45], to define a boundary between committing insecure and secure schemes. The line is drawn at 64-bit security, i.e., a scheme providing at least 64-bit committing security is called secure, while all others are called insecure. This number stems from the fact that the cost of a committing attack is bounded below by the cost of finding colliding tags. As most present-day schemes employ 128-bit tags, the latter can be done with about  $2^{64}$  queries using a birthday attack.

We divide the NIST LWC finalists into two groups: Firstly, ELEPHANT [16] and ISAP [29] follow the *Encrypt-then-MAC* (EtM) paradigm, as they start by

<sup>4</sup> For tidy schemes [47], an equivalent characterization of the notion requires the adversary to find  $(C, (K, N, A), (\overline{K}, \overline{N}, \overline{A}))$  such that  $(K, N, A) \neq (\overline{K}, \overline{N}, \overline{A})$  and  $\text{AE.DEC}(K, N, A, C), \text{AE.DEC}(\overline{K}, \overline{N}, \overline{A}, C) \neq \perp$  [9].

<sup>5</sup> Note that the Facebook protocol was changed to prevent the message franking attack.

<sup>6</sup> This covers all finalists except GRAIN-128AEAD [37].

encrypting the message and then authenticate the resulting ciphertext alongside the context. Secondly, ROMULUS [40], GIFT-COFB [3], PHOTON-BEETLE [5], XOODYAK [23], TINYJAMBU [56], ASCON [31], and SCHWAEMM [7] share a common structure, in the sense that all of them first process the context and then the message. We refer to schemes of the second type as *Context-pre-Processing* (CpP) schemes.

Surprisingly, even though the NIST finalists show strong structural similarities, our results regarding their committing security are of a very heterogeneous nature. As can be seen in Table 1, the results vary from attacks that require essentially no queries<sup>7</sup> to attacks that are costlier—while still using significantly less than  $2^{64}$  queries—and proofs showing about 64-bit committing security. In summary, there are four schemes we break completely, two schemes we break efficiently, and three schemes<sup>8</sup> for which we show committing security.

While our attacks exploit different vulnerabilities, some of them share the same fundamental idea. This is the case for ROMULUS and GIFT-COFB, which are both block-cipher-based AE schemes. Further, both feature a state-update-function, which is invoked in an alternating manner with the block-cipher. The attacks boil down to the fact that for a fixed ciphertext, key, and nonce, one can find associated data such that the ciphertext decrypts validly under this context. For this, starting from the target ciphertext, the component that processes the message is inverted. Then the fact that associated data blocks are XORed onto the whole state is used to connect the initial state with the state obtained from the reverse computation. This attack strategy depends heavily on the invertibility of the state-update-function. For ROMULUS, we show that such an inversion is always possible, while for GIFT-COFB it works with a probability of  $\frac{1}{2}$ . This implies that the attack cost for GIFT-COFB depends on the length of the ciphertext, however, by choosing a short ciphertext we obtain a very efficient attack. The XORing of an input onto the whole state is a vulnerability that is also exploited in our attack on ELEPHANT, a tweakable block-cipher-based AE scheme. In contrast to ROMULUS and GIFT-COFB, ELEPHANT is an Encrypt-then-MAC scheme. This structure simplifies the committing attack as we only need to find two different contexts that verify the ciphertext correctly—in this case the decryption of ELEPHANT will never return  $\perp$ .<sup>9</sup> Due to this, it suffices to concentrate on the MAC and, more precisely, finding a tag collision. The latter is easily achieved, as the associated data is XORed to the full state during the tag generation of ELEPHANT.

Except for these three schemes, none of the other NIST finalists carry out a full-state XOR. XOODYAK arguably comes very close to this situation, as it is based on a full-state sponge, which reserves only a few bits for padding, which are not directly accessible via the inputs. Therefore, we are able to control most of the state by a direct XOR, while for the remaining bits a birthday attack

<sup>7</sup> More precisely, these attacks need only the minimal cost of computing the respective encryption algorithm twice (once for each of the output tuples).

<sup>8</sup> Note that we consider a slightly modified version of SCHWAEMM.

<sup>9</sup> Menda et al. [45] coin this property as *NoFailDecrypt*.

is applied. Similarly, the attack on TINYJAMBU, a block-cipher-based scheme, also boils down to a birthday attack. We exploit that TINYJAMBU uses a tag of just 64 bits (the shortest one among all finalists<sup>10</sup>), hence a tag collision can be produced with reasonable cost.

Our attack on PHOTON-BEETLE, a sponge-based AE scheme, exploits the choice of the initial state. For most of the finalists, this state contains some fixed initialization vector, whereas for PHOTON-BEETLE it consists exclusively of key and nonce. However, this implies that the initial state can be controlled completely by a committing adversary, which will turn out to be the key ingredient of our attack. Simply speaking, the attack allows to choose an intermediate state (outcome of the context-pre-processing) that results in the same ciphertext. We can invert this intermediate state for different associated data and take the outcome as the key-nonce pair.

None of these attacks are applicable to any of the sponge-based schemes ISAP, ASCON, and SCHWAEMM. We give security proofs for these three schemes, showing that they achieve about 64-bit committing security. The high-level idea of all proofs is similar: we show that the schemes can be viewed as plain sponge constructions and give bounds for finding colliding tags, which directly translate to bounds on the committing security. Some extra care is necessary when dealing with the core features of the schemes—the re-keying mechanism deployed in ISAP and the state-/output-blinding applied in both ASCON and SCHWAEMM.

Our analysis reveals that sponges are better suited for achieving committing authenticated encryption. Simultaneously, our attacks against PHOTON-BEETLE and XOODYAK show that sponges are not always committing: Full-state sponges suffer from inherent weaknesses when it comes to committing security, regardless of how often the full-state property occurs.<sup>11</sup> To achieve security, a significant part of the state has to be kept “out-of-reach” of the committing adversary, meaning that it should not be affected by the inputs. Sponges are better suited for this due to their significantly larger states ( $\geq 256$  bits) compared to block-ciphers that typically use smaller states (128 bits). Since a committing adversary knows the key, the general advantage of block-ciphers—namely, the permutation remains concealed from the adversary—vanishes.

## 1.2 Related Work

Committing security can be traced back to [1,33] where the focus was on public-key encryption. In [34]—using the name *key-robustness*—Farshim et al. gave first definitions of committing security for symmetric encryption. Recently, Bellare and Hoang [9] introduced different variants of committing security for authenticated encryption, covering the prior variants where a ciphertext is a commitment to the key, but also stronger forms where a ciphertext is a commitment to all inputs. Ultimately, Menda et al. [45] developed a framework for fine-grained committing security notions. Instead of just having a ciphertext being a commitment

<sup>10</sup> ELEPHANT uses 64-bit tags as well, but also gives a parameter set with 128-bit tags.

<sup>11</sup> For PHOTON-BEETLE only the initial state exhibits the full-state property.

<sup>12</sup> The costs correlate with the length of the ciphertext.

Table 1: Overview of results: a  $\times$  indicates a committing attack with essentially no queries; a  $\blacklozenge$  indicates a committing (CMT) attack with significantly less than  $2^{64}$  queries; and a  $\checkmark$  indicates about 64-bit committing security.

Scheme	CMT	Theorem	Section	Illustration	Pseudocode
ELEPHANT [16]	$\times$	Theorem 3	Section 3.1	Fig. 5	Fig. 20
ROMULUS [40]	$\times$	Theorem 4	Section 3.2	Fig. 6	Fig. 22
GIFT-COFB [3]	$\times^{12}$	Theorem 5	Section 3.3	Fig. 7	Fig. 24
PHOTON-BEETLE [5]	$\times$	Theorem 6	Section 3.4	Fig. 8	Fig. 26
XOODYAK [23]	$\blacklozenge$	Theorem 7	Section 3.5	Fig. 9	Fig. 27
TINYJAMBU [56]	$\blacklozenge$	Theorem 8	Section 3.6	Fig. 10	Fig. 28
ISAP [29]	$\checkmark$	Theorem 9	Section 3.7	Fig. 12	Fig. 30
ASCON [31]	$\checkmark$	Theorem 10	Section 3.8	Fig. 13	Fig. 32
SCHWAEMM [7]	$\checkmark$	Theorem 11	Section 3.9	Fig. 14	Fig. 35

to either the key or all inputs, it allows for variants where it is a commitment to, say, the key and the nonce. Along with these committing notions, they also coin the term *context discovery attacks*. In contrast to committing attacks, which require the adversary to find two contexts that decrypt the same ciphertext, context discovery attacks require finding a context that decrypts a given ciphertext. Concurrently to [45], Chan and Rogaway [20] also developed a more fine-grained definitional framework for committing security, for instance, allowing for variants where the adversary has to use honest keys, i.e., randomly sampled ones.

## 2 Authenticated Encryption and the NIST LWC Finalists

In this section, we first introduce some notation and recall the definitions of authenticated encryption schemes and committing security. We then provide a general classification of the NIST LWC finalists and high-level approaches for the committing attacks.

### 2.1 Notation

Throughout this work, we write  $\{0, 1\}^*$  for the set of bit strings with arbitrary length. By  $\{0, 1\}^{\leq r}$  ( $\{0, 1\}^{\geq r}$ ) we denote the set of bit string with length at most  $r$  (at least  $r$ ). For a bit string  $S$  of length  $n$ , we write  $[S]_r$ ,  $[S]_c$ , and  $[S]_i^j$  for the first  $r$  bits, the last  $c$  bits, and the  $i$ -th to  $j$ -th bits of  $S$ , respectively. For bit strings  $X$ ,  $Y$ , and  $Z$ ,  $|X|$  describes the length of  $X$  and  $Y \parallel Z$  denotes the concatenation of  $Y$  and  $Z$ . For an integer  $k$ , the set  $\{1, \dots, k\}$  is abbreviated as  $[k]$ . We write  $X_1, \dots, X_l \stackrel{r}{\leftarrow} X$  to denote that  $X$  is split into bit strings  $X_1$  to  $X_l$  s.t.  $|X_i| = r$ , for  $i \in [l - 1]$  and  $|X_l| \leq r$ . Bit rotation/shift of  $x$  by  $k$  bits to the left is written as  $x \lll b / x \ll b$  ( $\ggg / \gg$ ) denote the same in the other direction). The encoding of  $x$  into one Byte is described by  $\text{enc}_8(x)$ . For

sake of simplicity, we use  $\iota$  as a generic value for domain separation in several schemes as our results are independent of it. Standard cryptographic background on sponges, block-ciphers (BC), and tweakable block-ciphers (TBC) as well as some results needed for our proofs are given in [Appendix A](#).

## 2.2 Definitions

We recall the relevant definitions of authenticated encryption schemes and committing security.

**Definition 1.** *An authenticated encryption (AE) scheme with associated data is a pair of two algorithms (ENC, DEC) such that*

- 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  $(C, T)$ .
- 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  $(C, T)$  as input and outputs a message  $M$  or a special symbol  $\perp$ .

The sets  $\mathcal{K}$ ,  $\mathcal{N}$ ,  $\mathcal{A}$ ,  $\mathcal{M}$ , and  $\mathcal{C}$  denote the key space, nonce space, associated data space, message space, and ciphertext space, 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  $\mathcal{C} = \{0, 1\}^* \times \{0, 1\}^\tau$ . An AE scheme is called *correct*, if  $\text{DEC}(K, N, A, \text{ENC}(K, N, A, M)) = M$ , for any  $(K, N, A, M)$ . We note further that all considered schemes are *tidy* [47], i.e.,  $M = \text{DEC}(K, N, A, C)$  implies that  $C = \text{ENC}(K, N, A, M)$ . Following the nomenclature from [45], we call the triple  $(K, N, A)$  a *context*.

Simply speaking, committing security requires the adversary to find two context-message pairs that encrypt to the same ciphertext. We recall some weaker forms of committing security in [Appendix A](#).

**Definition 2.** *Let  $\text{AE} = (\text{ENC}, \text{DEC})$  be an authenticated encryption scheme and the game CMT be defined as in [Fig. 1](#). For any adversary  $\mathcal{A}$ , its CMT advantage is defined as*

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

## 2.3 NIST LWC Finalists

The NIST LWC standardization process [48] required the submitted AE schemes to achieve the well-established notions of confidentiality and authenticity. For the former, the requirement was to maintain security as long as nonces are unique—security in case of repeating nonces can be mentioned as a special feature. Committing security is neither mentioned as a requirement nor a feature to be advertised. However, it is important to note that the call for algorithms was published the same year as the first attack [32] that exploited the absence

Game CMT (CMT-3 in [9])	Game CMT (CMT-4 in [9])
1: $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M}) \leftarrow \mathcal{A}()$	1: $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M}) \leftarrow \mathcal{A}()$
2: <b>if</b> $(K, N, A) = (\overline{K}, \overline{N}, \overline{A})$	2: <b>if</b> $(K, N, A, M) = (\overline{K}, \overline{N}, \overline{A}, \overline{M})$
3: <b>return</b> 0	3: <b>return</b> 0
4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$	4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$
5: $(\overline{C}, \overline{T}) \leftarrow \text{ENC}(\overline{K}, \overline{N}, \overline{A}, \overline{M})$	5: $(\overline{C}, \overline{T}) \leftarrow \text{ENC}(\overline{K}, \overline{N}, \overline{A}, \overline{M})$
6: <b>return</b> $((C, T) = (\overline{C}, \overline{T}))$	6: <b>return</b> $((C, T) = (\overline{C}, \overline{T}))$

Fig. 1: Security game CMT. The version on the left requires the contexts to differ, the one on the right requires the context-message pair to differ. Bellare and Hoang [9] showed the notions to be equivalent due to correctness. We use the version on the left side.

of committing security. Due to the more recent research in this area, it can be expected that committing security will either become a requirement or at least a feature considered relevant for cryptographic standards.<sup>13</sup>

The AE schemes that we study in this work are the NIST LWC finalists that are based on (tweakable) block-ciphers or sponges.<sup>14</sup> In the following, we provide some more information about the schemes with regards to similarities and differences. Table 2 gives an overview and details will be given in this section. For each candidate we focus on the main parameter set as described in Table 3.

**Classes of AE Schemes.** The considered schemes can be divided into two categories. The first class encompasses AE schemes that follow the *Encrypt-then-MAC* (EtM) paradigm [10]. These schemes first encrypt the message and subsequently authenticate the resulting ciphertext alongside the nonce and the associated data. The second class comprises AE schemes that follow what we call *Context-pre-Processing* (CpP). These schemes first process the context  $(K, N, A)$  via a function  $\text{ENC}_c$ . The result is then processed together with  $M$  and optionally  $K$  and  $N$ , yielding the ciphertext  $(C, T)$  via a function  $\text{ENC}_M$ . The former (EtM) is illustrated in Fig. 2, the latter (CpP) in Fig. 3. Out of the schemes that we analyze in this work, ELEPHANT and ISAP follow the EtM paradigm, whereas the others—ROMULUS, PHOTON-BEETLE, GIFT-COFB, XOODYAK, TINYJAMBU, ASCON, and SCHWAEMM—follow the CpP-approach.

*Attacking Encrypt-then-MAC AE Schemes* For the EtM schemes, we can focus on the underlying MAC. Once we have two contexts  $(K, N, A) \neq (\overline{K}, \overline{N}, \overline{A})$

<sup>13</sup> This can be seen in other standardization processes by NIST. While additional security properties for signature schemes [21] have not been considered in the initial NIST post-quantum cryptography standardization process [49], they are named as desirable features in the very recent call for additional post-quantum signatures [50].

<sup>14</sup> This covers all finalists except GRAIN-128AEAD [37] which does not fall into one of these categories.

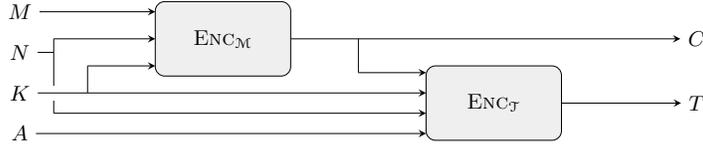


Fig. 2: Illustration of Encrypt-then-MAC AE schemes. Both ELEPHANT and ISAP follow this design.

that verify the same ciphertext  $(C, T)$ , we can immediately derive a committing attack. This is the case because for the described contexts, the decryption algorithm will return some messages  $M, \bar{M} \neq \perp$ .<sup>15</sup> Using the tidyness property, we get  $\text{ENC}(K, N, A, M) = (C, T) = \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ , which implies that we win the game CMT.

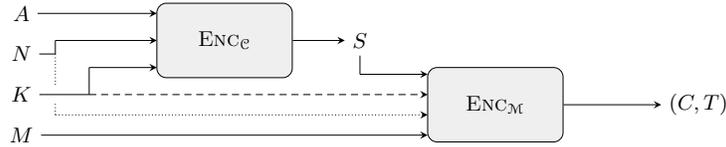


Fig. 3: Illustration of Context-pre-Processing AE schemes. The dotted/dashed arrows indicate that only some of the analyzed schemes exhibit these dependencies: PHOTON-BEETLE and XOODYAK have neither of the two; GIFT-COFB, TINYJAMBU, ASCON, and SCHWAEMM have only the dashed line; and ROMULUS has both lines.

*Attacking Context-pre-Processing AE Schemes* For the CpP schemes, we focus on the state  $S$  that is outputted by  $\text{ENC}_e$  and then fed into  $\text{ENC}_M$ . The general idea is to generate the first context  $(K, N, A)$  and ciphertext  $(C, T)$  at random. Then, we invert  $\text{ENC}_M$  for the same ciphertext  $(C, T)$  and a different key  $\bar{K}$  and nonce  $\bar{N}$ , which yields the state  $\bar{S}$  (along with the message  $\bar{M}$ ). In the last step, we find associated data  $\bar{A}$  such that  $\text{ENC}_e$  with input  $(\bar{K}, \bar{N}, \bar{A})$  results in  $\bar{S}$ , which ultimately yields a committing attack as  $\text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M}) = \text{ENC}_M(\bar{K}, \bar{N}, \text{ENC}_e(\bar{K}, \bar{N}, \bar{A}), \bar{M}) = \text{ENC}_M(\bar{K}, \bar{N}, \bar{S}, \bar{M}) = (C, T)$ . The step of finding  $\bar{A}$  is essentially what was recently coined a *context discovery attack* [45]. This is a stronger attack that easily translates to a committing attack as shown in [45]. Indeed our committing attacks against ROMULUS, GIFT-COFB, ELEPHANT and PHOTON-BEETLE can easily be translated into context discovery attacks; for the other committing attacks this is not the case.

<sup>15</sup> The underlying decryption algorithm never returns  $\perp$ , thus the AE scheme returns  $\perp$  iff the verification of the tag failed. For both ELEPHANT and ISAP, this is the case.

**State-Update-Function.** Out of the nine finalists, ROMULUS, GIFT-COFB, PHOTON-BEETLE, and SCHWAEMM deploy a so-called *state-update-function*<sup>16</sup>. This function takes as input a state  $S$  and some additional input data  $I$ , and outputs a new state  $Y$  and additional output data  $O$  (cf. Fig. 4). The state-update-functions work very similar for all four schemes: one of the outputs is the XOR of the inputs whereas the other is the XOR of the input data  $I$  and some underlying function—which depends on the respective scheme—applied to the input states.

Typically, the state-update-function is used to process the associated data and the message: The current state is used as the input state  $S$  while the associated data or the message—more precisely, a block of it—is used as the input data  $I$ . The output state  $Y$  is used as the new state while the output data  $O$  yields the ciphertext or is simply discarded when the associated data is processed. For decryption, the schemes use the inverse of  $\xi$ . Here it is important to note that, inverse is to be understood *only* in relation to the output data, i.e., for any  $(S, I)$ ,  $\xi(S, I) = (Y, O) \Rightarrow \xi^{-1}(S, O) = (Y, I)$ .

For our attacks against ROMULUS and GIFT-COFB, we need to invert the state-update-function with respect to *both* outputs, which is not obviously possible from the specifications. Our committing attack against PHOTON-BEETLE is independent of the used state-update-function and for SCHWAEMM the state-update-function is incorporated into the committing security proof.

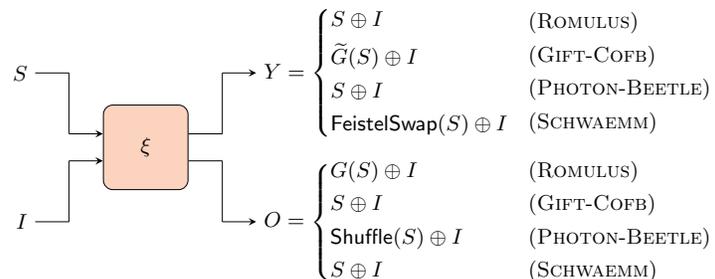


Fig. 4: Illustration of the state-update-function  $\xi$  for the different schemes. The components  $G$ ,  $\tilde{G}$ ,  $\text{Shuffle}$ , and  $\text{FeistelSwap}$ —if relevant for our results—are described along with the schemes in the respective sections.

**Achieving Committing Security via Transformations.** There are several transformations that turn an arbitrary AE scheme into one that is committing. Clearly, such transformations can be applied to the NIST LWC finalists to make them committing. However, there are several reasons against this: Firstly, these transformations often do not achieve CMT security as we target here but weaker

<sup>16</sup> We adopt this name from [40], the other three finalists use different names.

Table 2: Overview of the NIST LWC finalists regarding similarities and differences in their design. Here, CpP stands for Context-pre-Processing, EtM for Encrypt-then-MAC, and (T)BC for (tweakable) block-cipher.

Scheme	Class of scheme	Underlying Primitive	State-update-function
ELEPHANT [16]	EtM	TBC	No
ROMULUS [40]	CpP	TBC	Yes
GIFT-COFB [3]	CpP	BC	Yes
PHOTON-BEETLE [5]	CpP	Sponge	Yes
XOODYAK [23]	CpP	Sponge	No
TINYJAMBU [56]	CpP	BC	No
ISAP [29]	EtM	Sponge	No
ASCAN [31]	CpP	Sponge	No
SCHWAEMM [7]	CpP	Sponge	Yes

Table 3: Parameters of the NIST LWC finalists. Values for rate and capacity are only given for the sponge-based schemes. For ISAP, the parameters are for the version using KECCAK-P, the version using ASCAN-P has  $n = 320$  and  $r = 64$ . Note that for XOODYAK and ISAP, components of the schemes use rates deviating from the ones given above: The  $\text{ENC}_e$  component in XOODYAK is a full-state sponge and in ISAP’s re-keying mechanism a minimal rate of 1 is used.

Scheme	Key $\kappa$	Nonce $\nu$	Tag $\tau$	State $n$	Rate $r$	Capacity $c$
ELEPHANT [16]	128	96	64	160	-	-
ROMULUS [40]	128	128	128	128	-	-
GIFT-COFB [3]	128	128	128	128	-	-
PHOTON-BEETLE [5]	128	128	128	256	128	128
XOODYAK [23]	128	128	128	384	192	192
TINYJAMBU [56]	128	96	64	128	-	-
ISAP [29]	128	128	128	400	144	256
ASCAN [31]	128	128	128	320	64	256
SCHWAEMM [7]	128	256	128	384	256	128

notions [45]. Secondly, these transformations impose some overhead which—especially considering the lightweight aspect of these schemes—might render them impractical. Thirdly, consider, say, ISAP, which comes with a formal security proof incorporating side-channel leakage. Since none of the transformations are analyzed w.r.t. side-channel leakage, applying them to ISAP can render the leakage security guarantees obsolete.

### 3 Security Analysis

Here, we analyze the CMT security of the NIST LWC finalists. For ELEPHANT (cf. Section 3.1), ROMULUS (cf. Section 3.2), GIFT-COFB (cf. Section 3.3), and PHOTON-BEETLE (cf. Section 3.4), we give attacks that break committing security with essentially no cost—requiring the bare minimum of two encryptions. For XOODYAK (cf. Section 3.5) and TINYJAMBU (cf. Section 3.6), we give committing attacks requiring significantly less than  $2^{64}$  queries. For ISAP (cf. Section 3.7), ASCON (cf. Section 3.8), and SCHWAEMM (cf. Section 3.9), we give formal proofs showing that the schemes achieve committing security of about 64-bit. For all of the NIST schemes, we provide a figure illustrating the scheme as well as the pseudocode. The figures are in the respective subsections of the present section, while the pseudocodes can be found in Appendix B.

#### 3.1 ELEPHANT

The AE scheme ELEPHANT [15, 16] is based on tweakable block-ciphers. More precisely, it relies on a cryptographic permutation which gets masked using linear feedback shift registers similar to the masked Even-Mansour construction [35].

**Description of ELEPHANT.** The pseudocode of ELEPHANT is given in Fig. 20 and further illustration is provided in Fig. 5. ELEPHANT follows the Encrypt-then-MAC paradigm, i.e., it first encrypts the message  $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, M)$  and afterwards computes the tag  $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$ . Note that in  $\text{ENC}_{\mathcal{T}}$  the nonce and associated data are padded together, i.e., the first associated data block contains the nonce and the first bits of the associated data. This is in contrast to all other schemes, where the associated data blocks do not contain the nonce. Furthermore, note that the underlying encryption is an involution, i.e., to decrypt a ciphertext, we simply compute  $\text{ENC}_{\mathcal{M}}(K, N, C)$ .

**Committing Attack against ELEPHANT.** Since ELEPHANT follows the EtM paradigm, we only need to focus on the underlying MAC  $\text{ENC}_{\mathcal{T}}$ . If we can find two contexts that verify a ciphertext-tag pair, applying  $\text{ENC}_{\mathcal{M}}$  to the ciphertext and each context, gives back two valid messages. The following attack, which is the simplest one in this work, shows that ELEPHANT [16] does not achieve committing security.

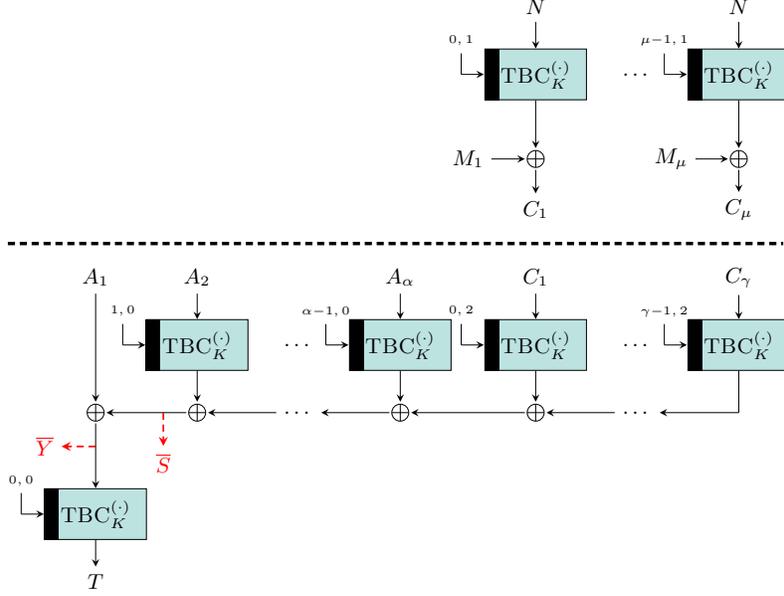


Fig. 5: Illustration of ELEPHANT in terms of  $\text{ENC}_{\mathcal{M}}$  (top) and  $\text{ENC}_{\mathcal{T}}$  (bottom). The states  $\bar{S}$  and  $\bar{Y}$ , marked in red, are used in our CMT attack.

**Theorem 3.** Consider ELEPHANT illustrated in Fig. 5. Let TBC be modeled as an ideal tweakable cipher  $\tilde{\text{E}}$ . Then there exists an adversary  $\mathcal{A}$ , making  $q$  queries to  $\tilde{\text{E}}$ , such that

$$\text{Adv}_{\text{ELEPHANT}}^{\text{CMT}}(\mathcal{A}) = 1,$$

where  $q = 2\mu + 2\gamma + \alpha + \bar{\alpha}$ . Here,  $\mu$  is the number of message blocks while computing  $\text{ENC}_{\mathcal{M}}$  and  $\gamma$  is the number of ciphertext blocks while computing  $\text{ENC}_{\mathcal{T}}$ .<sup>17</sup> Furthermore,  $\alpha$  and  $\bar{\alpha}$  are the number of associated data blocks for the two tuples that  $\mathcal{A}$  outputs.

*Proof (sketch).* The full proof is given in Appendix B.1, here we provide a sketch. Adversary  $\mathcal{A}$  generates a ciphertext  $(C, T)$  by encrypting an arbitrary context-message pair  $((K, N, A), M)$ . Then, it computes the states  $\bar{Y}$  and  $\bar{S}$  shown in Fig. 5 for  $(C, T)$ , a fresh key  $\bar{K}$ , and new associated data blocks  $\bar{A}_2, \dots, \bar{A}_{\bar{\alpha}}$ . It sets the remaining associated data block  $\bar{A}_1$  to the XOR of  $\bar{Y}$  and  $\bar{S}$ . Finally,  $\mathcal{A}$  extracts  $(\bar{N}, \bar{A})$  from the associated data blocks, computes  $\bar{M} \leftarrow \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, C)$  and outputs  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}))$ .  $\square$

<sup>17</sup> Note that  $\mu$  and  $\gamma$  might not be the same.

### 3.2 ROMULUS

ROMULUS [39, 40] is an authenticated encryption scheme based on tweakable block-ciphers. For the concrete instantiation of the TBC, they use SKINNY [8] and the authenticated encryption mode bears similarities with COFB [19]. ROMULUS comes in three different variants ROMULUS-N, ROMULUS-M, and ROMULUS-T. The former is the main candidate while the other two are designed with additional security guarantees in mind: ROMULUS-M achieves security against nonce-misuse while ROMULUS-T is designed to maintain security even in the presence of side-channel leakage. Throughout this work we only consider the main variant ROMULUS-N, which we simply refer to as ROMULUS.

**Description of ROMULUS.** The pseudocode of ROMULUS is given in Fig. 22. The scheme is further illustrated in Fig. 6. It follows the CpP approach, i.e., it first computes  $S \leftarrow \text{ENC}_e(K, N, A)$  and afterwards  $(C, T) \leftarrow \text{ENC}_M(K, N, S, M)$ . Both  $\text{ENC}_e$  and  $\text{ENC}_M$  apply the tweakable block-cipher and the state-update-function  $\xi$  in an alternating manner. The state-update-function

$$\xi : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n \times \{0, 1\}^n, \quad \xi(S, I) = (S \oplus I, G(S) \oplus I)$$

is an important component of ROMULUS and the matrix  $G$  it utilizes, is

$$G = \begin{pmatrix} G_s & 0 & 0 & \cdots & 0 \\ 0 & G_s & 0 & \cdots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \cdots & 0 & G_s & 0 \\ 0 & \cdots & 0 & 0 & G_s \end{pmatrix}, \quad \text{where } G_s = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

**Committing Attack Against ROMULUS.** We show that ROMULUS is insecure with respect to committing security. The attack is stated in the following theorem.

**Theorem 4.** *Consider ROMULUS illustrated in Fig. 6. Let TBC be modeled as an ideal tweakable cipher  $\tilde{E}$ . Then there exists an adversary  $\mathcal{A}$ , making  $q$  queries to  $\tilde{E}$ , such that*

$$\text{Adv}_{\text{ROMULUS}}^{\text{CMT}}(\mathcal{A}) = 1,$$

where  $q = 2\mu + \lfloor \frac{\alpha}{2} \rfloor + \lfloor \frac{\bar{\alpha}}{2} \rfloor + 2$ . Here,  $\mu$  is the number of blocks for the message,  $\alpha$  is the number of blocks for the associated data of the first tuple, and  $\bar{\alpha}$  is the number of blocks for the second tuple that  $\mathcal{A}$  outputs.

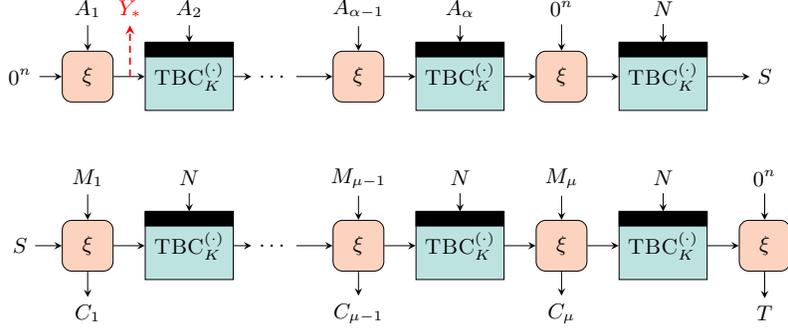


Fig. 6: Illustration of ROMULUS (for  $\alpha$  an even number) in terms of  $\text{ENC}_e$  (top) and  $\text{ENC}_M$  (bottom). The values that are input from the top into  $\text{TBC}_K^{(\cdot)}$  are used as tweaks (we drop the counters making the tweaks unique for simplicity). The state  $Y_*$ , marked in red, is used in our CMT attack.

*Proof (sketch).* The full proof is given in [Appendix B.2](#), here we provide a sketch. A key ingredient of the proof is to show that the state-update-function is invertible with respect to both the output and the state. Exploiting this, the following committing attack is possible. The adversary  $\mathcal{A}$  generates a ciphertext  $(C, T)$  by encrypting an arbitrary context-message pair  $((K, N, A), M)$ . It then inverts  $\text{ENC}_M$  for the ciphertext  $(C, T)$ , a new key  $\bar{K}$ , and new nonce  $\bar{N}$ , by inverting the state-update-function and the block-cipher one by one, resulting in some state  $S$ . Finally, it inverts  $\text{ENC}_e$  on  $S$  (using  $\bar{K}$ ,  $\bar{N}$ , and new associated data blocks  $\bar{A}_2, \dots, \bar{A}_\alpha$ ) up to the state  $Y_*$  (cf. [Fig. 6](#)). Setting the first associated data block  $\bar{A}_1$  to  $Y_*$  then guarantees that  $\mathcal{A}$  wins the game CMT by outputting  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}))$ .  $\square$

### 3.3 GIFT-COFB

The AE scheme GIFT-COFB uses the COFB mode [19] for authenticated encryption and instantiates the block-cipher using GIFT [4].

**Description of GIFT-COFB.** The pseudocode of GIFT-COFB is given in [Fig. 24](#) while [Fig. 7](#) provides an illustration of it. GIFT-COFB follows the CpP approach, i.e., it first computes  $S \leftarrow \text{ENC}_e(K, N, A)$  and afterwards  $(C, T) \leftarrow \text{ENC}_M(K, S, M)$ . The scheme features the state-update-function

$$\xi : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n \times \{0, 1\}^n, \quad \xi(S, I) = (\tilde{G}(S) \oplus I, S \oplus I)$$

that is invoked in an alternating manner with the block-cipher. This function makes use of the following matrix  $\tilde{G}$ , which gets a bit string of length  $n$ , swaps the two halves, and additionally applies a bit rotation to the (new) second half:

$$\tilde{G}: \{0, 1\}^{n/2} \times \{0, 1\}^{n/2} \rightarrow \{0, 1\}^{n/2} \times \{0, 1\}^{n/2}, \quad (S_1, S_2) \mapsto (S_2, S_1 \lll 1).$$

In between the state-update-function and the block-cipher, GIFT-COFB applies some masking by XORing some value  $\Delta$  to the state.

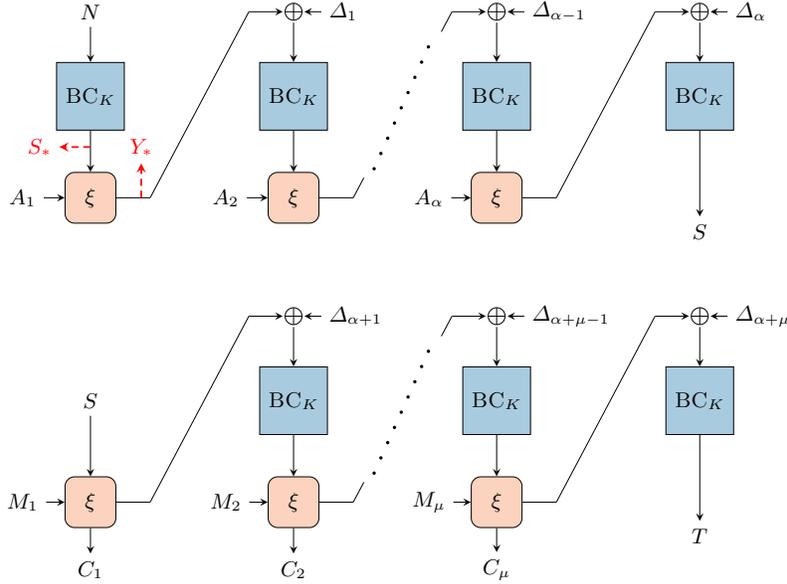


Fig. 7: Illustration of GIFT-COFB in terms of  $\text{ENC}_e$  (top) and  $\text{ENC}_M$  (bottom). The different indices for  $\Delta$  only indicate that the values are different—in the pseudocode the value  $\Delta$  is constantly updated. The states  $S_*$  and  $Y_*$ , marked in red, are used in our CMT attack.

**Committing Attack Against GIFT-COFB.** The AE scheme GIFT-COFB does not achieve committing security. The scheme is very similar to ROMULUS which allows to apply the same attack strategy. However, GIFT-COFB uses a different state-update-function  $\xi$  than ROMULUS. It turns out that we cannot always invert  $\xi$ —inversion only works with probability  $\frac{1}{2}$ . This is the reason, why the advantage depends on the number of ciphertext blocks.

**Theorem 5.** *Consider GIFT-COFB illustrated in Fig. 7. Let BC be modeled as an ideal cipher E. Then there exists an adversary  $\mathcal{A}$ , making  $q$  queries to E, such that*

$$\text{Adv}_{\text{GIFT-COFB}}^{\text{CMT}}(\mathcal{A}) = \frac{1}{2^\mu},$$

where  $q = 2\mu + \alpha + \bar{\alpha} + 2$ . Here,  $\mu$  is the number of message blocks,  $\alpha$  is the number of associated data blocks for the first tuple, and  $\bar{\alpha}$  for the second tuple that  $\mathcal{A}$  outputs.

*Proof (sketch).* The full proof is given in [Appendix B.3](#), here we provide a sketch. The gist is the same as for ROMULUS: generate a ciphertext, compute  $S_*$  and  $Y_*$  (cf. [Fig. 7](#)) for a new context, and set the associated data block  $\bar{A}_1$  accordingly. An important difference is that inverting  $\xi$  in GIFT-COFB works only with probability  $\frac{1}{2}$ . Another detail that has to be considered is the usage of the correct masking values.  $\square$

### 3.4 PHOTON-BEETLE

The authenticated encryption scheme PHOTON-BEETLE [5] is a (duplex) sponge-based AE scheme. It uses the PHOTON permutation [36] as the underlying permutation and the BEETLE mode of operation [18]. In contrast to the plain duplex, the BEETLE mode uses a feedback function to determine the next input to the underlying permutation of the sponge.

**Description of PHOTON-BEETLE.** PHOTON-BEETLE is described in [Fig. 26](#) and illustrated in [Fig. 8](#). It is a CpP scheme, i.e., it processes the context  $(K, N, A)$  via the function  $\text{ENC}_e$  and subsequently processes the message together with the output of  $\text{ENC}_e$ —an important property is that no part of the context is input to  $\text{ENC}_M$ . In  $\text{ENC}_M$ , the permutation and the state-update-function are applied in an alternating fashion. We omit the description of the latter, as our CMT attack is independent of it.

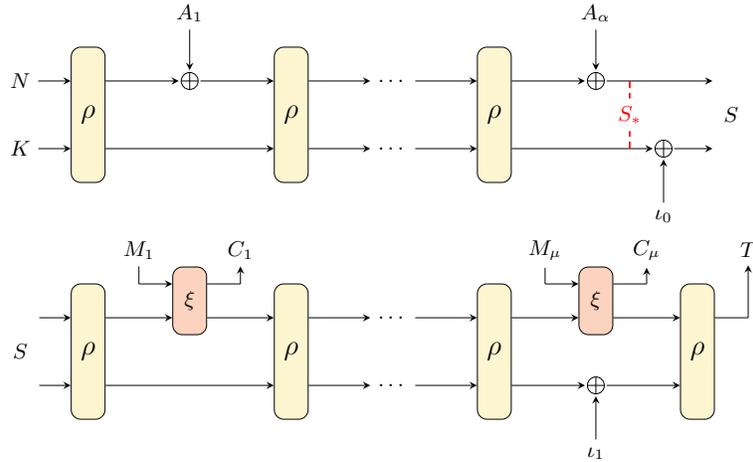


Fig. 8: Illustration of PHOTON-BEETLE in terms of  $\text{ENC}_e$  (top) and  $\text{ENC}_M$  (bottom). The state  $S_*$ , marked in red, is used in our CMT attack.

**Committing Attack Against PHOTON-BEETLE.** In this section, we show that PHOTON-BEETLE does not achieve committing security. The CMT attack is given in [Theorem 6](#) below.

**Theorem 6.** *Consider PHOTON-BEETLE illustrated in [Fig. 8](#). Let  $\rho$  be modeled as an ideal permutation. Then there exists an adversary  $\mathcal{A}$ , making  $q$  queries to  $\rho$ , such that*

$$\text{Adv}_{\text{PHOTON-BEETLE}}^{\text{CMT}}(\mathcal{A}) = 1,$$

where  $q = \alpha + \bar{\alpha}$ . Here,  $\alpha$  is the number of blocks for the associated data of the first tuple, and  $\bar{\alpha}$  is the number of blocks for the second tuple that  $\mathcal{A}$  outputs.

*Proof (sketch).* The full proof is given in [Appendix B.4](#), here we provide a sketch. The attack targets the state right before applying the domain separation at the end of  $\text{ENC}_e$  (the state  $S_*$  marked in red in [Fig. 8](#)). Since the initial state of PHOTON-BEETLE is entirely dependent on the context, we can take an arbitrary  $S_*$ , invert it for different associated data  $A$  and  $\bar{A}$ , and take the results as the key-nonce pairs  $(K, N)$  and  $(\bar{K}, \bar{N})$ . Since  $\text{ENC}_M$  only depends on the outcome of  $\text{ENC}_e$  and the message  $M$ ,  $\mathcal{A}$  wins by outputting  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, M)$  for an arbitrary message  $M$ .  $\square$

### 3.5 XOODYAK

The authenticated encryption scheme XOODYAK [23] is an AE scheme based on a full-state keyed duplex. XOODYAK uses the XOODOO permutation [22] as the underlying permutation and the CYCLIST mode of operation. The latter was introduced as part of the XOODYAK specification and is an adaption of the KEYAK mode [14] to the lightweight setting.

**Description of XOODYAK.** The pseudocode of XOODYAK is given in [Fig. 27](#) and further illustration is provided in [Fig. 9](#). XOODYAK is a CpP scheme, i.e., first  $S \leftarrow \text{ENC}_e(K, N, A)$  is computed, followed by the computation of ciphertext and tag as  $(C, T) \leftarrow \text{ENC}_M(S, M)$ . In  $\text{ENC}_e$ , the inputs are XORed onto the full-state (note that the last 32 bits are reserved for padding). In contrast to this,  $\text{ENC}_M$  uses a rate of 192 bits for the computation of ciphertext and tag.

XOODYAK exhibits a form of padding, that is used by none of the other NIST candidates and hence will be described shortly in the following: For a bit string  $X$  of length at most 352 and  $p \in \{0, 1\}^8$  define

$$\text{pad}_e(X, p) = X \parallel 00000001 \parallel 0^{368-|X|} \parallel p,$$

which is used for padding the context blocks. Further, for  $M \in \{0, 1\}^{\leq 192}$ , we define  $\text{pad}_M(M) = M \parallel 00000001$ , which will be used to pad the message blocks.

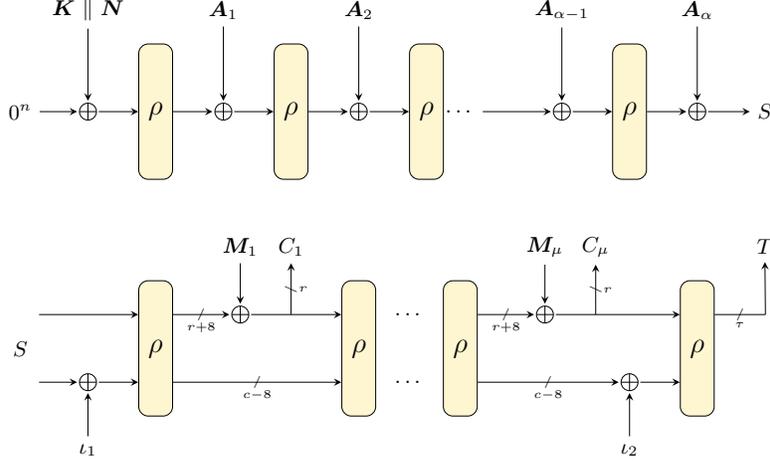


Fig. 9: Illustration of XOODYAK in terms of  $\text{ENC}_e$  (top) and  $\text{ENC}_M$  (bottom). Here,  $K \parallel N = \text{pad}_e((K \parallel N \parallel \text{enc}_8(|N|), 00000010))$ ,  $M_i = \text{pad}_M(M_i)$ ,  $A_1 = \text{pad}_e(A_1, 00000011)$ , and for  $i = 2, \dots, \alpha$ ,  $A_i = \text{pad}_e(A_i, 0^8)$ . In  $\text{ENC}_M$ , the increased rate  $(r + 8)$  is required for the padding  $\text{pad}_M$  which appends 8 bits to the message blocks.

**Committing Attack Against XOODYAK.** We show that XOODYAK does not achieve committing security. The attack is stated in the following theorem.

**Theorem 7.** *Consider XOODYAK illustrated in Fig. 9 Let  $\rho$  be modeled as an ideal permutation. Then there exists an adversary  $\mathcal{A}$  that makes  $q = 2^{17} + 1$  queries to  $\rho$  and fulfills*

$$\text{Adv}_{\text{XOODYAK}}^{\text{CMT}}(\mathcal{A}) \geq \frac{1}{2}.$$

*Proof (sketch).* The full proof is given in Appendix B.5, here we provide a sketch. The attack exploits the fact that  $\text{ENC}_M$  only depends on the output of  $\text{ENC}_e$  and the message, i.e., two different contexts  $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$  with  $\text{ENC}_e(K, N, A) = \text{ENC}_e(\bar{K}, \bar{N}, \bar{A})$  easily yield a committing attack by outputting  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, M)$  for an arbitrary  $M$ . Though XOODYAK is a full-state sponge, one cannot simply choose contexts yielding a collision, as the last 32 bits are reserved for padding. With a birthday attack, the adversary can find a collision on the reserved bits and then choose the context accordingly.  $\square$

### 3.6 TINYJAMBU

TINYJAMBU [56] is a block-cipher-based authenticated encryption scheme. The specification introduces the TINYJAMBU mode, which is a lightweight variant of the JAMBU mode [55]. The latter was part of the CAESAR competition [11]. For the permutation underlying TINYJAMBU, a keyed permutation based on non-linear feedback shift registers is defined.

**Description of TINYJAMBU.** The pseudocode of TINYJAMBU is given in Fig. 28 and an illustration of the scheme can be found in Fig. 10. TINYJAMBU follows the CpP approach, i.e., it first processes the context  $(K, N, A)$  via the function  $\text{ENC}_E$  and then passes the output on to  $\text{ENC}_M$ , where it is processed together with the message. Note that, as for PHOTON-BEETLE, no part of the context is directly given as input into  $\text{ENC}_M$ . TINYJAMBU uses two keyed permutations  $\text{BC}_1$  and  $\text{BC}_2$ , both based on the same keyed permutation that is applied 640 and 1024 times for  $\text{BC}_1$  and  $\text{BC}_2$ , respectively.

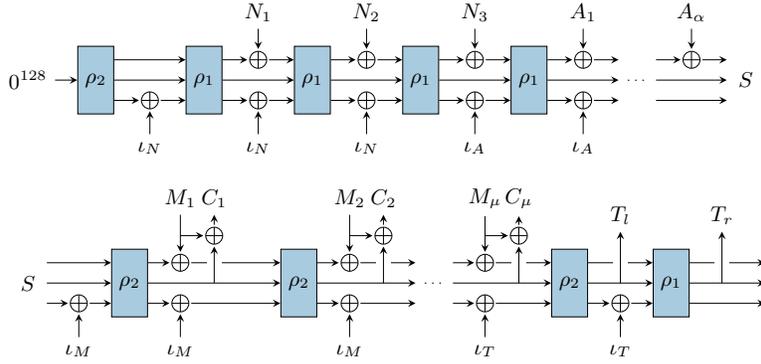


Fig. 10: Illustration of TINYJAMBU in terms of  $\text{ENC}_E$  (top) and  $\text{ENC}_M$  (bottom), where  $\rho_1 = \text{BC}_1(K, \cdot)$  and  $\rho_2 = \text{BC}_2(K, \cdot)$ .

**Committing Attack Against TINYJAMBU.** In this section, we show that TINYJAMBU does not achieve CMT security. The attack exploits the short tag length of 64 bits in TINYJAMBU, which enables an efficient deployment of the birthday bound. In the security proof of TINYJAMBU (see [56, Section 6]), this setting is modeled with only one permutation  $\rho$ . We adopt the same for the TINYJAMBU attack given in the following.

**Theorem 8.** *Consider TINYJAMBU illustrated in Fig. 10. Let  $\text{BC}_1$  and  $\text{BC}_2$  be modeled as an ideal cipher  $E$ . Then there exists an adversary  $A$  that makes  $q$  queries to  $E$  such that*

$$\text{Adv}_{\text{TINYJAMBU}}^{\text{CMT}}(A) \geq \frac{3}{8}.$$

Here,  $q = 2(2^{32} + 1)(6 + \alpha + \mu)$  for  $\alpha$  and  $\mu$  the number of associated data and message blocks, respectively, that  $A$  outputs.

*Proof (sketch).* The full proof is given in Appendix B.6, here we provide a sketch. The core idea is to apply a birthday attack against the tag, which requires about  $2^{32}$  queries due to the tag length of 64 bits. Some extra care is needed to ensure

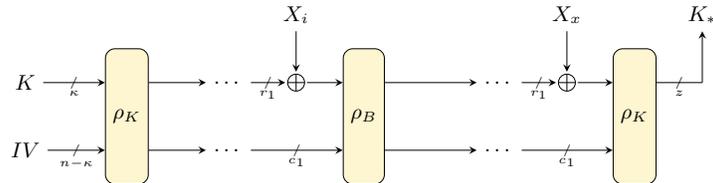


Fig. 11: Illustration of ISAP.RK.

that also the actual ciphertexts (and not just the tags) collide, as a tag collision is not sufficient for a CMT attack. Since the ciphertext is the XOR of the message and a keystream, this poses no hindrance to the attack, as the adversary can simply choose the messages in a way that yields the same ciphertext.  $\square$

### 3.7 ISAP

The authenticated encryption scheme ISAP [28–30] is a sponge-based scheme designed to withstand side-channel leakage. It features a re-keying approach that guarantees that for each input a different session key is used. The re-keying function uses a small rate to prevent adversaries from obtaining too much leakage.

**Description of ISAP.** The pseudocode of ISAP is given in Fig. 30 and further illustration is provided in Fig. 12. ISAP follows the EtM approach, i.e., first the message is encrypted via  $\text{ENC}_{\mathcal{M}}$  resulting in a ciphertext  $C$  and afterwards, the tag  $T$  is computed using  $\text{ENC}_{\mathcal{T}}$ , which processes the context and the ciphertext. Both  $\text{ENC}_{\mathcal{M}}$  and  $\text{ENC}_{\mathcal{T}}$  internally uses the re-keying function ISAP.RK to derive the session key.

In ISAP, the underlying permutation is applied several times between two absorptions, the precise number depending on the position in the ISAP sponge. For ISAP.RK, for instance, more rounds are applied when the key is processed and when the tag is generated, while fewer rounds are used in between. In summary, ISAP uses four permutations  $\rho_K$ ,  $\rho_H$ ,  $\rho_B$ , and  $\rho_E$ , each based on the same permutation but applied a different number of times.

**Committing Security of ISAP.** We show that ISAP achieves CMT security. In the security proof for ISAP [28], the two permutations used in ISAP.RK (namely  $\rho_K$  and  $\rho_B$ ) are modeled as one permutation. We adopt this for our proof and denote the permutation used in ISAP.RK by  $\rho_1$  and the one used in  $\text{ENC}_{\mathcal{T}}$  by  $\rho_2$ .<sup>18</sup> In conformity with this, the rate in ISAP.RK is denoted by  $r_1$  and the rate in  $\text{ENC}_{\mathcal{T}}$  by  $r_2$ .

Further, we consider a slightly different domain separation in  $\text{ENC}_{\mathcal{T}}$ , by XOR-ing  $1 \parallel 0^*$  instead of  $0^* \parallel 1$ . This neither influences the security of ISAP nor

<sup>18</sup> The proof is independent of  $\text{ENC}_{\mathcal{M}}$  which is why we do not need a third permutation.

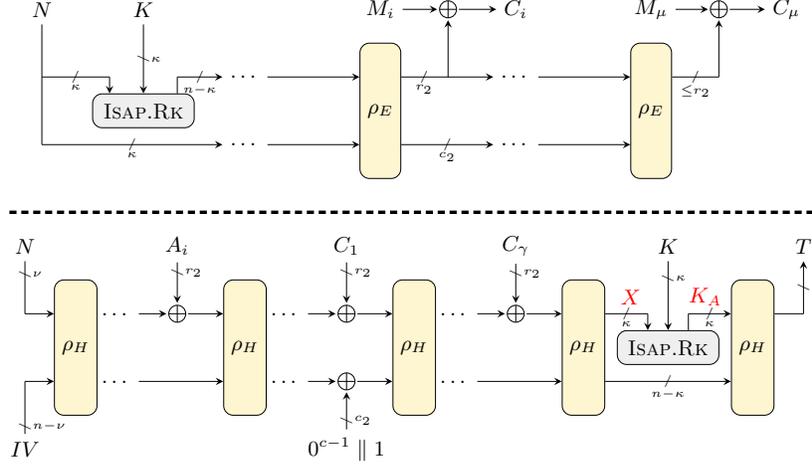


Fig. 12: Illustration of ISAP in terms of  $\text{ENC}_{\mathcal{M}}$  (top) and  $\text{ENC}_{\mathcal{T}}$  (bottom), which both rely on the re-keying function ISAP.RK. The values  $X$  and  $K_A$ , marked in red, are used in our CMT proof.

interferes with the purpose of the domain separation but it allows us to view  $\text{ENC}_{\mathcal{T}}$  as a sponge with increased rate of  $r + 1$ .<sup>19</sup>

Lastly, note that the scheme comes in two variants, using either ASCON-P or KECCAK-P as a permutation. The theorem given below holds for both instances, considering for both of them the main parameter sets as described in Table 3. Since for all parameter sets under consideration the key-length equals the tag-length ( $\kappa = \tau = 128$ ), we exclusively use the variable  $\kappa$  in the following proof.

**Theorem 9.** *Consider ISAP illustrated in Fig. 12. Let  $\rho_1$  and  $\rho_2$  be modeled by ideal permutations  $\rho_1$  and  $\rho_2$ , respectively. Then for any adversary  $\mathcal{A}$  making  $q_1$  and  $q_2$  queries to  $\rho_1$  and  $\rho_2$ , respectively, it holds that*

$$\text{Adv}_{\text{ISAP}}^{\text{CMT}}(\mathcal{A}) \leq \frac{q_1(q_1 - 1)}{2^\kappa} + \frac{q_1(q_1 + 1)}{2^{n-\kappa}} + \frac{q_2(q_2 - 1)}{2^\kappa} + \frac{q_2(q_2 + 1)}{2^{n-\max\{\kappa, r_2+1\}}}.$$

*Proof (sketch).* The full proof is given in Appendix B.7, here we provide a sketch. The idea is to view  $\text{ENC}_{\mathcal{T}}$  as a sponge-based hash function and leverage Theorem 15 to upper bound the advantage of finding two distinct contexts and a ciphertext that result in the same tag. For this, we adapt the rate in  $\text{ENC}_{\mathcal{T}}$  in a way, that allows us to absorb the nonce as the first input. With this technical trick, we obtain sponges with constant initial states at the price of a decreased capacity. Further, we model the usage of ISAP.RK in  $\text{ENC}_{\mathcal{T}}$  by a suitable XOR-operation with input denoted by  $Z$ . The proof is composed of two parts: Firstly,

<sup>19</sup> The same argument was also used for the sponge-based AE scheme SLAE [26] and is also mentioned for XOODYAK in [23, Section 4.2.1].

we consider the case that at least one of the inputs (nonce, associated data, and  $Z$ —from the replacement of ISAP.RK) differs for the two tuples the adversary outputs. Then, [Theorem 15](#) provides an upper bound for the probability that the hash values, i.e., the tags, agree. Secondly, we have to consider a special case, in which all of the inputs (nonce, associated data, and  $Z$ ) are the same across both tuples.<sup>20</sup> Such inputs do not constitute a collision of  $\text{ENC}_{\mathcal{T}}$  and hence the above argument does not apply. However, in this case we can show that the committing adversary has found a collision in ISAP.RK. The latter can also be viewed as a sponge-based hash function (by modifying the rate) and a second application of [Theorem 15](#) finishes the proof.  $\square$

### 3.8 ASCON

ASCON [31] is a sponge-based authenticated encryption scheme. The scheme was chosen as the primary candidate for lightweight applications in the CAESAR competition. Furthermore, ASCON was selected to be standardized as part of the NIST LWC standardization process. As part of the CAESAR competition and the NIST LWC standardization process, ASCON enjoys a long line of research, in particular, with respect to the underlying permutation ASCON-P. For the authenticated encryption mode, no formal security analysis existed until recently, when Lefevre and Mennink [42] gave the first security proof for ASCON.<sup>21</sup>

**Description of ASCON.** The pseudocode of ASCON is given in [Fig. 32](#) and further illustration is provided in [Fig. 13](#). Similar to the other schemes, ASCON can be viewed as a CpP scheme which first processes the context using  $\text{ENC}_{\mathcal{C}}$  before the message is processed using  $\text{ENC}_{\mathcal{M}}$ . A core feature of ASCON is that at the very start (first permutation of  $\text{ENC}_{\mathcal{C}}$ ) and the very end (last permutation of  $\text{ENC}_{\mathcal{M}}$ ), it uses more rounds of the underlying permutation for security ( $\rho^a$  and  $\rho^b$  for  $a = 12$  and  $b = 6$ ). Note that ASCON XORs the key three additional times: After the first permutation as well as before and after the last permutation. We call the former two instances state-blinding and the latter output-blinding.

**Committing Security of ASCON.** We show that ASCON achieves CMT security. We model the two permutations  $\rho^a$  and  $\rho^b$  by one ideal permutation  $\rho$ , as we did in the committing security proof of ISAP for the re-keying function. Further, we consider a slightly different order of inputs at two points in the ASCON encryption. Firstly, the initial state is changed by moving the initialization vector from the beginning of the state to the end. Secondly, the state-blinding is changed so that it affects the first bits of the inner state rather than the last bits.<sup>22</sup> Note that these are cosmetic changes. They do not influence the overall

<sup>20</sup> The careful reader might notice that in this case the CMT adversary outputs two tuples that differ solely in the keys but result in the same session keys.

<sup>21</sup> While an earlier work [41] argued that their proof covers ASCON, they actually only show security for a simplified version (namely without the state-/output-blinding).

<sup>22</sup> This only affects the first state-blinding; the second one is already of that form.

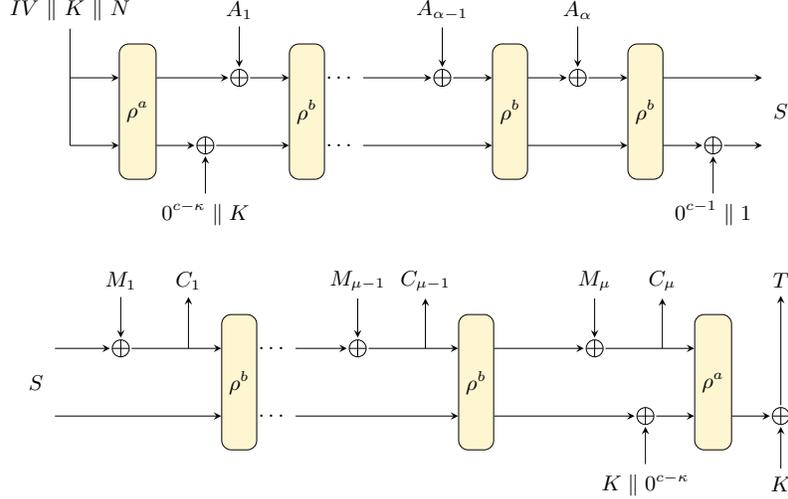


Fig. 13: Illustration of ASCON in terms of  $\text{ENC}_E$  (top) and  $\text{ENC}_M$  (bottom).

security of ASCON but greatly simplify our proof, as we can capture the state-blinding by considering a larger rate. Lastly, for sake of simplicity, we drop the domain separation of ASCON in the proof. It can, however, be easily incorporated by moving the domain separation to the first bits of the capacity, as we already did for ISAP.

Similarly to ISAP, we could show committing security by leveraging [Theorem 15](#). However, for ASCON this would yield only about 32-bit committing security. This is due to the short initialization vector of 64 bits, which corresponds to a 64-bit capacity. By recasting the proof of [Theorem 15](#), we can show that this limitation is not inherent, i.e., a smaller capacity in just the first round does not fundamentally affect the quality of the bound.

**Theorem 10.** *Consider ASCON illustrated in [Fig. 13](#). Let  $\rho^a$  and  $\rho^b$  be modeled as a random permutation  $\rho$ . Then for any adversary  $\mathcal{A}$  making  $q \leq 2^{127}$  queries to  $\rho$ , it holds that*

$$\mathbf{Adv}_{\text{ASCON}}^{\text{CMT}}(\mathcal{A}) \leq 1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right) + \frac{q}{2^{63}} + \frac{q(q-1)}{2^{128}}.$$

*Proof (sketch).* The full proof is given in [Appendix B.8](#), here we provide a sketch. The main idea is to adapt the proof of [Theorem 15](#) to our particular setting. While we structurally follow the original proof, several changes are necessary and a more fine-grained analysis gives us a better bound.

The proof of [Theorem 15](#) visualizes the attack via a graph, built from the adversary's queries. It defines colliding paths, which correspond to collisions in the hash function, and so-called problematic paths. The collision resistance is upper bounded by the probability of finding either (1) colliding paths that are not problematic or (2) finding problematic paths.

Our proof focuses on bounding the probability of finding colliding tags—which directly yields a bound on the CMT security of ASCON. We define paths in the adversary’s graph that correspond to ASCON evaluations and model tag collisions by introducing the notion of colliding ASCON paths. Furthermore, we define problematic ASCON paths and, using both the state- and output-blinding deployed in ASCON, we give a bound on finding colliding ASCON paths that are not problematic. It is left to bound the probability of finding problematic paths. The gist of this part is to view ASCON paths as plain sponge paths, which comes at the cost of an increased rate—for the first round of absorption, the rate is increased to 256 while for the rest it is increased to 192. This allows us to understand the state blinding as part of the normal absorption into the rate part of the sponge. We then show that problematic ASCON paths translate to problematic plain sponge paths and continue by giving a bound on finding the latter. A careful analysis shows that the first absorption step is treated differently: The adversary has to hit the initialization vector with a query, whereas for the other absorption steps, it has to find a collision. Thus, despite the very large rate in the first round, we are still able to achieve about 64-bit of committing security.  $\square$

### 3.9 SCHWAEMM

SCHWAEMM [6, 7] is a sponge-based AE scheme. The permutation used to instantiate SCHWAEMM is SPARKLE which is inspired by the block-cipher SPARX [27]. The authentication mode is a variant of the BEETLE mode [18].

**Description of SCHWAEMM.** The pseudocode of SCHWAEMM is given in Fig. 35 and further illustration is provided in Fig. 14. SCHWAEMM follows the CpP approach, i.e., first the context is processed resulting in  $S \leftarrow \text{ENC}_E(K, N, A)$  and afterwards the ciphertext is computed as  $(C, T) \leftarrow \text{ENC}_M(K, S, M)$ . In SCHWAEMM the underlying permutation  $\rho$  is applied a varying number of times, depending on the position in the sponge ( $\rho^a$  and  $\rho^b$  for  $a = 11$  and  $b = 7$ ), similar to ISAP and ASCON.

Like some of the other schemes, SCHWAEMM features a state-update-function that is defined as follows:

$$\begin{aligned} \xi: \{0, 1\}^r \times \{0, 1\}^r &\rightarrow \{0, 1\}^r \times \{0, 1\}^r, \\ (S, I) &\mapsto (\xi_1(S, I), \xi_2(S, I)) = (\text{FeistelSwap}(S) \oplus I, S \oplus I). \end{aligned}$$

for  $\text{FeistelSwap} : \{0, 1\}^r \rightarrow \{0, 1\}^r$ ,  $\text{FeistelSwap}(S) = S_2 \parallel (S_2 \oplus S_1)$  with  $S_1 = \lceil S \rceil_{\frac{r}{2}}$  and  $S_2 = \lfloor S \rfloor_{\frac{r}{2}}$ . Furthermore, SCHWAEMM deploys a so-called rate-whitening function, given by

$$\omega_{c,r}: \{0, 1\}^c \rightarrow \{0, 1\}^r, \quad \omega_{c,r}(I) = (I_1, I_2, I_1, I_2) \text{ for } I_1 = \lceil I \rceil_{\frac{c}{2}}, I_2 = \lfloor I \rfloor_{\frac{c}{2}}.$$

In each round, it is applied between the state-update-function and the permutation. After the final permutation, the last  $\kappa$  bits are XORed with the key to yield the tag. As for ASCON, we refer to this as output-blinding.

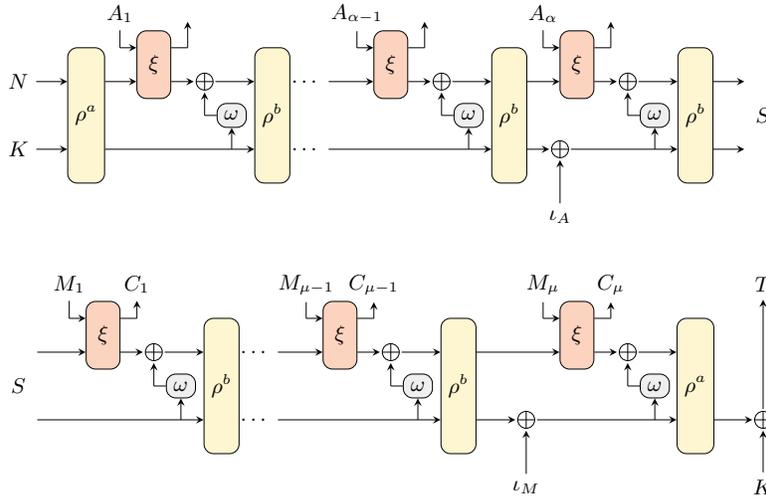


Fig. 14: Illustration of SCHWAEMM in terms of  $ENC_e$  (top) and  $ENC_M$  (bottom).

**Committing Security of SCHWAEMM.** We show that SCHWAEMM achieves committing security. At the first glance, it looks like one can apply the same attack used against PHOTON-BEETLE: Invert  $ENC_e$  for some  $S$  and take the result as the concatenation of key and nonce. However, SCHWAEMM deploys output-blinding in  $ENC_M$  (the last XOR of the key in Fig. 14), that makes the attack unlikely to succeed. Output-blinding is a feature we have also encountered in ASCON, as one of the central properties making it committing secure. Despite that, we cannot show committing security in the same way, as SCHWAEMM lacks the state-blinding, that ASCON has, and SCHWAEMM’s initial state does not contain a fixed component. However, we noticed that introducing an  $IV$  to SCHWAEMM’s initial state suffices to obtain a committing secure scheme—despite the weaker blinding mechanism. More precisely, we decrease the length of the nonce from 256 to 128 bits<sup>23</sup> and instead incorporate a fixed  $IV$  (of length 128 bit) into the initial state. For the resulting modified version of SCHWAEMM, denoted by  $SCHWAEMM_{IV}$ , we can show about 64-bit committing security. For the proof, we model the two permutations  $\rho^a$  and  $\rho^b$  by one ideal permutation  $\rho$ , as it was done for ISAP and ASCON. We further drop the domain separation in our proof for sake of simplicity. This part can easily be incorporated at the cost of reducing the committing security by the number of bits required for the domain separation.

**Theorem 11.** Consider SCHWAEMM illustrated in Fig. 14 and  $SCHWAEMM_{IV}$ , its modified version described above. Let  $\rho^a$  and  $\rho^b$  be modeled as a random permutation  $\rho$ . Then for any adversary  $\mathcal{A}$  making  $q \leq 2^{127}$  queries to  $\rho$ , it holds

<sup>23</sup> Note that the modified scheme is still in accordance to the NIST requirements [48] that nonces are at least 96 bits long.

that

$$\mathbf{Adv}_{\text{SCHWAEMM}_{IV}}^{\text{CMT}}(\mathcal{A}) \leq 1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right) + \epsilon,$$

$$\text{for } \epsilon > \frac{(1-2^{-256})q^2 + (1+2^{-256})q}{2^{129}}.$$

*Proof (Sketch).* The full proof is given in [Appendix B.9](#), here we provide a sketch. The idea is to view  $\text{SCHWAEMM}_{IV}$  as a plain sponge and show that it is hard to find colliding tags, which, in turn, yields that it is hard to break committing security. Due to the output-blinding, tag collisions are not equal to collisions of the plain sponge, as the keys—used to blind the output—can be different. To deal with that, we define shifted collisions, which encompass these tag collisions for the plain sponge. We conclude the proof by utilizing the indistinguishability of the plain sponge from a random function and providing a bound on finding shifted collisions for the latter.  $\square$

## 4 Conclusion

Out of the nine considered NIST finalists, we have shown that six do not achieve committing security while the remaining three do. For the former, we gave concrete attacks, while the others are backed up by formal security proofs. For  $\text{ELEPHANT}$ ,  $\text{GIFT-COFB}$ , and  $\text{ROMULUS}$ , the attacks can be traced back to the fact that inputs are XORed onto the entire state (we call this full-state XOR). Similarly to this,  $\text{XOODYAK}$  exhibits the same property as part of it ( $\text{ENC}_{\mathbb{C}}$ ) is a full-state sponge: The whole state—except for a few bits reserved for padding—can be influenced by the adversary. Furthermore, also the attack against  $\text{PHOTON-BEETLE}$ , which exploits the scheme’s context-dependent initial state, can be regarded as a “full-state-attack”. For this, note that an initial state that can be controlled entirely by the adversary, can easily be represented as a full-state XOR of some input onto the all-zero string. Our attack against  $\text{TINYJAMBU}$  is based on the scheme’s short tag, which enables an efficient deployment of the birthday attack. Note that in order to obtain 64-bit committing security, a tag length of at least 128 bits is necessary.<sup>24</sup>

Overall, we observe that all block-cipher-based NIST finalists are broken with respect to committing security. At the same time, the three schemes for which we formally prove committing security are all sponge-based. This suggests that sponges are better suited to building committing authenticated encryption schemes, whereas block-cipher-based schemes seem to exhibit some inherent vulnerabilities. However, note that full-state sponges—which have proven to yield AE schemes that achieve confidentiality and authenticity [46]—are also vulnerable with respect to committing security. Our overall insight, i.e., that sponges are more favorable than block-ciphers when designing committing AE schemes, is also supported by prior results: While there is a number of committing attacks

<sup>24</sup> Though, longer tags are not sufficient to obtain CMT security for  $\text{TINYJAMBU}$ .

against block-cipher-based schemes [9,38,45], we are not aware of any committing attacks against sponge-based schemes. Moreover, very recently, positive results regarding committing AE from sponges emerged [24,25].

Looking at the positive results (ISAP, ASCON, and SCHWAEMM), a common feature is that a significant part of the state is “out-of-reach” for the adversary, i.e., not manipulable by the input. All of these schemes have a capacity of at least 128 bits. Interestingly, TINYJAMBU features a similar design in the sense that a part of its state (96 bits) is unaffected by the inputs. However, due to the common, but comparably small, block size of  $n = 128$ , it is not possible for TINYJAMBU to achieve 64-bit committing security. This would require the entire 128-bit state to not be affected by the inputs throughout the whole scheme, which would render it pointless. In contrast to this, sponge constructions typically rely on permutations over much larger states: ASCON-P ( $n = 320$ ), KECCAK-P ( $n = 400$ ), SPARKLE ( $n = 384$ ), XOODOO ( $n = 384$ ), and PHOTON ( $n = 256$ ). These allow for efficient constructions while maintaining a large part of the state unaffected by the inputs, which is favorable for committing security.

## References

1. Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 480–497. Springer, Heidelberg, February 2010.
2. 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.
3. Subhadeep Banik, Avik Chakraborti, Tetsu Iwata, Kazuhiko Minematsu, Mridul Nandi, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Yosuke Todo. GIFT-COFB. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
4. Subhadeep Banik, Sumit Kumar Pandey, Thomas Peyrin, Yu Sasaki, Siang Meng Sim, and Yosuke Todo. GIFT: A small present - towards reaching the limit of lightweight encryption. In Wieland Fischer and Naofumi Homma, editors, *CHES 2017*, volume 10529 of *LNCS*, pages 321–345. Springer, Heidelberg, September 2017.
5. Zhenzhen Bao, Avik Chakraborti, Nilanjan Datta, Jian Guo, Mridul Nandi, Thomas Peyrin, and Kan Yasuda. PHOTON-Beetle. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
6. Christof Beierle, Alex Biryukov, Luan Cardoso dos Santos, Johann Großschädl, Léo Perrin, Aleksei Udovenko, Vesselin Velichkov, and Qingju Wang. Lightweight AEAD and hashing using the Sparkle permutation family. *IACR Trans. Symm. Cryptol.*, 2020(S1):208–261, 2020.
7. Christof Beierle, Alex Biryukov, Luan Cardoso dos Santos, Johann Großschädl, Léo Perrin, Aleksei Udovenko, Vesselin Velichkov, Qingju Wang, Amir Moradi, and Aein Rezaei Shahmirzadi. Schwaemm and Esch. Technical report, National

- Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
8. Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi, Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The SKINNY family of block ciphers and its low-latency variant MANTIS. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 123–153. Springer, Heidelberg, August 2016.
  9. 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, Heidelberg, May / June 2022.
  10. 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, Heidelberg, December 2000.
  11. Daniel J. Bernstein. CAESAR: Competition for authenticated encryption: Security, applicability, and robustness. <https://competitions.cr.yp.to/caesar.html>, 2014.
  12. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indistinguishability of the sponge construction. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197. Springer, Heidelberg, April 2008.
  13. Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge functions. In *ECRYPT Hash Workshop*, 2007.
  14. Guido Bertoni, Joan Daemen, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Keyak v2. Technical report, Submission to the CAESAR Competition, 2016. Available at <https://keccak.team/files/Keyakv2-doc2.2.pdf>.
  15. Tim Beyne, Yu Long Chen, Christoph Dobraunig, and Bart Mennink. Dumbo, Jumbo, and Delirium: Parallel authenticated encryption for the lightweight circus. *IACR Trans. Symm. Cryptol.*, 2020(S1):5–30, 2020.
  16. 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>.
  17. Dan Boneh and Victor Shoup. *A Graduate Course in Applied Cryptography*. 2023. Draft 0.6, <http://toc.cryptobook.us/>.
  18. Avik Chakraborti, Nilanjan Datta, Mridul Nandi, and Kan Yasuda. Beetle family of lightweight and secure authenticated encryption ciphers. *IACR TCHES*, 2018(2):218–241, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/881>.
  19. Avik Chakraborti, Tetsu Iwata, Kazuhiko Minematsu, and Mridul Nandi. Blockcipher-based authenticated encryption: How small can we go? In Wieland Fischer and Naofumi Homma, editors, *CHES 2017*, volume 10529 of *LNCS*, pages 277–298. Springer, Heidelberg, September 2017.
  20. 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, Heidelberg, September 2022.
  21. Cas Cremers, Samed Düzlülü, Rune Fiedler, Marc Fischlin, and Christian Janson. BUFFing signature schemes beyond unforgeability and the case of post-quantum signatures. In *2021 IEEE Symposium on Security and Privacy*, pages 1696–1714. IEEE Computer Society Press, May 2021.

22. Joan Daemen, Seth Hoffert, Gilles Van Assche, and Ronny Van Keer. The design of Xoodoo and Xooff. *IACR Trans. Symm. Cryptol.*, 2018(4):1–38, 2018.
23. Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, Ronny Van Keer, and Silvia Mella. Xoodyak. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
24. Joan Daemen, Silvia Mella, and Gilles Van Assche. Committing authenticated encryption based on SHAKE. *IACR Cryptol. ePrint Arch.*, 2023:1494, 2023.
25. Jean Paul Degabriele, Marc Fischlin, and Jérôme Govinden. The indistinguishability of the duplex and its practical applications. In *ASIACRYPT 2023*, 2023.
26. Jean Paul Degabriele, Christian Janson, and Patrick Struck. Sponges resist leakage: The case of authenticated encryption. In Steven D. Galbraith and Shihō Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 209–240. Springer, Heidelberg, December 2019.
27. Daniel Dinu, Léo Perrin, Aleksei Udovenko, Vesselin Velichkov, Johann Großschädl, and Alex Biryukov. Design strategies for ARX with provable bounds: Sparx and LAX. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 484–513. Springer, Heidelberg, December 2016.
28. 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.
29. Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, Bart Mennink, Robert Primas, and Thomas Unterluggauer. ISAP. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
30. 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.
31. Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer. Ascon. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
32. 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, Heidelberg, August 2018.
33. Pooya Farshim, Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Robust encryption, revisited. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 352–368. Springer, Heidelberg, February / March 2013.
34. Pooya Farshim, Claudio Orlandi, and Răzvan Roşie. Security of symmetric primitives under incorrect usage of keys. *IACR Trans. Symm. Cryptol.*, 2017(1):449–473, 2017.
35. Robert Granger, Philipp Jovanovic, Bart Mennink, and Samuel Neves. Improved masking for tweakable blockciphers with applications to authenticated encryption. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 263–293. Springer, Heidelberg, May 2016.
36. Jian Guo, Thomas Peyrin, and Axel Poschmann. The PHOTON family of lightweight hash functions. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 222–239. Springer, Heidelberg, August 2011.

37. Martin Hell, Thomas Johansson, Willi Meier, Jonathan Sönnnerup, Hirotaka Yoshida, and Alexander Maximov. Grain-128AEAD. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
38. Takanori Isobe and Mostafizar Rahman. Key committing security analysis of AEGIS. *IACR Cryptol. ePrint Arch.*, 2023:1495, 2023.
39. Tetsu Iwata, Mustafa Khairallah, Kazuhiko Minematsu, and Thomas Peyrin. Duel of the titans: The Romulus and Remus families of lightweight AEAD algorithms. *IACR Trans. Symm. Cryptol.*, 2020(1):43–120, 2020.
40. 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>.
41. Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond  $2^{c/2}$  security in sponge-based authenticated encryption modes. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 85–104. Springer, Heidelberg, December 2014.
42. Charlotte Lefevre and Bart Mennink. Generic security of the Ascon mode: On the power of key blinding. *IACR Cryptol. ePrint Arch.*, 2023:796, 2023.
43. 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.
44. Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 31–46. Springer, Heidelberg, August 2002.
45. 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, Heidelberg, April 2023.
46. Bart Mennink, Reza Reyhanitabar, and Damian Vizár. Security of full-state keyed sponge and duplex: Applications to authenticated encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 465–489. Springer, Heidelberg, November / December 2015.
47. 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, Heidelberg, May 2014.
48. National Institute of Standards and Technology. Lightweight cryptography standardization process. <https://csrc.nist.gov/projects/lightweight-cryptography>, 2015.
49. National Institute of Standards and Technology. Post-quantum cryptography standardization process. <https://csrc.nist.gov/projects/post-quantum-cryptography>, 2017.
50. National Institute of Standards and Technology. Call for additional digital signature schemes for the post-quantum cryptography standardization process. <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf>, 2022.
51. Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.

52. Meltem Sönmez Turan, Kerry McKay, Donghoon Chang, Lawrence E. Bassham, Jinkeon Kang, Noah D. Waller, John M. Kelsey, and Deukjo Hong. Status report on the final round of the NIST lightweight cryptography standardization process. <https://nvlpubs.nist.gov/nistpubs/ir/2023/NIST.IR.8454.pdf>, 2023.
53. Serge Vaudenay. Security flaws induced by CBC padding - applications to SSL, IPSEC, WTLS... In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 534–546. Springer, Heidelberg, April / May 2002.
54. David Wagner. A generalized birthday problem. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 288–303. Springer, Heidelberg, August 2002.
55. Hongjun Wu and Tao Huang. The JAMBU lightweight authentication encryption mode (v2.1). Technical report, Submission to the CAESAR Competition, 2016. Available at <https://competitions.cr.yp.to/round3/jambuv21.pdf>.
56. Hongjun Wu and Tao Huang. TinyJambu. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.

## A Additional Preliminaries

In [Appendix A.1](#) we describe the used paddings and define more security notions. We give some background on (tweakable) block-ciphers and sponges in [Appendix A.2](#) and [Appendix A.3](#), respectively. Finally, in [Appendix A.4](#), we provide some results that are relevant for our attacks and proofs.

### A.1 Paddings and Security Notions

The authenticated encryption schemes considered in this work, use common paddings which we recall below. The one-zero padding  $\text{pad}_{10^*}(\cdot, r)$ , appends a 1, followed by 0s until the desired length  $r$  is reached. Simply padding with 0s to length  $r$  is denoted by  $\text{pad}_{0^*}(\cdot, r)$ . By  $\text{pad}_L(\cdot, r)$ , we denote the padding which appends 0, followed by appending the length of the input.

Below we define collision resistance of a hash function.

**Definition 12.** *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*

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

Menda et al. [45] defined several variants of committing security. These variants require different parts of the contexts to disagree (and sometimes also others to agree). Below we recall their security notions.<sup>25</sup>

**Definition 13.** *Let  $\text{AE} = (\text{ENC}, \text{DEC})$  be an authenticated encryption scheme and the games  $\text{CMT}_X$  and  $\text{CMT}_X^*$  for  $X \in \{\text{K}, \text{N}, \text{A}\}$  be defined as in [Fig. 1](#). For any adversary  $\mathcal{A}$ , its  $\text{CMT}_X$  and  $\text{CMT}_X^*$  advantages are defined as*

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

Next, we define the advantage of finding colliding tags. At its core, this is a weakened version of committing security as the ciphertexts are not required to agree. We use this to bound the committing security of ASCON and SCHWAEMM.

**Definition 14.** *Let  $\text{AE} = (\text{ENC}, \text{DEC})$  be an authenticated encryption scheme and the game  $\text{TagColl}$  be defined as in [Fig. 16](#). For any adversary  $\mathcal{A}$ , its  $\text{TagColl}$  advantage is defined as*

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

<sup>25</sup> Note, however, that  $\text{CMT}_K$  originates from [9].

<p>Game <math>\text{CMT}_K</math></p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: <b>if</b> $K = \bar{K}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$	<p>Game <math>\text{CMT}_K^*</math></p> <hr/> 1: $((K, \bar{K}), N, A, (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: <b>if</b> $K = \bar{K}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(\bar{K}, N, A, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$
<p>Game <math>\text{CMT}_N</math></p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: <b>if</b> $N = \bar{N}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$	<p>Game <math>\text{CMT}_N^*</math></p> <hr/> 1: $(K, (N, \bar{N}), A, (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: <b>if</b> $N = \bar{N}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(K, \bar{N}, A, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$
<p>Game <math>\text{CMT}_A</math></p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: <b>if</b> $A = \bar{A}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$	<p>Game <math>\text{CMT}_A^*</math></p> <hr/> 1: $(K, N, (A, \bar{A}), (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: <b>if</b> $A = \bar{A}$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(K, N, \bar{A}, \bar{M})$ 6: <b>return</b> $((C, T) = (\bar{C}, \bar{T}))$

Fig. 15: Security games  $\text{CMT}_K$ ,  $\text{CMT}_N$ ,  $\text{CMT}_A$ ,  $\text{CMT}_K^*$ ,  $\text{CMT}_N^*$ , and  $\text{CMT}_A^*$  for authenticated encryption schemes. Here,  $((K, \bar{K}), N, A, (M, \bar{M}))$  is an abbreviation for  $(K, N, A, M), (\bar{K}, N, A, \bar{M})$ , likewise used for the other context components.

<p>Game <math>\text{TagColl}</math></p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: <b>if</b> $(K, N, A) = (\bar{K}, \bar{N}, \bar{A})$ 3: <b>return</b> 0 4: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 5: $(\bar{C}, \bar{T}) \leftarrow \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ 6: <b>return</b> $(T = \bar{T})$
--

Fig. 16: Security Game  $\text{TagColl}$  for authenticated encryption schemes used in the proof of [Theorem 10](#).

## A.2 (Tweakable) Block-Ciphers

A block-cipher  $\text{BC}: \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  takes as input a key  $K$  of length  $\kappa$  and a message  $M$  of length  $n$ , and outputs a ciphertext  $C$  of the same length as the message. For every  $K \in \{0, 1\}^\kappa$ ,  $\text{BC}(K, \cdot)$  is a permutation over  $\{0, 1\}^n$ . A tweakable block-cipher [44]  $\text{TBC}: \{0, 1\}^\kappa \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  takes as input a key  $K$  of length  $\kappa$ , a tweak  $T$  (from some set of tweaks  $\mathcal{T}$ ), and a message  $M$  of length  $n$ , and outputs a ciphertext  $C$  of the same length as the message. For every pair  $(K, T) \in \{0, 1\}^\kappa \times \mathcal{T}$ ,  $\text{TBC}(K, T, \cdot)$  is a permutation over  $\{0, 1\}^n$ . We also use  $\text{TBC}^T(K, \cdot)$  as an alternative notation for  $\text{TBC}(K, T, \cdot)$ .

For our results, we model the block-ciphers  $\text{BC}$  and tweakable block-ciphers  $\text{TBC}$  by an ideal cipher  $\text{E}$  and ideal tweakable cipher  $\tilde{\text{E}}$ , respectively.



Fig. 17: Block-cipher (left) and tweakable block-cipher (right). For tweakable block-ciphers, the black bar indicates that the incoming arrow ( $T$ ) is used as a tweak.

## A.3 Sponges

Sponges [13] are a versatile tool for cryptographic primitives. Rather than just being relevant for cryptographic hash functions—as was their main design goal—they turned out to be more powerful as one can construct numerous cryptographic primitives from sponges.

The underlying component of a sponge is a permutation  $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$ . Here,  $n$  is the size of the sponge state. The sponge operates in a round-wise fashion, where each round it absorbs a part of the input and applies  $\rho$ . The rate  $r$  describes how many bits of the input can be absorbed in each round by XORing them to the first  $r$  bits of the sponge state. The higher the rate the faster the sponge as fewer rounds, hence fewer invocations of  $\rho$ , are required to absorb the input. The part of the sponge state that is not affected by the input absorption is called the inner state and its size is denoted by the capacity  $c$ , thus we have  $r + c = n$ . The capacity is related to the security of the sponge, the higher the capacity the better the security of the sponge.

We refer to sponges of the form described above by *plain sponges* and provide an illustration in Fig. 18. It was shown that—especially in the context of AE schemes—one can also deploy *full-state sponges* and *duplex sponges*. The former XORs the input to the entire state, i.e.,  $r = n$  and  $c = 0$ . The latter absorbs and squeezes in each round, in contrast to the plain sponge which squeezes only after

the absorption is finished. XOODYAK uses a full-state sponge, while a duplex sponge is used, for instance, by ASCON.

Below we recall two results for the plain sponge construction that we will use later: First, a bound on the collision resistance of a simple sponge-based hash function and, second, the indistinguishability of sponges from a random function.

**Theorem 15 ([17, Theorem 8.6]).** *Let  $\mathcal{H}$  be a hash function obtained from a permutation  $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$ , with capacity  $c$ , rate  $r$  (so  $n = r + c$ ), and output length  $w \leq r$ . For every adversary  $\mathcal{A}$ , if the number of ideal permutation queries plus the number of  $r$ -bit blocks in the output of  $\mathcal{A}$  is bounded by  $q$ , it holds that*

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

**Theorem 16 ([12, Theorem 2]).** *Let  $\mathcal{H}$  be a (padded) sponge construction obtained from a permutation  $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$ , with capacity  $c$  and rate  $r$  (so  $n = r + c$ ). Then, for any adversary  $\mathcal{A}$ , making significantly less than  $2^c$  queries to  $\rho$ ,  $\mathcal{H}$  is indistinguishable from a random oracle  $F$ , except with probability at most  $\frac{(1-2^{-256})q^2 + (1+2^{-256})q}{2^{129}}$ .*

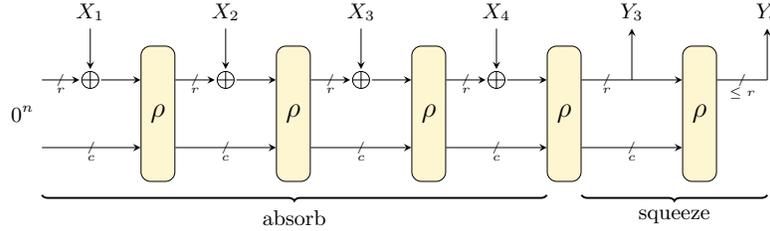


Fig. 18: Illustration of a plain sponge construction with four rounds of absorbing and two rounds of squeezing

#### A.4 Existing Results

The theorem below gives both upper and lower bounds on finding collisions for independent random variables.

**Theorem 17 ([17, Theorem B.1]).** *Let  $\mathcal{M}$  be a set of size  $n$  and  $X_1, \dots, X_k$  be  $k$  independent random variables uniform in  $\mathcal{M}$ . Let  $C$  be the event that for some distinct  $i, j \in \{1, \dots, k\}$  we have that  $X_i = X_j$ . Then*

$$\Pr[C] \geq 1 - \exp\left(\frac{-k(k-1)}{2n}\right) \geq \min\left\{\frac{-k(k-1)}{4n}, 0.63\right\} \text{ and}$$

$$\Pr[C] \leq 1 - \exp\left(\frac{-k(k-1)}{n}\right) \text{ when } k < \frac{n}{2}.$$

We use the formulation of the birthday problem presented in [54], but provide a more formal description using [Theorem 17](#).

**Lemma 18.** *Consider two lists  $L_1, L_2$  of elements drawn uniformly and independently at random from  $\{0, 1\}^\tau$ . We denote the size of  $L_1$  and  $L_2$  by  $l_1$  and  $l_2$ , respectively, and define  $l = l_1 + l_2$ . Then one finds  $x_k \in L_1$  and  $x_j \in L_2$  such that  $x_k \oplus x_j = 0$  with a probability of at least*

$$\left(1 - \exp\left(\frac{-l(l-1)}{2^{\tau+1}}\right)\right) \cdot \frac{2l_1l_2}{l^2 - l}.$$

*Proof.* Consider the concatenation of the two lists  $L = L_1 \parallel L_2$  and write  $x_1, \dots, x_l$  for its elements. Denote by  $C$  the event that  $x_k = x_j$  holds for some  $k \neq j$ . Since the  $x_i$  are drawn uniformly and independently, [Theorem 17](#) yields the bound  $\Pr[C] \geq 1 - \exp\left(\frac{-l(l-1)}{2^{\tau+1}}\right)$ . However, this probability also counts internal collisions of  $L_1$  and  $L_2$ , respectively. For two elements of  $L$ , the probability that they are not both from either  $L_1$  or  $L_2$  is  $\frac{l_1}{l} \frac{l_2}{l-1} + \frac{l_2}{l} \frac{l_1}{l-1} = \frac{2l_1l_2}{l^2 - l}$ . Taking this into account, the probability that  $x_k = x_j$  holds for some  $x_k \in L_1$  and  $x_j \in L_2$  is thus bounded above by  $\left(1 - \exp\left(\frac{-l(l-1)}{2^{\tau+1}}\right)\right) \cdot \frac{2l_1l_2}{l^2 - l}$ .  $\square$

The following lemma contains two technical results needed for the committing security proof of ASCON ([Theorem 10](#)). While the computations are not hard, we give them here to provide a complete presentation.

**Lemma 19.** *Let  $n, c \in \mathbb{N}$  such that  $c \leq n$  and  $IV \in \{0, 1\}^c$ . Let further  $\rho$  be a random permutation over  $\{0, 1\}^n$  and  $\mathcal{A}$  be an adversary making queries to  $\rho$ . Consider the following events:*

1. *Event  $E_t$  (target hitting query):*  
 $\mathcal{A}$  makes a query  $Y$  to  $\rho$  such that  $\lfloor \rho(Y) \rfloor_{64} = IV$  or  $\mathcal{A}$  makes a query  $S$  to  $\rho^{-1}$  such that  $\lfloor \rho^{-1}(S) \rfloor_{64} = IV$ .
2. *Event  $E_c$  (colliding queries):*  
 $\mathcal{A}$  makes queries  $Y \neq \bar{Y}$  to  $\rho$  such that  $\lfloor \rho(Y) \rfloor_{128} = \lfloor \rho(\bar{Y}) \rfloor_{128}$  or  $\mathcal{A}$  makes queries  $Y$  to  $\rho$  and  $\bar{S}$  to  $\rho^{-1}$  such that  $\lfloor \rho(Y) \rfloor_{128} = \lfloor \rho^{-1}(\bar{S}) \rfloor_{128}$ .

If  $\mathcal{A}$  makes  $q \leq 2^{n-1}$  queries, then

$$\Pr[E_t] \leq \frac{q}{2^{c-1}}.$$

and

$$\Pr[E_c] \leq \frac{q(q-1)}{2^c}.$$

*Proof.* We start with the bound for  $E_t$ . Let  $X_i$  be the event that  $\mathcal{A}$  triggers event  $E_t$  with its  $i$ -th query. It holds that

$$\Pr[E_t] \leq \sum_{i=1}^q \Pr[X_i] = \sum_{i=1}^q \frac{2^r}{2^n - i + 1} \leq \frac{2^r q}{2^n - q} \leq \frac{2^r q}{2^{n-1}} = \frac{q}{2^{c-1}},$$

where  $q \leq 2^{n-1}$  is used for the last inequality. Next, we bound event  $\mathbf{E}_c$ . Let  $\mathbf{X}_{ij}$  be the event that the  $j$ -th query by  $\mathcal{A}$  forms a collision with the  $i$ -th query. Then it holds that

$$\Pr[\mathbf{E}_c] \leq \sum_{j=i}^q \sum_{i=1}^{j-1} \Pr[\mathbf{X}_{ij}] \leq \sum_{j=1}^q \frac{(j-1)2^r}{2^n - j + 1} \leq 2^r \left( \frac{q(q-1)}{2(2^n - q)} \right),$$

and using  $q \leq 2^{n-1}$  again, we obtain

$$\Pr[\mathbf{E}_c] \leq 2^r \left( \frac{q(q-1)}{2^n} \right) = \frac{q(q-1)}{2^c},$$

which finishes the proof. □

## B Deferred Proofs

### B.1 Proof of Theorem 3 (ELEPHANT)

*Proof.* We construct a CMT adversary  $\mathcal{A}$  against ELEPHANT as shown in Fig. 19. As a first step it samples a key  $K$ , a nonce  $N$ , associated data  $A$ , and a message  $M$  at random from the respective sets. It computes the ciphertext  $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, M)$  and the tag  $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$ . The ciphertext is parsed into blocks  $C_1, \dots, C_\gamma \xleftarrow{n} \text{pad}_{10^*}(C, n)$ . Next, the adversary samples a second, different key  $\bar{K} \xleftarrow{s} \mathcal{K} \setminus \{K\}$  and associated data blocks  $\bar{A}_2, \dots, \bar{A}_\alpha \xleftarrow{s} \{0, 1\}^n$ .<sup>26</sup> The adversary then computes the state

$$\bar{S} \leftarrow \bigoplus_{i=2}^{\bar{\alpha}} (\tilde{\text{E}}(K, (i-1, 0), \bar{A}_i)) \oplus \bigoplus_{i=1}^{\gamma} (\tilde{\text{E}}(K, (i-1, 2), C_i)),$$

shown in Fig. 5. The value  $\bar{Y}$  is computed by querying  $\tilde{\text{E}}^{-1}$  on  $\bar{K}$ ,  $(0, 0)$ , and  $T$  (padded with 0s to length  $n$ ). Adversary  $\mathcal{A}$  then computes  $\bar{A}_1$  as the XOR of  $\bar{S}$  and  $\bar{Y}$ . Together with the other associated data blocks,  $\mathcal{A}$  computes  $(\bar{N}, \bar{A}) \leftarrow \text{pad}_{10^*}^{-1}(\bar{A}_1 \parallel \dots \parallel \bar{A}_\alpha)$ , i.e., removes the padding. It remains to compute the message  $\bar{M}$  to which the ciphertext  $C$  decrypts under the context  $(\bar{K}, \bar{N}, \bar{A})$ . This can easily be achieved by setting  $\bar{M} \leftarrow \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, C)$ . Finally,  $\mathcal{A}$  outputs  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . Observe that  $\mathcal{A}$  wins the game CMT, as we have

$$\text{ELEPHANT.ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M}) = (C, T) = \text{ELEPHANT.ENC}(K, N, A, M).$$

As for the queries to  $\tilde{\text{E}}$ ,  $\mathcal{A}$  makes  $\mu$  queries to compute  $C$  and  $\alpha + \gamma$  to compute  $T$ . Additionally,  $\mathcal{A}$  makes  $\mu$  queries to compute  $\bar{M}$  and  $\bar{\alpha} + \gamma$  queries to compute  $\bar{S}$  and  $\bar{Y}$ , totalling up to  $q = 2\mu + 2\gamma + \alpha + \bar{\alpha}$  queries.  $\square$

The gist of the attack, is finding a second associated data  $\bar{A}$  which yields the target ciphertext. The attack easily extends to a context discovery attack (CDY<sub>A</sub><sup>\*</sup>) [45]. Hence, we can conclude that ELEPHANT is also vulnerable with respect to the weaker security notions CMT<sub>K</sub> and CMT<sub>N</sub> by using [45, Corollar 3]. Furthermore, the attack can be translated to one against CMT<sub>A</sub> by observing that the adversary can choose  $\bar{A}$  to differ from  $A$  at some point (note that  $\mathcal{A}$  can freely choose all but one block). Finally, the attack is also extendable to the more restricted notion CMT<sub>A</sub><sup>\*</sup> by choosing the second key-nonce pair  $(\bar{K}, \bar{N})$  not at random but equal to the first pair  $(K, N)$ .<sup>27</sup>

<sup>26</sup> We assume that  $\mathcal{A}$  chooses the last block to exhibit a valid padding.

<sup>27</sup> In this case, the adversary needs to target a different associated data block, which grants the freedom to choose the nonce  $\bar{N}$ .

ELEPHANT adversary $\mathcal{A}$	$\mathcal{B}(C_1, \dots, C_\gamma, \bar{A}_2, \dots, \bar{A}_\alpha)$
1: $K, N, A, M \leftarrow_{\$} \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$	13: $\bar{S} \leftarrow 0^n$
2: $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, M)$	14: <b>for</b> $i = 1, \dots, \gamma$
3: $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$	15: $\bar{S} \leftarrow \bar{S} \oplus \tilde{\text{E}}^{-1}(\bar{K}, (i-1, 2), C_i)$
4: $C_1, \dots, C_\gamma \xleftarrow{n} \text{pad}_{10^*}(C, n)$	16: <b>for</b> $i = 2, \dots, \bar{\alpha}$
5: $\bar{K} \leftarrow_{\$} \mathcal{K} \setminus \{K\}$	17: $\bar{S} \leftarrow \bar{S} \oplus \tilde{\text{E}}^{-1}(\bar{K}, (i-1, 0), \bar{A}_i)$
6: $\bar{A}_2, \dots, \bar{A}_\alpha \leftarrow_{\$} \{0, 1\}^n$	18: <b>return</b> $\bar{S}$
7: $\bar{S} \leftarrow \mathcal{B}(C_1, \dots, C_\gamma, \bar{A}_2, \dots, \bar{A}_\alpha)$	
8: $\bar{Y} \leftarrow \tilde{\text{E}}^{-1}(\bar{K}, (0, 0), \text{pad}_{0^*}(T, n))$	
9: $\bar{A}_1 \leftarrow \bar{Y} \oplus \bar{S}$	
10: $(\bar{N}, \bar{A}) \leftarrow \text{pad}_{10^*}^{-1}(\bar{A}_1, \dots, \bar{A}_\alpha)$	
11: $\bar{M} \leftarrow \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, C)$	
12: <b>return</b> $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$	

Fig. 19: ELEPHANT adversary  $\mathcal{A}$  from [Theorem 3](#).

ELEPHANT.ENC( $K, N, A, M$ )	ENC $_{\mathcal{T}}(K, N, A, C)$
1: $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, M)$	9: $A_1, \dots, A_\alpha \xleftarrow{n} \text{pad}_{10^*}(N \parallel A, n)$
2: $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$	10: $C_1, \dots, C_\gamma \xleftarrow{n} \text{pad}_{10^*}(C, n)$
3: <b>return</b> $(C, T)$	11: $T \leftarrow A_1$
	12: <b>for</b> $i = 2, \dots, \alpha$
ENC $_{\mathcal{M}}(K, N, M)$	13: $T \leftarrow T \oplus \text{TBC}^{(i-1, 0)}(K, A_i)$
4: $M_1, \dots, M_\mu \xleftarrow{n} \text{pad}_{0^*}(M, n)$	14: <b>for</b> $i = 1, \dots, \gamma$
5: <b>for</b> $i = 1, \dots, \mu$	15: $T \leftarrow T \oplus \text{TBC}^{(i-1, 2)}(K, C_i)$
6: $C_i \leftarrow M_i \oplus \text{TBC}^{(i-1, 1)}(K, N)$	16: $T \leftarrow \text{TBC}^{(0, 0)}(K, T)$
7: $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$	17: <b>return</b> $[T]_\tau$
8: <b>return</b> $C$	

Fig. 20: Pseudocode of ELEPHANT [16] in terms of ENC $_{\mathcal{M}}$  and ENC $_{\mathcal{T}}$ .

## B.2 Proof of Theorem 4 (ROMULUS)

For the proof of Theorem 4, we formulate and prove two lemmas. Firstly, we show that the state-update-function  $\xi$  is invertible (Lemma 20). Secondly, we prove that we can invert both  $\text{ENC}_e$  and  $\text{ENC}_M$  (Lemma 21 and Lemma 22), where, for the latter, we make use of the invertibility of  $\xi$ .

Recall that the state-update-function  $\xi$  of ROMULUS maps a state  $S$  and an input  $I$  to a new state  $Y$  and an output  $O$ . In ROMULUS.DEC, the inverse of  $\xi$  is considered, however, inverse is understood only in relation to the output data. This means that the inverse function will *not* invert the state. When looking at ROMULUS, one can see that the output of  $\xi$  is discarded in  $\text{ENC}_e$ —a fact that will be exploited later. For our attack against  $\text{ENC}_M$ , this no longer works, as we have to invert the output of  $\xi$  while maintaining equal ciphertexts. The following lemma shows that we can invert  $\xi$  with respect to *both* its output and state. We write  $M$  for the input and  $C$  for the output of  $\xi$  (instead of  $I$  and  $O$ ), which is the case for our scenario.

**Lemma 20.** *Let  $\xi$  be the state-update-function of ROMULUS. Let further  $\mathcal{A}_\xi$  be the algorithm displayed in Fig. 21. Then for any  $(Y, C) \in \{0, 1\}^n \times \{0, 1\}^n$ , it holds that*

$$\xi(\mathcal{A}_\xi(Y, C)) = (Y, C).$$

*Proof.* We first observe that the matrix  $G$  works block-wise on blocks of 8 bits. Since  $\mathcal{A}_\xi$  splits the state and ciphertext into blocks of 8 bits each and runs  $\mathcal{A}_G$  (cf. Fig. 21) on each of these blocks, we merely need to show that  $\mathcal{A}_G$  correctly inverts an 8-bit block.

Let  $(Y, C) \in \{0, 1\}^8 \times \{0, 1\}^8$  and  $(S, M) \leftarrow \mathcal{A}_G(Y, C)$ . For  $i \in \{1, \dots, 8\}$ , denote the  $i$ -th bit of  $Y$ ,  $C$ ,  $S$ , and  $M$  by  $Y[i]$ ,  $C[i]$ ,  $S[i]$ , and  $M[i]$ , respectively. Algorithm  $\mathcal{A}_G$  first computes  $S[1] \leftarrow Y[8] \oplus C[8]$  and  $M[1] \leftarrow Y[1] \oplus S[1]$ . Subsequently, for  $i \in \{2, \dots, 8\}$ , it computes  $S[i] \leftarrow M[i-1] \oplus C[i-1]$  and  $M[i] \leftarrow Y[i] \oplus S[i]$ . Denote the output of  $\xi$  on input  $(S, M)$  by  $(\bar{Y}, \bar{C})$ . By construction we have:

$$\begin{aligned} \bar{Y} &= M \oplus S \\ &= (M[1] \oplus S[1]) \parallel \dots \parallel (M[8] \oplus S[8]) \\ &= (Y[1] \oplus S[1] \oplus S[1]) \parallel \dots \parallel (Y[8] \oplus S[8] \oplus S[8]) \\ &= Y[1] \parallel \dots \parallel Y[8] \\ &= Y \end{aligned}$$

and

$$\begin{aligned}
\bar{C} &= M \oplus G_s(S) \\
&= (M[1] \oplus S[2]) \parallel \dots \parallel (M[7] \oplus S[8]) \parallel (M[8] \oplus S[8] \oplus S[1]) \\
&= (M[1] \oplus M[1] \oplus C[1]) \parallel \dots \parallel (M[7] \oplus M[7] \oplus C[7]) \parallel (M[8] \oplus S[8] \oplus S[1]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel (M[8] \oplus S[8] \oplus S[1]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel (M[8] \oplus M[7] \oplus C[7] \oplus S[1]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel (M[8] \oplus M[7] \oplus C[7] \oplus Y[8] \oplus C[8]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel (Y[8] \oplus S[8] \oplus M[7] \oplus C[7] \oplus Y[8] \oplus C[8]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel (Y[8] \oplus M[7] \oplus C[7] \oplus M[7] \oplus C[7] \oplus Y[8] \oplus C[8]) \\
&= C[1] \parallel \dots \parallel C[7] \parallel C[8] \\
&= C.
\end{aligned}$$

Thus we obtain  $\xi(\mathcal{A}_\xi(Y, C)) = \xi(S, M) = (\bar{Y}, \bar{C}) = (Y, C)$ , which concludes the proof.  $\square$

Next, we give an adversary that inverts  $\text{ENC}_e$ , i.e., for a given output  $S$  of  $\text{ENC}_e$  and a partial context  $(K, N)$ , it finds matching associated data. We exploit the fact that the associated blocks are XORed to the full state in  $\text{ENC}_e$ .

**Lemma 21.** *Consider ROMULUS described in Fig. 22. There exists an adversary  $\mathcal{A}_e$ , making  $q$  queries to  $\tilde{\text{E}}$  such that for any  $(K, N, S) \in \mathcal{K} \times \mathcal{N} \times \{0, 1\}^n$ , it holds that*

$$\Pr[\text{ENC}_e(K, N, A) = S \mid A \leftarrow \mathcal{A}_e(K, N, S)] = 1.$$

*The number of ideal tweakable cipher queries by  $\mathcal{A}_e$  is  $q = \lfloor \frac{\alpha}{2} \rfloor + 1$  for  $\alpha$  being the number of associated data blocks that  $\mathcal{A}_e$  outputs.*

*Proof.* We construct  $\mathcal{A}_e$  as shown in Fig. 21. As input it receives  $(K, N, S)$ . It chooses an arbitrary even number of associated data block  $\alpha$  and chooses all except the first block at random, i.e.,  $A_2, \dots, A_\alpha \leftarrow \ast \{0, 1\}^n$ .<sup>28</sup> In addition,  $\mathcal{A}$  sets  $A_{\alpha+1} \leftarrow 0^n$ .<sup>29</sup> Then  $\mathcal{A}$  proceeds by inverting  $S$  using the ideal tweakable cipher to obtain  $Y$ . For  $i \in \{1, \dots, \frac{\alpha}{2}\}$ ,  $\mathcal{A}$  first computes  $S \leftarrow Y \oplus A_{2i+1}$  followed by computing  $Y \leftarrow \tilde{\text{E}}^{-1}(K, A_{2i}, S)$ , inverting the state-update-function  $\xi$  and the ideal tweakable cipher. Denote the resulting state by  $Y_*$  (see also the state marked in red in Fig. 6). Adversary  $\mathcal{A}$  sets the first associated data as  $A_1 \leftarrow Y_*$  which ensures that  $\text{ENC}_e(K, N, A) = S$ .

The number of queries to the ideal tweakable cipher  $\tilde{\text{E}}$  by  $\mathcal{A}$  is  $\lfloor \frac{\alpha}{2} \rfloor + 1$ : the initial one plus  $\lfloor \frac{\alpha}{2} \rfloor$  for the for loop.  $\square$

We now give an adversary that inverts  $\text{ENC}_M$ , i.e., for a given ciphertext  $(C, T)$  and a partial input  $(K, N)$ , it finds a matching pair of state  $S$  and message  $M$ . The attack relies heavily on the invertibility of  $\xi$  as shown in Lemma 20.

<sup>28</sup> We assume that these blocks are chosen to exhibit a valid padding.

<sup>29</sup> This corresponds to the input of the last application of  $\xi$  in the upper part of Fig. 6.

**Lemma 22.** Consider ROMULUS as described in Fig. 22. There exists an adversary  $\mathcal{A}_{\mathcal{M}}$ , making  $q$  queries to  $\tilde{\mathbf{E}}$  such that for any  $(K, N, (C, T)) \in \mathcal{K} \times \mathcal{N} \times \mathcal{C}$ , it holds that

$$\Pr[\text{ENC}_{\mathcal{M}}(K, N, S, M) = (C, T) \mid (S, M) \leftarrow \mathcal{A}_{\mathcal{M}}(K, N, (C, T))] = 1.$$

The number of ideal tweakable cipher queries by  $\mathcal{A}_{\mathcal{M}}$  is  $q = \mu$  for  $\mu$  being the number of ciphertext blocks that  $\mathcal{A}_{\mathcal{M}}$  receives as input.

*Proof.* We construct adversary  $\mathcal{A}_{\mathcal{M}}$  as shown in Fig. 21.  $\mathcal{A}_{\mathcal{M}}$  gets as input  $(K, N, (C, T))$ . For ease of exposition, we assume that the length of  $C$  is a multiple of the block size  $n$ .<sup>30</sup> Let  $C_1, \dots, C_\gamma \stackrel{n}{\leftarrow} C$ . The adversary  $\mathcal{A}_{\mathcal{M}}$  first sets  $Y \leftarrow T$  and then computes  $S \leftarrow G^{-1}(Y)$ . For  $i \in \{1, \dots, \gamma\}$ ,  $\mathcal{A}_{\mathcal{M}}$  computes  $Y \leftarrow \tilde{\mathbf{E}}^{-1}(K, N, S)$ <sup>31</sup> followed by the computation of  $(S, M_i) \leftarrow \mathcal{A}_\xi(Y, C_i)$  from Lemma 20. Finally,  $\mathcal{A}_{\mathcal{M}}$  outputs  $(S, M)$ , where  $M = M_1 \parallel \dots \parallel M_\gamma$ . By construction, it holds that  $\text{ENC}_{\mathcal{M}}(K, N, S, M) = (C, T)$  as  $\mathcal{A}_{\mathcal{M}}$  inverted all invocations of  $\tilde{\mathbf{E}}$  and  $\xi$  during  $\text{ENC}_{\mathcal{M}}$ —using  $\mathcal{A}_\xi$  from Lemma 20 for the latter.

$\mathcal{A}$  queries the ideal tweakable cipher  $\tilde{\mathbf{E}}$  a total of  $\mu$  times.  $\square$

Having established Lemma 20, Lemma 21, and Lemma 22, we can now prove Theorem 4.

*Proof (of Theorem 4).* We construct the following adversary  $\mathcal{A}$  against ROMULUS as shown in Fig. 21. It samples a context  $(K, N, A)$  together with a message  $M$  at random and computes the ciphertext  $(C, T) \leftarrow \text{ROMULUS.ENC}(K, N, A, M)$ . It then samples  $(\bar{K}, \bar{N})$  at random, computes  $(\bar{S}, \bar{M}) \leftarrow \mathcal{A}_{\mathcal{M}}(\bar{K}, \bar{N}, (C, T))$ , and  $\bar{A} \leftarrow \mathcal{A}_e(\bar{K}, \bar{N}, \bar{S})$ . Finally,  $\mathcal{A}$  outputs  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . By using Lemma 21 and Lemma 22, we obtain

$$\begin{aligned} \text{ROMULUS.ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M}) &= \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, \text{ENC}_e(\bar{K}, \bar{N}, \bar{A}), \bar{M}) \\ &= \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, \bar{S}, \bar{M}) && \text{(Lemma 21)} \\ &= (C, T) && \text{(Lemma 22)} \\ &= \text{ROMULUS.ENC}(K, N, A, M). \end{aligned}$$

As for the number of queries to the ideal tweakable cipher  $\tilde{\mathbf{E}}$ ,  $\mathcal{A}$  makes  $\mu + \lfloor \frac{\alpha}{2} \rfloor + 1$  while computing the ciphertext  $(C, T)$  for the first tuple and additionally  $\mu$  and  $\lfloor \frac{\alpha}{2} \rfloor + 1$  queries while running  $\mathcal{A}_{\mathcal{M}}$  and  $\mathcal{A}_e$ , respectively. This accumulates to  $q = 2\mu + \lfloor \frac{\alpha}{2} \rfloor + \lfloor \frac{\alpha}{2} \rfloor + 2$  queries in total and concludes the proof.  $\square$

The attack easily extends to  $\text{CMT}_{\mathcal{K}}$ ,  $\text{CMT}_{\mathcal{N}}$ ,  $\text{CMT}_{\mathcal{A}}$ , and  $\text{CMT}_{\mathcal{A}}^*$  attacks. The reasoning follows the one given for ELEPHANT.

<sup>30</sup> This is justified by letting  $\mathcal{A}$  choose a message satisfying this.

<sup>31</sup> Note that we drop the counter which is part of the tweak for simplicity.

<hr/> ROMULUS adversary $\mathcal{A}()$ <hr/> 1 : $(K, N, A, M) \leftarrow_{\$} \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ 2 : $(C, T) \leftarrow \text{ENC}(K, N, A, M)$ 3 : $(\overline{K}, \overline{N}) \leftarrow_{\$} \mathcal{K} \times \mathcal{N}$ 4 : $(\overline{S}, \overline{M}) \leftarrow \mathcal{A}_{\mathcal{M}}(\overline{K}, \overline{N}, (C, T))$ 5 : $\overline{A} \leftarrow \mathcal{A}_e(\overline{K}, \overline{N}, \overline{S})$ 6 : <b>return</b> $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ <hr/> $\mathcal{A}_{\xi}(Y, C)$ <hr/> 7 : $Y_1, \dots, Y_{\frac{n}{8}} \xleftarrow{\$} S$ 8 : $C_1, \dots, C_{\frac{n}{8}} \xleftarrow{\$} C$ 9 : <b>for</b> $i = 1, \dots, \frac{n}{8}$ <b>do</b> 10 : $(S_i, M_i) \leftarrow \mathcal{A}_G(Y_i, C_i)$ 11 : $S \leftarrow S_1 \parallel \dots \parallel S_{\frac{n}{8}}$ 12 : $M \leftarrow M_1 \parallel \dots \parallel M_{\frac{n}{8}}$ 13 : <b>return</b> $(S, M)$ <hr/> $\mathcal{A}_G(Y, C)$ <hr/> 14 : $Y[1], \dots, Y[8] \xleftarrow{\$} Y$ 15 : $C[1], \dots, C[8] \xleftarrow{\$} C$ 16 : $S[1] \leftarrow Y[8] \oplus C[8]$ 17 : $M[1] \leftarrow Y[1] \oplus S[1]$ 18 : <b>for</b> $i = 2, \dots, 8$ <b>do</b> 19 : $S[i] \leftarrow M[i-1] \oplus C[i-1]$ 20 : $M[i] \leftarrow Y[i] \oplus S[i]$ 21 : $S \leftarrow S_1 \parallel \dots \parallel S_8$ 22 : $M \leftarrow M_1 \parallel \dots \parallel M_8$ 23 : <b>return</b> $(S, M)$	<hr/> ENC <sub>e</sub> adversary $\mathcal{A}_e(K, N, S)$ <hr/> 24 : $\alpha \leftarrow_{\$} 2\mathbb{N}$ 25 : $A_2, \dots, A_{\alpha} \leftarrow_{\$} \{0, 1\}^n$ 26 : $A_{\alpha+1} \leftarrow 0^n$ 27 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, N, S)$ 28 : <b>for</b> $i = \frac{\alpha}{2}, \dots, 1$ <b>do</b> 29 : $S \leftarrow Y \oplus A_{2i+1}$ 30 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, A_{2i}, S)$ 31 : $A_1 \leftarrow Y$ 32 : $A \leftarrow A_1 \parallel \dots \parallel A_{\alpha}$ 33 : <b>return</b> $A$ <hr/> ENC <sub>M</sub> adversary $\mathcal{A}_{\mathcal{M}}(K, N, (C, T))$ <hr/> 34 : $C_1, \dots, C_{\gamma} \xleftarrow{\$} C$ 35 : $Y \leftarrow T$ 36 : $S \leftarrow G^{-1}(Y)$ 37 : <b>for</b> $i = \mu, \dots, 1$ <b>do</b> 38 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, N, S)$ 39 : $(S, M_i) \leftarrow \mathcal{A}_{\xi}(Y, C_i)$ 40 : $M \leftarrow M_1 \parallel \dots \parallel M_{\mu}$ 41 : <b>return</b> $(S, M)$
---	--

Fig. 21: ROMULUS adversary  $\mathcal{A}$  from Theorem 4 and the state-update-function adversary  $\mathcal{A}_{\xi}$  from Lemma 20.

ROMULUS.ENC( $K, N, A, M$ )	ENC $_M(K, N, S, M)$
1: $S \leftarrow \text{ENC}_e(K, N, A)$	15: $M_1, \dots, M_\mu \xleftarrow{n} \text{pad}_L(M, n)$
2: $(C, T) \leftarrow \text{ENC}_M(K, N, S, M)$	16: $S \leftarrow S$
3: <b>return</b> $(C, T)$	17: <b>for</b> $i = 1, \dots, \mu - 1$
<hr style="width: 100%; border: 0.5px solid black;"/>	18: $(Y, C_i) \leftarrow \xi(S, M_i)$
ENC $_e(K, N, A)$	19: $S \leftarrow \text{TBC}^N(K, Y)$
4: $A_1, \dots, A_\alpha \xleftarrow{n} \text{pad}_L(A, n)$	20: $(Y, C_\mu) \leftarrow \xi(S, M_\mu)$
5: $S \leftarrow 0^n$	21: $S \leftarrow \text{TBC}^N(K, Y)$
6: <b>for</b> $i = 1, \dots, \lfloor \frac{\alpha}{2} \rfloor$	22: $(\cdot, O) \leftarrow \xi(S, 0^n)$
7: $(Y, \cdot) \leftarrow \xi(S, A_{2i-1})$	23: $T \leftarrow [O]_\tau$
8: $S \leftarrow \text{TBC}^{A_{2i}}(K, Y)$	24: $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$
9: $V \leftarrow 0^n$	25: <b>return</b> $(C, T)$
10: <b>if</b> $\alpha \bmod 2 \neq 0$	<hr style="width: 100%; border: 0.5px solid black;"/>
11: $V \leftarrow A_\alpha$	$\xi(S, I)$
12: $(Y, \cdot) \leftarrow \xi(S, V)$	<hr style="width: 100%; border: 0.5px solid black;"/>
13: $S \leftarrow \text{TBC}^N(K, Y)$	26: $Y \leftarrow S \oplus I$
14: <b>return</b> $S$	27: $O \leftarrow G(S) \oplus I$
	28: <b>return</b> $(Y, O)$

Fig. 22: Pseudocode of ROMULUS [40] in terms of ENC $_e$  and ENC $_M$ . For sake of simplicity, we drop the counter that is part of the tweak.

### B.3 Proof of Theorem 5 (GIFT-COFB)

For the proof of Theorem 5 we drop the XOR of the masking values. This avoids some very cumbersome and tedious bookkeeping. Subsequent to the proof, we discuss why the results also hold if the masking values are used.

Before giving the proof of Theorem 5, we give three lemmas that we will use to prove it. The first lemma, Lemma 23, shows that there is an algorithm that inverts the state-update-function of GIFT-COFB for random outputs with probability  $\frac{1}{2}$ . The second lemma, Lemma 24, shows that  $\text{ENC}_C$  can be inverted, i.e., given an arbitrary key  $K$ , a nonce  $N$ , and an output state  $S$ , there is an algorithm that outputs associated data  $A$ , such that  $\text{ENC}_C(K, N, A) = S$ . The third lemma, Lemma 25, shows that there is an algorithm that inverts  $\text{ENC}_M$ . The latter relies heavily on Lemma 23 which is why the success probability drops exponentially in the number of ciphertext blocks.

**Lemma 23.** *Let  $\xi$  be the state-update-function of GIFT-COFB. Let further  $\mathcal{A}_\xi$  be the algorithm displayed in Fig. 23. Let  $C$  be an arbitrary bit string of length  $n$ . Then for  $Y \leftarrow_s \{0, 1\}^n$ , it holds that*

$$\Pr[\xi(\mathcal{A}_\xi(Y, C)) = (Y, C)] = \frac{1}{2}.$$

*Proof.* We start by describing the algorithm  $\mathcal{A}_\xi$  that is displayed in Fig. 23. We divide  $C$  and  $Y$  in two  $\frac{n}{2}$ -sized blocks, denoted by  $C_1, C_2$  and  $Y_1, Y_2$ , respectively. Further, we denote the bits of  $C_1 \oplus Y_1 \oplus C_2 \oplus Y_2$  by  $z_1, \dots, z_{\frac{n}{2}}$ . Then, we randomly sample a bit  $s_1$  and set  $s_i \leftarrow s_{i-1} \oplus z_{i-1}$  for all  $i = 2, \dots, \frac{n}{2}$ . By  $S_1$  we denote the bit string resulting from concatenating the  $s_i$ 's. Next, we define  $S_2 \leftarrow C_1 \oplus Y_1 \oplus S_1$  and  $S \leftarrow S_1 \parallel S_2$ , which will be the first output of  $\mathcal{A}_\xi$ . Further, we set  $M_1 \leftarrow Y_1 \oplus S_2$  and  $M_2 \leftarrow Y_2 \oplus (S_1 \lll 1)$ . The concatenated bit string  $M \leftarrow M_1 \parallel M_2$  is the second output of  $\mathcal{A}_\xi$ , i.e.,  $\mathcal{A}_\xi(Y, C) = (S, M)$ .

Next, we check that  $\xi(\mathcal{A}_\xi(Y, C)) = (Y, C)$  holds with probability  $\frac{1}{2}$ . By definition of the state-update-function, the first component of  $\xi(\mathcal{A}_\xi(Y, C)) = \xi(S, M)$  is given by

$$\begin{aligned} \tilde{G}(S) \oplus M &= (S_2, S_1 \lll 1) \oplus M \\ &= (S_2 \oplus M_1, (S_1 \lll 1) \oplus M_2) \\ &= (Y_1, Y_2) = Y. \end{aligned}$$

The second component of  $\xi(\mathcal{A}_\xi(Y, C)) = \xi(S, M)$  computes as

$$\begin{aligned} M \oplus S &= (M_1 \oplus S_1, M_2 \oplus S_2) \\ &= (Y_1 \oplus S_2 \oplus S_1, Y_2 \oplus (S_1 \lll 1) \oplus S_2). \end{aligned}$$

Note that this expression is equal to  $C$  if and only if  $S_1 \oplus (S_1 \lll 1) = C_1 \oplus Y_1 \oplus C_2 \oplus Y_2$ , which results from solving  $C_1 = Y_1 \oplus S_2 \oplus S_1$  for  $S_2$  and plugging the result into  $C_2 = Y_2 \oplus (S_1 \lll 1) \oplus S_2$ . Breaking down this equation to the bit-level, yields

$$(s_1, s_2, \dots, s_{\frac{n}{2}}) \oplus (s_2, s_3, \dots, s_{\frac{n}{2}}, s_1) = (z_1, \dots, z_{\frac{n}{2}}).$$

This gives the following equations

$$\begin{aligned}
s_1 \oplus s_2 &= z_1 \\
s_2 \oplus s_3 &= z_2 \\
&\vdots \\
s_{\frac{n}{2}-1} \oplus s_{\frac{n}{2}} &= z_{\frac{n}{2}-1} \\
s_{\frac{n}{2}} \oplus s_1 &= z_{\frac{n}{2}}
\end{aligned}$$

of which the first  $\frac{n}{2} - 1$  equations hold by construction of the algorithm  $\mathcal{A}_\xi$ . Finally, we need to determine in which cases the last equation holds. For this, we replace  $s_{\frac{n}{2}}$  by  $s_{\frac{n}{2}-1} \oplus z_{\frac{n}{2}-1}$  which we get from the second to last equation. Then, in turn, we replace  $s_{\frac{n}{2}-1}$  with  $s_{\frac{n}{2}-2} \oplus z_{\frac{n}{2}-2}$  and continue this process until we get an equation depending only on the  $z_i$ 's and  $s_1$ , namely

$$z_{\frac{n}{2}-1} \oplus \dots \oplus z_2 \oplus z_1 \oplus s_1 \oplus s_1 = z_{\frac{n}{2}},$$

which is equivalent to  $z_{\frac{n}{2}} \oplus z_{\frac{n}{2}-1} \oplus \dots \oplus z_2 \oplus z_1 = 0$ . Thus we get  $\xi(\mathcal{A}_\xi(Y, C)) = (\cdot, C)$  if and only if  $C_1 \oplus Y_1 \oplus C_2 \oplus Y_2 = z_1 \parallel \dots \parallel z_{\frac{n}{2}}$  contains an even number of 1s. Since  $Y$  is chosen uniformly at random, this is the case with probability  $\frac{1}{2}$ . This finishes the proof of the claim.  $\square$

**Lemma 24.** *Consider GIFT-COFB as described in Fig. 24. There exists an adversary  $\mathcal{A}_e$ , making  $q$  queries to the ideal cipher  $\mathbf{E}$  such that for any  $(K, N, S) \in \mathcal{K} \times \mathcal{N} \times \{0, 1\}^n$ , it holds that*

$$\Pr[\text{ENC}_e(K, N, A) = S \mid A \leftarrow \mathcal{A}_e(K, N, S)] = 1.$$

*The number of ideal cipher queries by  $\mathcal{A}_e$  is  $q = \alpha + 1$  for  $\alpha$  being the number of associated data blocks that  $\mathcal{A}_e$  outputs.*

*Proof.* We construct the following adversary  $\mathcal{A}_e$  against GIFT-COFB as shown in Fig. 23. Its input is  $(K, N, S)$ . First,  $\mathcal{A}_e$  randomly picks associated data blocks  $A_2, \dots, A_\alpha$ , i.e., all except the first one.<sup>32</sup> Next,  $\mathcal{A}_e$  computes both  $S_*$  and  $Y_*$  (cf. Fig. 7): For the former,  $\mathcal{A}_e$  computes  $S_* \leftarrow \mathbf{E}(K, N)$  and for the latter,  $\mathcal{A}_e$  consecutively inverts the ideal cipher  $\mathbf{E}$  (starting from the input  $S$ ) and the state-update-function  $\xi$  (by XORing an associated data block and inverting  $\tilde{G}$ —note that we merely need to invert  $\tilde{G}$  as the output of  $\xi$  is discarded). Finally,  $\mathcal{A}_e$  computes  $A_1 \leftarrow S_* \oplus Y_*$  and outputs  $A$ . By construction, it holds that

$$\text{ENC}_e(K, N, A) = S.$$

The number of ideal cipher queries by  $\mathcal{A}_e$  is  $\alpha + 1$ .  $\square$

<sup>32</sup> We assume that  $\mathcal{A}_e$  picks the blocks such that they exhibit a valid padding according to GIFT-COFB.

**Lemma 25.** Consider GIFT-COFB described in Fig. 24. There exists an adversary  $\mathcal{A}_{\mathcal{M}}$ , making  $q$  queries to the ideal cipher  $\mathbf{E}$  such that

$$\Pr[\text{ENC}_{\mathcal{M}}(K, N, S, M) = (C, T) \mid (S, M) \leftarrow \mathcal{A}_{\mathcal{M}}(K, N, (C, T))] = \frac{1}{2^\mu},$$

holds for any  $(K, N, (C, T)) \in \mathcal{K} \times \mathcal{N} \times \mathcal{C}$ . The number of ideal cipher queries by  $\mathcal{A}_{\mathcal{M}}$  is  $q = \mu$  for  $\mu$  being the number of ciphertext blocks that  $\mathcal{A}_{\mathcal{M}}$  receives as input.

*Proof.* We construct an adversary  $\mathcal{A}_{\mathcal{M}}$  against GIFT-COFB as shown in Fig. 23. It gets  $(K, N, (C, T))$  as input. For ease of exposition, we assume that the length of  $C$  is a multiple of the block size  $n$ , hence  $C_1, \dots, C_\mu \stackrel{n}{\leftarrow} C$  yields  $\mu$  full blocks of length  $n$ . First, Adversary  $\mathcal{A}_{\mathcal{M}}$  sets  $S \leftarrow T$ . Next,  $\mathcal{A}_{\mathcal{M}}$  consecutively inverts the ideal cipher  $Y \leftarrow \mathbf{E}^{-1}(K, S)$  and computes  $(S, M_i) \leftarrow \mathcal{A}_\xi(Y, C_i)$ . Finally, it sets  $M \leftarrow M_1 \parallel \dots \parallel M_\mu$  and outputs  $(S, M)$ . Provided that  $\mathcal{A}_\xi$  correctly inverts  $\xi$ , we obtain

$$\text{ENC}_{\mathcal{M}}(K, N, S, M) = (C, T).$$

By Lemma 23—using  $Y$  is uniformly random due to  $\mathbf{E}$  being an ideal cipher—every inversion of  $\xi$  succeeds with probability  $\frac{1}{2}$ . Since  $\mathcal{A}_\xi$  is run  $\mu$  times by  $\mathcal{A}_{\mathcal{M}}$ , we get

$$\Pr[\text{ENC}_{\mathcal{M}}(K, N, S, M) = (C, T) \mid (S, M) \leftarrow \mathcal{A}_{\mathcal{M}}(K, N, (C, T))] = \frac{1}{2^\mu}.$$

The number of queries to the ideal cipher  $\mathbf{E}$  by  $\mathcal{A}_{\mathcal{M}}$  is  $\mu$ . □

We are now ready to give the proof of Theorem 5.

*Proof (of Theorem 5).* We construct  $\mathcal{A}$  against GIFT-COFB as displayed in Fig. 23. It starts by sampling a context  $(K, N, A)$  together with a message  $M$  at random and computes  $(C, T) \leftarrow \text{GIFT-COFB.ENC}(K, N, A, M)$ . Next, it samples a fresh key-nonce pair  $(\bar{K}, \bar{N})$ , computes  $(\bar{S}, \bar{M}) \leftarrow \mathcal{A}_{\mathcal{M}}(\bar{K}, \bar{N}, (C, T))$ , and  $\bar{A} \leftarrow \mathcal{A}_e(\bar{K}, \bar{N}, \bar{S})$ . Finally,  $\mathcal{A}$  outputs  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ .

By Lemma 24 and Lemma 25, we have the following equalities with probability  $\frac{1}{2^\mu}$ :

$$\begin{aligned} \text{GIFT-COFB.ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M}) &= \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, \text{ENC}_e(\bar{K}, \bar{N}, \bar{A}), \bar{M}) \\ &= \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, \bar{S}, \bar{M}) && \text{(Lemma 24)} \\ &= (C, T) && \text{(Lemma 25)} \\ &= \text{GIFT-COFB.ENC}(K, N, A, M). \end{aligned}$$

This yields

$$\mathbf{Adv}_{\text{GIFT-COFB}}^{\text{CMT}}(\mathcal{A}) = \frac{1}{2^\mu}.$$

Adversary  $\mathcal{A}$  makes a total of  $q = 2\mu + \alpha + \bar{\alpha} + 2$  queries to  $\mathbf{E}$ :  $\mu + \alpha + 1$  for computing the first tuple,  $\mu$  for inverting  $\text{ENC}_{\mathcal{M}}$ , and  $\bar{\alpha} + 1$  for inverting  $\text{ENC}_e$ . □

When considering the masking values, inverting  $\text{ENC}_{\mathcal{M}}$  might seem to be a problem, as the adversary needs to invert the function using the correct masking values. We observe that the masking values depend on the key, the nonce, and the length of associated data and message/ciphertext. In particular, they are independent of the exact values of  $A$ ,  $M$ , and  $C$ . Thus, when inverting  $\text{ENC}_{\mathcal{M}}$ , the adversary merely has to choose how long the associated data will be, as this allows to use the correct masking values.

The attack easily extends to  $\text{CMT}_{\mathcal{K}}$ ,  $\text{CMT}_{\mathcal{N}}$ ,  $\text{CMT}_{\mathcal{A}}$ , and  $\text{CMT}_{\mathcal{A}}^*$  attacks, following the argument we gave for **ELEPHANT**.

GIFT-COFB adversary $\mathcal{A}()$	$\text{ENC}_{\mathcal{M}}$ adversary $\mathcal{A}_{\mathcal{M}}(K, N, (C, T))$																												
1: $(K, N, A, M) \leftarrow_{\$} \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$	17: $S \leftarrow T$																												
2: $(C, T) \leftarrow \text{ENC}(K, N, A, M)$	18: <b>for</b> $i = \mu, \dots, 1$																												
3: $(\overline{K}, \overline{N}) \leftarrow_{\$} \mathcal{K} \times \mathcal{N}$	19: $Y \leftarrow \mathbf{E}^{-1}(K, S)$																												
4: $(\overline{S}, \overline{M}) \leftarrow \mathcal{A}_{\mathcal{M}}(\overline{K}, \overline{N}, (C, T))$	20: $(S, M_i) \leftarrow \mathcal{A}_{\xi}(Y, C_i)$																												
5: $\overline{A} \leftarrow \mathcal{A}_{\mathcal{C}}(\overline{K}, \overline{N}, \overline{S})$	21: $M \leftarrow M_1 \parallel \dots \parallel M_{\mu}$																												
6: <b>return</b> $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$	22: <b>return</b> $(S, M)$																												
<table border="0" style="width: 100%;"> <thead> <tr> <th style="border-bottom: 1px solid black;"><math>\text{ENC}_{\mathcal{C}}</math> adversary <math>\mathcal{A}_{\mathcal{C}}(K, N, S)</math></th> <th style="border-bottom: 1px solid black;"><math>\mathcal{A}_{\xi}(Y, C)</math></th> </tr> </thead> <tbody> <tr> <td>7: <math>A_2, \dots, A_{\alpha} \leftarrow_{\\$} \{0, 1\}^n</math></td> <td>23: <math>C_1, C_2 \leftarrow_{\frac{n}{2}} C</math></td> </tr> <tr> <td>8: <math>S_* \leftarrow \mathbf{E}(K, N)</math></td> <td>24: <math>Y_1, Y_2 \leftarrow_{\frac{n}{2}} Y</math></td> </tr> <tr> <td>9: <math>Y \leftarrow \mathbf{E}^{-1}(K, S)</math></td> <td>25: <math>z_1, \dots, z_{\frac{n}{2}} \leftarrow^1 C_1 \oplus Y_1 \oplus C_2 \oplus Y_2</math></td> </tr> <tr> <td>10: <b>for</b> <math>i = \alpha, \dots, 2</math></td> <td>26: <math>s_1 \leftarrow_{\\$} \{0, 1\}</math></td> </tr> <tr> <td>11: <math>S \leftarrow \tilde{\mathbf{G}}^{-1}(Y \oplus A_i)</math></td> <td>27: <b>for</b> <math>i = 2, \dots, \frac{n}{2}</math> <b>do</b></td> </tr> <tr> <td>12: <math>Y \leftarrow \mathbf{E}^{-1}(K, S)</math></td> <td>28: <math>s_i \leftarrow s_{i-1} \oplus z_{i-1}</math></td> </tr> <tr> <td>13: <math>Y_* \leftarrow Y</math></td> <td>29: <math>S_1 \leftarrow s_1 \parallel \dots \parallel s_{\frac{n}{2}}</math></td> </tr> <tr> <td>14: <math>A_1 \leftarrow \tilde{\mathbf{G}}(S_*) \oplus Y_*</math></td> <td>30: <math>S_2 \leftarrow C_1 \oplus Y_1 \oplus S_1</math></td> </tr> <tr> <td>15: <math>A \leftarrow A_1 \parallel \dots \parallel A_{\alpha}</math></td> <td>31: <math>S \leftarrow S_1 \parallel S_2</math></td> </tr> <tr> <td>16: <b>return</b> <math>A</math></td> <td>32: <math>M_1 \leftarrow Y_1 \oplus S_2</math></td> </tr> <tr> <td></td> <td>33: <math>M_2 \leftarrow Y_2 \oplus (S_1 \lll 1)</math></td> </tr> <tr> <td></td> <td>34: <math>M \leftarrow M_1 \parallel M_2</math></td> </tr> <tr> <td></td> <td>35: <b>return</b> <math>(S, M)</math></td> </tr> </tbody> </table>		$\text{ENC}_{\mathcal{C}}$ adversary $\mathcal{A}_{\mathcal{C}}(K, N, S)$	$\mathcal{A}_{\xi}(Y, C)$	7: $A_2, \dots, A_{\alpha} \leftarrow_{\$} \{0, 1\}^n$	23: $C_1, C_2 \leftarrow_{\frac{n}{2}} C$	8: $S_* \leftarrow \mathbf{E}(K, N)$	24: $Y_1, Y_2 \leftarrow_{\frac{n}{2}} Y$	9: $Y \leftarrow \mathbf{E}^{-1}(K, S)$	25: $z_1, \dots, z_{\frac{n}{2}} \leftarrow^1 C_1 \oplus Y_1 \oplus C_2 \oplus Y_2$	10: <b>for</b> $i = \alpha, \dots, 2$	26: $s_1 \leftarrow_{\$} \{0, 1\}$	11: $S \leftarrow \tilde{\mathbf{G}}^{-1}(Y \oplus A_i)$	27: <b>for</b> $i = 2, \dots, \frac{n}{2}$ <b>do</b>	12: $Y \leftarrow \mathbf{E}^{-1}(K, S)$	28: $s_i \leftarrow s_{i-1} \oplus z_{i-1}$	13: $Y_* \leftarrow Y$	29: $S_1 \leftarrow s_1 \parallel \dots \parallel s_{\frac{n}{2}}$	14: $A_1 \leftarrow \tilde{\mathbf{G}}(S_*) \oplus Y_*$	30: $S_2 \leftarrow C_1 \oplus Y_1 \oplus S_1$	15: $A \leftarrow A_1 \parallel \dots \parallel A_{\alpha}$	31: $S \leftarrow S_1 \parallel S_2$	16: <b>return</b> $A$	32: $M_1 \leftarrow Y_1 \oplus S_2$		33: $M_2 \leftarrow Y_2 \oplus (S_1 \lll 1)$		34: $M \leftarrow M_1 \parallel M_2$		35: <b>return</b> $(S, M)$
$\text{ENC}_{\mathcal{C}}$ adversary $\mathcal{A}_{\mathcal{C}}(K, N, S)$	$\mathcal{A}_{\xi}(Y, C)$																												
7: $A_2, \dots, A_{\alpha} \leftarrow_{\$} \{0, 1\}^n$	23: $C_1, C_2 \leftarrow_{\frac{n}{2}} C$																												
8: $S_* \leftarrow \mathbf{E}(K, N)$	24: $Y_1, Y_2 \leftarrow_{\frac{n}{2}} Y$																												
9: $Y \leftarrow \mathbf{E}^{-1}(K, S)$	25: $z_1, \dots, z_{\frac{n}{2}} \leftarrow^1 C_1 \oplus Y_1 \oplus C_2 \oplus Y_2$																												
10: <b>for</b> $i = \alpha, \dots, 2$	26: $s_1 \leftarrow_{\$} \{0, 1\}$																												
11: $S \leftarrow \tilde{\mathbf{G}}^{-1}(Y \oplus A_i)$	27: <b>for</b> $i = 2, \dots, \frac{n}{2}$ <b>do</b>																												
12: $Y \leftarrow \mathbf{E}^{-1}(K, S)$	28: $s_i \leftarrow s_{i-1} \oplus z_{i-1}$																												
13: $Y_* \leftarrow Y$	29: $S_1 \leftarrow s_1 \parallel \dots \parallel s_{\frac{n}{2}}$																												
14: $A_1 \leftarrow \tilde{\mathbf{G}}(S_*) \oplus Y_*$	30: $S_2 \leftarrow C_1 \oplus Y_1 \oplus S_1$																												
15: $A \leftarrow A_1 \parallel \dots \parallel A_{\alpha}$	31: $S \leftarrow S_1 \parallel S_2$																												
16: <b>return</b> $A$	32: $M_1 \leftarrow Y_1 \oplus S_2$																												
	33: $M_2 \leftarrow Y_2 \oplus (S_1 \lll 1)$																												
	34: $M \leftarrow M_1 \parallel M_2$																												
	35: <b>return</b> $(S, M)$																												

Fig. 23: GIFT-COFB adversary  $\mathcal{A}$  from [Theorem 5](#) and the state-update-function adversary  $\mathcal{A}_{\xi}$  from [Lemma 23](#).



#### B.4 Proof of Theorem 6 (PHOTON-BEETLE)

*Proof.* We construct the following CMT adversary  $\mathcal{A}$  against PHOTON-BEETLE as shown in Fig. 25. It chooses  $(K, N, A)$  uniformly at random from the respective sets and computes  $S \leftarrow \text{ENC}_e(K, N, A)$ . Let  $S_*$  denote the state before the domain separation is applied (see Fig. 8). Adversary  $\mathcal{A}$  then chooses different associated data  $\bar{A}$  and inverts  $\text{ENC}_e$  starting from  $S_*$  up to the initial state. This initial state is then used as the concatenation of nonce  $N$  and key  $K$ .<sup>33</sup> As the final step,  $\mathcal{A}$  picks a message  $M$  at random and outputs  $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, M))$ . By construction, we have  $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$  and it holds that

$$\begin{aligned} \text{PHOTON-BEETLE.ENC}(K, N, A, M) &= \text{ENC}_{\mathcal{M}}(\text{ENC}_e(K, N, A), M) \\ &= \text{ENC}_{\mathcal{M}}(S_*, M) \\ &= \text{ENC}_{\mathcal{M}}(\text{ENC}_e(\bar{K}, \bar{N}, \bar{A}), M) \\ &= \text{PHOTON-BEETLE.ENC}(\bar{K}, \bar{N}, \bar{A}, M), \end{aligned}$$

thus  $\mathcal{A}$  wins the game CMT.

$\mathcal{A}$  makes  $\alpha$  queries to compute  $S_*$  and additional  $\bar{\alpha}$  queries to compute  $(\bar{K}, \bar{N})$ , resulting in  $q = \alpha + \bar{\alpha}$  queries in total.  $\square$

The attack is by construction also a valid  $\text{CMT}_{\mathbf{A}}$  attack as the associated data are chosen to be different. While the adversary does not choose the key  $\bar{K}$  and nonce  $\bar{N}$  for the second tuple, it is clear that the attack easily extends to  $\text{CMT}_{\mathbf{K}}$  and  $\text{CMT}_{\mathbf{N}}$ —if  $(\bar{K}, \bar{N}) = (K, N)$ , the adversary simply chooses different associated data  $\bar{A}$  and repeats the attack until the keys and nonces differ.

PHOTON-BEETLE adversary $\mathcal{A}$	$\mathcal{B}(S_*, \bar{A})$
1 : $(K, N) \leftarrow_{\$} \mathcal{K} \times \mathcal{N}$	13 : $\bar{A}_1, \dots, \bar{A}_{\bar{\alpha}} \leftarrow^r \text{pad}_{10^*}(\bar{A}, r)$
2 : $S \leftarrow N \parallel K$	14 : <b>for</b> $i = \bar{\alpha}, \dots, 1$ <b>do</b>
3 : $A \leftarrow_{\$} \mathcal{A}$	15 : $S_* \leftarrow S_* \oplus (\bar{A}_i \parallel 0^c)$
4 : $A_1, \dots, A_{\alpha} \leftarrow^r \text{pad}_{10^*}(A, r)$	16 : $S_* \leftarrow \rho^{-1}(S_*)$
5 : <b>for</b> $i = 1, \dots, \alpha$ <b>do</b>	17 : $\bar{N} \parallel \bar{K} \leftarrow S_*$
6 : $S \leftarrow \rho(S)$	18 : <b>return</b> $(\bar{K}, \bar{N})$
7 : $S \leftarrow S \oplus (A_i \parallel 0^c)$	
8 : $S_* \leftarrow S$	
9 : $\bar{A} \leftarrow_{\$} \mathcal{A} \setminus \{A\}$	
10 : $(\bar{K}, \bar{N}) \leftarrow \mathcal{B}(S_*, \bar{A})$	
11 : $M \leftarrow_{\$} \mathcal{M}$	
12 : <b>return</b> $((K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, M))$	

Fig. 25: PHOTON-BEETLE adversary  $\mathcal{A}$  from Theorem 6.

<sup>33</sup> Note that these are likely to be different than  $K$  and  $N$  but not guaranteed to be.

PHOTON-BEETLE.ENC( $K, N, A, M$ )	ENC <sub>e</sub> ( $K, N, A$ )
1: $S \leftarrow \text{ENC}_e(K, N, A)$	13: $S \leftarrow N \parallel K$
2: $(C, T) \leftarrow \text{ENC}_M(S, M)$	14: $A_1, \dots, A_\alpha \xleftarrow{r} \text{pad}_{10^*}(A, r)$
3: <b>return</b> $(C, T)$	15: <b>for</b> $i = 1, \dots, \alpha$
ENC <sub>M</sub> ( $S, M$ )	16: $S \leftarrow \rho(S)$
4: $M_1, \dots, M_\mu \xleftarrow{r} \text{pad}_{10^*}(M, r)$	17: $S \leftarrow S \oplus (A_i \parallel 0^c)$
5: <b>for</b> $i = 1, \dots, \mu$	18: $S \leftarrow S \oplus (0^r \parallel \iota_0)$
6: $S \leftarrow \rho(S)$	19: <b>return</b> $S$
7: $(\lceil S \rceil_r, C_i) \leftarrow \xi(\lceil S \rceil_r, M_i)$	$\xi(S, I)$
8: $S \leftarrow S \oplus (0^r \parallel \iota_1)$	20: $O \leftarrow \text{Shuffle}(S) \oplus I$
9: $S \leftarrow \rho(S)$	21: $Y \leftarrow S \oplus I$
10: $T \leftarrow \lceil S \rceil_\tau$	22: <b>return</b> $(Y, O)$
11: $C \leftarrow \lceil C_1 \parallel \dots \parallel C_\mu \rceil_{ M }$	
12: <b>return</b> $(C, T)$	

Fig. 26: Pseudocode of PHOTON-BEETLE [5] in terms of ENC<sub>e</sub> and ENC<sub>M</sub>. Here,  $\text{Shuffle}(S) = S_2 \parallel (S_1 \ggg 1)$  for  $S_1, S_2 \xleftarrow{\frac{r}{2}} S$ .

### B.5 Proof of Theorem 7 (XOODYAK)

*Proof.* We construct a CMT adversary  $\mathcal{A}$  against XOODYAK. It uses a birthday attack to find a collision in the last 32 bits of the sponge state after the first application of  $\rho$ . For this,  $q = 2^{17} + 1$  different keys and  $q$  nonces are sampled randomly. For  $i \in [q]$ , we denote them by  $K_i$  and  $N_i$  and write  $\mathbf{K}_i \parallel \mathbf{N}_i = \text{pad}_e((K_i \parallel N_i \parallel \text{enc}_8(N_i)), 00000010)$  for their padded concatenation. Further, we consider the following random function

$$f : \{0, 1\}^{384} \rightarrow \{0, 1\}^{32}, \quad f(X) = \lfloor \rho(X) \rfloor_{32},$$

and compute  $f(\mathbf{K}_i \parallel \mathbf{N}_i)$  for each  $i \in [q]$ . By the birthday attack [17, Section 8.3]<sup>34</sup>, a collision of  $f$  is found with probability at least  $\frac{1}{2}$ . Assume that such a collision has been found and write  $(K, N)$  and  $(\bar{K}, \bar{N})$  for the key-nonce pairs that lead to it, i.e., have  $\lfloor \rho(\mathbf{K} \parallel \mathbf{N}) \rfloor_{32} = \lfloor \rho(\bar{\mathbf{K}} \parallel \bar{\mathbf{N}}) \rfloor_{32}$ , where  $\mathbf{K} \parallel \mathbf{N}$  and  $\bar{\mathbf{K}} \parallel \bar{\mathbf{N}}$  denote the corresponding padded values. Next, adversary  $\mathcal{A}$  picks  $A \in \{0, 1\}^{352}$  at random and computes  $\bar{A} \leftarrow \lceil \rho(\bar{\mathbf{K}} \parallel \bar{\mathbf{N}}) \rceil_{352} \oplus \lceil \rho(\mathbf{K} \parallel \mathbf{N}) \rceil_{352} \oplus A$ . Together with the collision on the last 32 bits,  $\mathcal{A}$  has produced a collision on the whole state. Then, the adversary wins the game CMT against XOODYAK by outputting  $(K, N, A, M)$  and  $(\bar{K}, \bar{N}, \bar{A}, M)$  for  $M$  some randomly sampled message. This is the case, as  $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$  and the states after the associated data is absorbed agree, after which point only the same input (namely  $M$ ) is processed for both tuples. Hence, we obtain  $\text{XOODYAK.ENC}(K, N, A, M) = \text{XOODYAK.ENC}(\bar{K}, \bar{N}, \bar{A}, M)$  and in total have shown that  $\mathcal{A}$  wins with probability at least  $\frac{1}{2}$  for  $q = 2^{17} + 1$  queries.  $\square$

Note that, using the same strategy as presented above, we obtain an attacker that wins with probability 1 after making  $2^{32} + 1$  queries. Further, the above attack is by construction also a valid  $\text{CMT}_K$  and  $\text{CMT}_N$  attack as keys and nonces, respectively, are chosen to be different. Moreover, it can be shown to be a  $\text{CMT}_A$  attack: The same associated data blocks are only chosen if the states after the first application of  $\rho$  already coincide in their rate part. Since they agree in the last 32 bits by construction, this would constitute a full-state collision of  $\rho$ , which is impossible for a permutation.

<sup>34</sup> Note that the prerequisite, regarding the size of domain and codomain of  $f$ , is fulfilled as  $2^{384} \geq 100 \cdot 2^{32}$ .

XOODYAK.ENC( $K, N, A, M$ )	ENC $_{\mathcal{M}}$ ( $S, M$ )
1 : $S \leftarrow \text{ENC}_e(K, N, A)$ 2 : $(C, T) \leftarrow \text{ENC}_{\mathcal{M}}(S, M)$ 3 : <b>return</b> $(C, T)$	15 : $Y \leftarrow S \oplus (0^{r+8} \parallel \iota_1)$ 16 : $M_1, \dots, M_\mu \xleftarrow{r} M$ 17 : <b>for</b> $i = 1, \dots, \mu$ 18 : $\mathbf{M}_i \leftarrow \text{pad}_{\mathcal{M}}(M_i)$ 19 : $S \leftarrow \rho(Y)$ 20 : $Y \leftarrow S \oplus (\mathbf{M}_i \parallel 0^{c-8})$ 21 : $C_i \leftarrow [Y]_r$ 22 : $Y \leftarrow Y \oplus (0^{r+8} \parallel \iota_2)$ 23 : $S \leftarrow \rho(Y)$ 24 : $T \leftarrow [S]_\tau$ 25 : $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$ 26 : <b>return</b> $(C, T)$
ENC $_e$ ( $K, N, A$ )	
4 : $S \leftarrow 0^n$ 5 : $X \leftarrow K \parallel N \parallel \text{enc}_8( N )$ 6 : $Y \leftarrow S \oplus \text{pad}_e(X, 0^6 \parallel 10)$ 7 : $S \leftarrow \rho(Y)$ 8 : $A_1, \dots, A_\alpha \xleftarrow{352} A$ 9 : $Y \leftarrow S \oplus \text{pad}_e(A_1, 0^6 \parallel 11)$ 10 : <b>for</b> $i = 2, \dots, \alpha$ 11 : $S \leftarrow \rho(Y)$ 12 : $Y \leftarrow S \oplus \text{pad}_e(A_i, 0^8)$ 13 : $S \leftarrow Y$ 14 : <b>return</b> $S$	

Fig. 27: Pseudocode of XOODYAK [23] in terms of ENC $_e$  and ENC $_{\mathcal{M}}$ .

## B.6 Proof of Theorem 8 (TINYJAMBU)

*Proof.* We construct a CMT adversary  $\mathcal{A}$  against TINYJAMBU as follows: First, we randomly choose two different keys  $K \neq \bar{K}$  and a target ciphertext  $C$ . Note that, due to the structure of TINYJAMBU, the context produces a key stream which is XORed with the message to obtain the ciphertext. Hence, for a random context it is always possible to find a  $M$  such that the TINYJAMBU encryption results in the initially chosen target ciphertext  $C$ . In the following, we will implicitly consider this “matching” message for each context that occurs. Hence it suffices to find two different contexts that—together with their matching message—yield colliding tags.

We start by building two lists of tags, where for one we use  $K$  as key and in the other  $\bar{K}$ . For this, sample pairwise different  $(N_i, A_i)$  for  $i \in \{1, \dots, 2^{32} + 1\}$  and compute the corresponding tags  $T_i$  using the key  $K$ . We then set  $L = (T_i)_{i \in [2^{32} + 1]}$ . Analogously, we sample pairwise different  $(\bar{N}_i, \bar{A}_i)$ <sup>35</sup> for  $i \in [2^{32} + 1]$  and write the corresponding tags  $\bar{T}_i$  (computed using  $\bar{K}$ ) into the list  $\bar{L} = (\bar{T}_i)_{i \in [2^{32} + 1]}$ . Building the two lists, takes a total of  $q = 2(2^{32} + 1)(6 + \alpha + \mu)$  queries to  $\rho$ .

For the context  $(K, N_i, A_i)$ , denote the states before the second to last permutation application (see Fig. 10) by  $S_i$ , and analogously for  $(\bar{K}, \bar{N}_i, \bar{A}_i)$  by  $\bar{S}_i$ . Note that for a fixed key, TINYJAMBU can be considered a sponge-based function with rate 32 and capacity 96. Therefore, the event  $S_i = S_j$  for  $i \neq j$  (and analogously  $\bar{S}_i \neq \bar{S}_j$  for  $i \neq j$ ), constitutes an inner collision, which is—for a sponge with capacity 96—highly unlikely<sup>36</sup>. As we model both  $\text{BC}_1$  and  $\text{BC}_2$  by an ideal cipher  $E$  and the states  $S_i$  (and respectively  $\bar{S}_i$ ) collide with negligible probability, we can assume the list elements to be distributed uniformly and independently.

This puts us in the situation of Lemma 18 (for  $l_1 = l_2 = 2^{32} + 1$  and  $\tau = 64$ ), hence we obtain the following lower bound for finding a collision  $T_i = \bar{T}_j$ :

$$\left(1 - \exp\left(\frac{-(2^{33} + 2)(2^{33} + 1)}{2^{65}}\right)\right) \cdot \frac{2 \cdot (2^{32} + 1)^2}{(2^{33} + 2)^2 - (2^{33} + 2)} \quad (1)$$

Since

$$\frac{-(2^{33} + 2)(2^{33} + 1)}{2^{65}} \leq \frac{-2^{33} \cdot 2^{33}}{2^{65}} = -2,$$

the first factor in Eq. (1) can be bounded below by  $1 - \exp\left(\frac{-(2^{33} + 2)(2^{33} + 1)}{2^{65}}\right) \geq 1 - e^{-2} \geq \frac{3}{4}$ . The second factor in Eq. (1) simplifies to  $\frac{2^{32} + 1}{2^{33} + 1}$  which is lower bound by  $\frac{1}{2}$ . In total, the probability for finding a tag collision (and hence winning the game CMT) is at least  $\frac{3}{8}$ .  $\square$

<sup>35</sup> For sake of simplicity, we assume that  $\mathcal{A}$  chooses all associated data to have the same number of blocks  $\alpha$ .

<sup>36</sup> More precisely, the probability is  $\frac{q(q+1)}{2^{97}} - \frac{q(q-1)}{2^{129}}$  [13].

The attack exploits the fact that TINYJAMBU uses a very short tag (64 bits) compared to the other schemes—the only other scheme with a 64-bit tag is ELEPHANT, though they also provide a parameter set with a larger tag. Increasing the tag length of TINYJAMBU would render our attack impractical. Note, however, that increasing the tag length to 128 does not make TINYJAMBU committing secure. For such a variant of TINYJAMBU, we can similarly apply a birthday attack to find a collision on the capacity part, while the associated data is processed. Such a 96-bit collision can be found with about  $2^{48}$  queries and, by properly choosing the associated data, results in a full collision. One can modify the parameters such that a 127-bit collision has to be found—though this variant is impractical as the inputs would have to be processed bit by bit.

By construction, the above attack is a  $\text{CMT}_K$  attack, as  $K$  and  $\bar{K}$  were chosen to be different. Moreover, by requiring not only the tuples  $(N_i, A_i)$  to differ for all  $i$ , but the individual nonces and associated data, we also obtain a  $\text{CMT}_N$  and a  $\text{CMT}_A$  attack.

TINYJAMBU.ENC( $K, N, A, M$ )	ENC $_M(K, S, M)$
1 : $S \leftarrow \text{ENC}_e(K, N, A)$	17 : $M_1, \dots, M_\mu \xleftarrow{32} \text{pad}_{0^*}(M, 32)$
2 : $(C, T) \leftarrow \text{ENC}_M(K, S, M)$	18 : <b>for</b> $i = 1, \dots, \mu$
3 : <b>return</b> $(C, T)$	19 : $S \leftarrow S \oplus (0^{64} \parallel \iota_M)$
<hr style="width: 100%; border: 0.5px solid black;"/>	20 : $S \leftarrow \text{BC}_2(K, S)$
4 : $N_1, N_2, N_3 \xleftarrow{32} N$	21 : $S \leftarrow (\lceil S \rceil_{32} \oplus M_i) \parallel \lfloor S \rfloor_{96}$
5 : $A_1, \dots, A_\alpha \xleftarrow{32} \text{pad}_{0^*}(A, 32)$	22 : $C_i \leftarrow \lfloor S \rfloor_{33}^{64}$
6 : $S \leftarrow 0^{128}$	23 : $S \leftarrow S \oplus (0^{64} \parallel \iota_T)$
7 : $S \leftarrow \text{BC}_2(K, S)$	24 : $S \leftarrow \text{BC}_2(K, S)$
8 : <b>for</b> $i = 1, \dots, 3$	25 : $T_i \leftarrow \lfloor S \rfloor_{33}^{64}$
9 : $S \leftarrow S \oplus (0^{64} \parallel \iota_N)$	26 : $S \leftarrow S \oplus (0^{64} \parallel \iota_T)$
10 : $S \leftarrow \text{BC}_1(K, S)$	27 : $S \leftarrow \text{BC}_1(K, S)$
11 : $S \leftarrow (\lceil S \rceil_{32} \oplus N_i) \parallel \lfloor S \rfloor_{96}$	28 : $T_r \leftarrow \lfloor S \rfloor_{33}^{64}$
12 : <b>for</b> $i = 1, \dots, \alpha$	29 : $C \leftarrow \lceil C_1 \parallel \dots \parallel C_\mu \rceil_{ M }$
13 : $S \leftarrow S \oplus (0^{64} \parallel \iota_A)$	30 : $T \leftarrow T_i \parallel T_r$
14 : $S \leftarrow \text{BC}_1(K, S)$	31 : <b>return</b> $(C, T)$
15 : $S \leftarrow (\lceil S \rceil_{32} \oplus A_i) \parallel \lfloor S \rfloor_{96}$	
16 : <b>return</b> $S$	

Fig. 28: Pseudocode of TINYJAMBU [56] in terms of  $\text{ENC}_e$  and  $\text{ENC}_M$ . If the last block of associated data or message is not of full length, TINYJAMBU XORs the respective lengths into the last bits (as part of  $\iota_A$  and  $\iota_M$ ).

## B.7 Proof of Theorem 9 (ISAP)

*Proof (of Theorem 9).* Let  $\mathcal{A}$  be a CMT adversary against ISAP with output denoted by  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ . Further note that  $IV$  denotes the initialization vector used in ISAP. We assume that  $\mathcal{A}$  makes queries to  $\rho_1$  and  $\rho_2$  that correspond to its output, i.e., querying all states that occur during the evaluation of ISAP for the two output tuples of  $\mathcal{A}$ . This assumption is without loss of generality, as we can easily transform any adversary into one that runs  $\mathcal{A}$  to obtain  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$  and—before outputting the same—makes all queries to  $\rho$  corresponding to  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ .

The authentication component  $\text{ENC}_{\mathcal{T}}$  uses a session key denoted by  $K_A$  (resp.  $\overline{K}_A$ ), which results from an application of  $\text{ISAP.RK}$  to the key  $K$  (resp.  $\overline{K}$ ) and the intermediate state  $X$  (resp.  $\overline{X}$ ) computed during  $\text{ENC}_{\mathcal{T}}$ . We consider the event  $\mathbf{E}$  that  $(N, A) = (\overline{N}, \overline{A})$  and  $K_A = \overline{K}_A$ . Using this, the CMT advantage can be divided up as follows

$$\begin{aligned} \text{Adv}_{\text{ISAP}}^{\text{CMT}}(\mathcal{A}) &= \Pr[\text{CMT}(\mathcal{A}) \rightarrow 1] \\ &= \Pr[\mathbf{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] + \Pr[\neg \mathbf{E} \wedge \text{CMT}(\mathcal{A}) \rightarrow 1]. \end{aligned}$$

We start by giving an upper bound for the second summand. For this, we construct a CR (see Definition 12) adversary  $\mathcal{B}$  against a sponge hash function  $\mathcal{H}_2$  obtained from the permutation  $\rho_2$  with rate  $\overline{r}_2 = \max\{\kappa, r_2 + 1\}$ <sup>37</sup>, capacity  $\overline{c}_2 = n - \overline{r}_2$ , and output length  $\kappa$ . Further  $0^\kappa \parallel IV$  is chosen as the initial state of  $\mathcal{H}_2$ . First,  $\mathcal{B}$  runs  $\mathcal{A}$ , which outputs  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ . For every query that  $\mathcal{A}$  makes to  $\rho_2$ , the adversary  $\mathcal{B}$  makes the same query to its own permutation and sends the response back to  $\mathcal{A}$ . Further, Adversary  $\mathcal{B}$  simulates  $\rho_1$  for  $\mathcal{A}$ . Using this,  $\mathcal{B}$  computes for both output tuples of  $\mathcal{A}$ , the state in  $\text{ENC}_{\mathcal{T}}$  after the associated data and the ciphertext blocks are absorbed. We denote the states for the first tuple and second tuple by  $X$  and  $\overline{X}$ , respectively and the session keys for  $\text{ENC}_{\mathcal{T}}$  by  $K_A$  and  $\overline{K}_A$ , respectively (cf. Fig. 29). The states obtained after XORing these together are denoted by  $Z = X \oplus K_A$  and  $\overline{Z} = \overline{X} \oplus \overline{K}_A$ .<sup>38</sup>

Let  $A_1, \dots, A_\alpha \stackrel{r_2}{\leftarrow} \text{pad}_{10^*}(A, r_2)$  and  $\overline{A}_1, \dots, \overline{A}_\alpha \stackrel{r_2}{\leftarrow} \text{pad}_{10^*}(\overline{A}, r_2)$  be the division of  $A$  and  $\overline{A}$  into blocks of length  $r_2$ . Analogously, the ciphertext  $C = \text{ENC}_{\mathcal{M}}(K, N, M) = \text{ENC}_{\mathcal{M}}(\overline{K}, \overline{N}, \overline{M})$  is parsed as  $C_1, \dots, C_\gamma \stackrel{r_2}{\leftarrow} \text{pad}_{10^*}(C, r_2)$ . The adversary  $\mathcal{B}$  then outputs

$$\begin{aligned} O &= (N \parallel 0^*, A_1 \parallel 0^*, \dots, A_\alpha \parallel 0^*, C_1 \parallel 1 \parallel 0^*, \dots, C_\gamma \parallel 0^*, Z \parallel 0^*) \\ \overline{O} &= (\overline{N} \parallel 0^*, \overline{A}_1 \parallel 0^*, \dots, \overline{A}_\alpha \parallel 0^*, C_1 \parallel 1 \parallel 0^*, \dots, C_\gamma \parallel 0^*, \overline{Z} \parallel 0^*), \end{aligned}$$

<sup>37</sup> The definition of the rate over the maximum ensures that the argument works for both ISAP variants (ISAP-K and ISAP-A). More precisely, it guarantees that all inputs can still be fully absorbed after the adjustment of the rate.

<sup>38</sup> Note that  $\mathcal{B}$  can compute these values by looking up the queries and responses from  $\mathcal{A}$ 's queries—using the assumption that it makes permutation queries corresponding to its output. Thus, this step does not require any additional permutation queries by  $\mathcal{B}$ .

where  $\|0^*$  denotes the padding with 0s up to length  $\bar{r}$ . A visualization for this is provided in Fig. 29. We show that if  $\mathcal{A}$  wins the game CMT against ISAP and the event  $\neg E$  holds, then the constructed adversary  $\mathcal{B}$  wins the game CR against  $\mathcal{H}_2$ . Note that  $\mathcal{A}$  winning the game CMT implies that  $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$  and  $\text{ENC}_{\mathcal{T}}(K, N, A, C) = T = \text{ENC}_{\mathcal{T}}(\bar{K}, \bar{N}, \bar{A}, C)$ . Hence, the output tuples of  $\mathcal{B}$  are mapped to the same result under  $\mathcal{H}_2$  (namely  $T$ ) and it is only left to check that  $O \neq \bar{O}$  to guarantee a collision. As event  $\neg E$  holds, we know that  $(N, A) \neq (\bar{N}, \bar{A})$  or  $K_A \neq \bar{K}_A$ . In the case that  $(N, A) \neq (\bar{N}, \bar{A})$ , we have  $O \neq \bar{O}$ . Hence, we can assume from now on that  $(N, A) = (\bar{N}, \bar{A})$ . Next, consider the case that  $K_A \neq \bar{K}_A$ . As  $(N, A) = (\bar{N}, \bar{A})$ , we know that  $X = \bar{X}$  and hence  $Z = [X]_{\kappa} \oplus K_A \neq [\bar{X}]_{\kappa} \oplus \bar{K}_A = \bar{Z}$ . Thus,  $O \neq \bar{O}$  is also given in this case. We have shown that

$$\Pr[\neg E \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{CR}(\mathcal{B}) \rightarrow 1] \leq \frac{q_2(q_2 - 1)}{2^{\kappa}} + \frac{q_2(q_2 + 1)}{2^{n - \bar{r}_2}},$$

where the last inequality holds by Theorem 15, which bounds the probability of finding a collision in a general sponge hash function. Here, we exploit the fact that  $\mathcal{B}$  makes the same number of queries to  $\rho_2$  as  $\mathcal{A}$ .

Next, we give a bound for the first summand  $\Pr[E \wedge \text{CMT}(\mathcal{A}) \rightarrow 1]$ . For this, construct a CR adversary  $\mathcal{C}$  against a sponge hash function  $\mathcal{H}_1$  obtained from the permutation  $\rho_1$  with rate  $\bar{r}_1 = \kappa$ , capacity  $\bar{c}_1 = n - \bar{r}_2$  and output length  $\kappa$ . Further its initial state is given by  $0^{\kappa} \parallel IV$ . The adversary  $\mathcal{C}$  starts by running  $\mathcal{A}$ , which outputs  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . For every query that  $\mathcal{A}$  makes to  $\rho_1$ , the adversary  $\mathcal{C}$  makes the same query to its own permutation and sends the response back to  $\mathcal{A}$ . Further, adversary  $\mathcal{C}$  simulates  $\rho_2$  for  $\mathcal{A}$ . Then, it computes  $X$  and  $\bar{X}$  (analogously to adversary  $\mathcal{B}$ ), and we write  $X_1, \dots, X_{\kappa} \stackrel{\perp}{\leftarrow} X$  and  $\bar{X}_1, \dots, \bar{X}_{\kappa} \stackrel{\perp}{\leftarrow} \bar{X}$ . Lastly, adversary  $\mathcal{C}$  outputs  $(K, X_1 \parallel 0^{\kappa-1}, \dots, X_{\kappa} \parallel 0^{\kappa-1})$  and  $(\bar{K}, \bar{X}_1 \parallel 0^{\kappa-1}, \dots, \bar{X}_{\kappa} \parallel 0^{\kappa-1})$ . Next, we show that if  $\mathcal{A}$  wins the game CMT against ISAP and event  $E$  holds, then the constructed adversary  $\mathcal{C}$  wins the game CR against  $\mathcal{H}_1$ . First observe that  $\mathcal{A}$  winning the game CMT, implies that  $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$ , and  $\text{ISAP.ENC}(K, N, A, M) = \text{ISAP.ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . If at the same time event  $E$  holds, i.e.,  $(N, A) = (\bar{N}, \bar{A})$  and  $K_A = \bar{K}_A$  hold, then  $K \neq \bar{K}$ , as otherwise  $\mathcal{A}$  would not be a valid CMT adversary. Further note that  $X = \bar{X}$  as  $(N, A) = (\bar{N}, \bar{A})$ . Hence,  $\mathcal{C}$  wins the game CR, because the tuples he outputs are different, but their image under  $\mathcal{H}_1$  agrees (as  $K_A = \bar{K}_A$ ). This implies that

$$\Pr[E \wedge \text{CMT}(\mathcal{A}) \rightarrow 1] \leq \Pr[\text{CR}(\mathcal{C}) \rightarrow 1] \leq \frac{q_1(q_1 - 1)}{2^{\kappa}} + \frac{q_1(q_1 + 1)}{2^{n - \kappa}},$$

where the last inequality holds by Theorem 15. Using  $\bar{r}_2 = \max\{\kappa, r_2 + 1\}$ , we obtain in total

$$\text{Adv}_{\text{ISAP}}^{\text{CMT}}(\mathcal{A}) \leq \frac{q_1(q_1 - 1)}{2^{\kappa}} + \frac{q_1(q_1 + 1)}{2^{n - \kappa}} + \frac{q_2(q_2 - 1)}{2^{\kappa}} + \frac{q_2(q_2 + 1)}{2^{n - \max\{\kappa, r_2 + 1\}}},$$

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

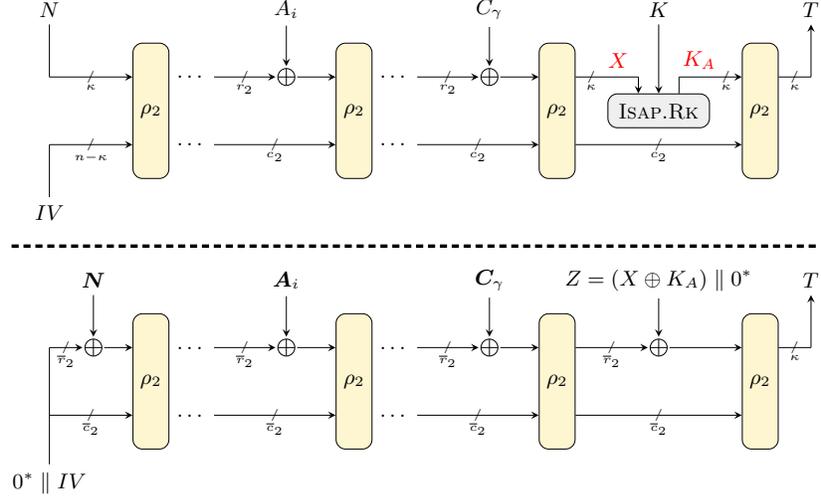


Fig. 29: Illustration of a proof step for ISAP (Theorem 9).. Here,  $\bar{r}_2 = \max\{\kappa, r_H + 1\}$  (i.e., 129 for ISAP-A and 145 for ISAP-K) and  $\bar{c}_2 = n - \bar{r}_2$ ; further write  $\mathbf{N} = N \parallel 0^*$ ,  $\mathbf{A}_i = A_i \parallel 0^*$ ,  $\mathbf{C}_1 = C_1 \parallel 1 \parallel 0^*$  and  $\mathbf{C}_i = C_i \parallel 0^*$  for  $i \in \{2, \dots, \gamma\}$ .

The dominant term in the bound from Theorem 9 is  $\frac{q_1(q_1-1)}{2^\kappa} + \frac{q_2(q_2-1)}{2^\kappa}$ , thus by increasing  $\kappa$  (i.e., the tag and key length), we can increase the committing security. Note however, that—for ISAP-A—we can only increase  $\kappa$  up to 160 as for larger values the other term  $\frac{q_1(q_1+1)}{2^{n-\kappa}} + \frac{q_2(q_2+1)}{2^{n-\max\{\kappa, r_2+1\}}}$  becomes the dominant term. This would result in about 80-bit committing security. A similar argument applies for ISAP-K, which deploys KECCAK-P as the underlying permutation. For this variant, we have  $n = 400$  which allows to increase  $\kappa$  up to 200, allowing for about 100-bit committing security.

ISAP.ENC( $K, N, A, M$ )	ENC $_{\mathcal{M}}$ ( $K, N, M$ )
1 : $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, M)$	19 : $M_1, \dots, M_\mu \xleftarrow{r_2} \text{pad}_{0^*}(M, r_2)$
2 : $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$	20 : $K_E \leftarrow \text{ISAP.RK}(K, N)$
3 : <b>return</b> ( $C, T$ )	21 : $S \leftarrow K_E \parallel N$
	22 : <b>for</b> $i = 1, \dots, \mu$
ENC $_{\mathcal{T}}$ ( $K, N, A, C$ )	23 : $S \leftarrow \rho_E(S)$
4 : $A_1, \dots, A_\alpha \xleftarrow{r_2} \text{pad}_{10^*}(A, r_2)$	24 : $C_i \leftarrow [S]_{r_2} \oplus M_i$
5 : $C_1, \dots, C_\gamma \xleftarrow{r_2} \text{pad}_{10^*}(C, r_2)$	25 : $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$
6 : $Y \leftarrow N \parallel IV$	26 : <b>return</b> $C$
7 : $S \leftarrow \rho_H(Y)$	
8 : <b>for</b> $i = 1, \dots, \alpha$	ISAP.RK( $K, X$ )
9 : $Y \leftarrow S \oplus (A_i \parallel 0^{c_2})$	27 : $X_1, \dots, X_z \xleftarrow{r_1} X$
10 : $S \leftarrow \rho_H(Y)$	28 : $Y \leftarrow K \parallel IV$
11 : $S \leftarrow S \oplus 0^{n-1} \parallel 1$	29 : $S \leftarrow \rho_K(Y)$
12 : <b>for</b> $i = 1, \dots, \gamma$	30 : <b>for</b> $i = 1, \dots, z - 1$
13 : $Y \leftarrow S \oplus (C_i \parallel 0^{c_2})$	31 : $Y \leftarrow S \oplus (X_i \parallel 0^{n-r_1})$
14 : $S \leftarrow \rho_H(Y)$	32 : $S \leftarrow \rho_B(Y)$
15 : $K_A \leftarrow \text{ISAP.RK}(K, [S]_\kappa)$	33 : $Y \leftarrow S \oplus (X_z \parallel 0^{n-r_1})$
16 : $S \leftarrow \rho_H(K_A, [S]_\kappa)$	34 : $S \leftarrow \rho_K(Y)$
17 : $T \leftarrow [S]_\tau$	35 : <b>return</b> $[S]_\kappa$
18 : <b>return</b> $T$	

Fig. 30: Pseudocode of ISAP [29] in terms of ENC $_{\mathcal{M}}$  and ENC $_{\mathcal{T}}$ .

## B.8 Proof of Theorem 10 (ASCON)

*Proof (of Theorem 10).* Let  $\mathcal{A}$  be a CMT adversary against ASCON with output denoted by  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . Further note that  $IV$  denotes the initialization vector used in ASCON. As a first step, we observe that finding different inputs to the ASCON encryption that give the same ciphertext is easy due to the duplex construction used in the sponge. The difficulty in breaking CMT security for ASCON lies in finding a tag collision, which is why we focus our attention on this task. An adversary that wins the game CMT against ASCON, in particular finds a tag collision, i.e., it wins the game  $\text{TagColl}$  (see Fig. 16), which allows the following reduction step

$$\mathbf{Adv}_{\text{ASCON}}^{\text{CMT}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{ASCON}}^{\text{TagColl}}(\mathcal{A}).$$

To bound the advantage on the right, we adapt the proof of Theorem 15 to fit our particular needs.

We build a directed graph  $G$  from the ideal permutation queries the adversary makes, in the following way: The nodes in  $G$  are the  $2^{320}$  bit strings of length 320 and an edge from  $Y$  to  $S$  is added if  $\mathcal{A}$  makes a query of the form  $\rho(Y) = S$  or  $\rho^{-1}(S) = Y$  (the graph starts with no edges). The edges resulting from  $\rho$  queries are called *forward edges* and the ones resulting from  $\rho^{-1}$  queries are referred to as *backward edges*.

We assume  $\mathcal{A}$  to make queries to  $\rho$  that correspond to its output, i.e., querying all states that occur during the evaluation of ASCON for the two output tuples of  $\mathcal{A}$ . This assumption is without loss of generality, as we can easily transform any adversary into one that runs  $\mathcal{A}$  to obtain  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$  and—before outputting the same—makes all queries to  $\rho$  corresponding to  $(K, N, A, M)$  and  $(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ . Additionally, we assume  $\mathcal{A}$  to make no redundant queries, i.e., once two values  $(Y, S)$  are known to be connected via an edge, no further  $\rho$  queries are made on  $Y$  and no further  $\rho^{-1}$  queries are made on  $S$ .

For this graph, we define a special kind of path, the *A-path*, which models an ASCON evaluation. An A-path  $P_A$  of length  $l$ <sup>39</sup> is a sequence of  $2l$  nodes

$$Y_0, S_1, Y_1, S_2, \dots, S_{l-1}, Y_{l-1}, S_l$$

with

1.  $Y_0 = K \parallel N \parallel IV$ ,
2.  $\lfloor Y_1 \rfloor_{256} = \lfloor S_1 \rfloor_{256} \oplus (K \parallel 0^{128})$  and  $\lfloor Y_{l-1} \rfloor_{256} = \lfloor S_{l-1} \rfloor_{256} \oplus (K \parallel 0^{128})$ ,
3.  $\lfloor Y_i \rfloor_{256} = \lfloor S_i \rfloor_{256}$  for all  $i \in \{2, \dots, l-2\}$ , and
4.  $G$  contains edges from  $Y_{i-1}$  to  $S_i$  for all  $i \in \{1, \dots, l\}$ ,

for some  $K, N \in \{0, 1\}^{128}$ . We define the *input* of an A-path  $P_A$  as  $I := (K, N, X_1, \dots, X_{l-1}) \in \{0, 1\}^{128} \times \{0, 1\}^{128} \times \{0, 1\}^{64} \dots \times \{0, 1\}^{64}$  for  $K \parallel N :=$

<sup>39</sup> Note that  $l \geq 3$ , as ASCON involves at least three applications of  $\rho$  (in case both associated data and message consist of a single block).

$[Y_0]_{256}$  and  $X_i = [S_i]_{64} \oplus [Y_i]_{64}$  for all  $i \in \{1, \dots, l-1\}$ . The *result* of  $P_A$  is defined as  $R = [S_l]_{128} \oplus K$ . By construction, this models the tag generation of ASCON for key  $K$ , nonce  $N$ , and the tuple of associated data and message  $(A, M) = (X_1, \dots, X_{l-1})$ . As a notation for A-paths that incorporates the input, we write

$$(K \parallel N) | Y_0 \rightarrow \dots \rightarrow S_{l-2} | X_{l-2} | Y_{l-2} \rightarrow S_{l-1} | X_{l-1} | Y_{l-1} \rightarrow S_l. \text{ }^{40}$$

Next, we define two properties a pair  $(P_A, \bar{P}_A)$  of A-paths can have. For this, denote the nodes in  $\bar{P}_A$  by  $\bar{Y}_0, \bar{S}_1, \bar{Y}_1, \dots, \bar{Y}_{l-1}, \bar{S}_l$ . Firstly, the paths  $P_A$  and  $\bar{P}_A$  are *colliding* if their inputs differ but their results agree. Secondly, the paths  $P_A$  and  $\bar{P}_A$  are *problematic* if their inputs differ and

1.  $Y_{l-1} = \bar{Y}_{l-1}$  or
2. at least one of the edges in  $P$  or  $\bar{P}$  is a backward edge.

We are interested in the event CP that  $\mathcal{A}$  finds a pair of colliding paths. Note that finding such paths means that  $\mathcal{A}$  wins the game **TagColl**. In order to compute the probability of CP, we define the auxiliary event  $\text{PP}_A$  that  $\mathcal{A}$  finds a pair of problematic A-paths. Using this, we obtain

$$\mathbf{Adv}_{\text{ASCON}}^{\text{TagColl}}(\mathcal{A}) = \Pr[\text{CP}] \leq \Pr[\text{CP} \wedge \neg \text{PP}_A] + \Pr[\text{PP}_A],$$

and proceed by deriving upper bounds for both of the above summands.

We start with the easier case, which is giving an upper bound for the probability that  $\text{CP} \wedge \neg \text{PP}_A$  holds, i.e., that  $\mathcal{A}$  finds a pair of colliding A-paths that is not problematic. Hence,  $\mathcal{A}$  finds two different inputs  $I = (K, N, X_1, \dots, X_{l-1})$  and  $\bar{I} = (\bar{K}, \bar{N}, \bar{X}_1, \dots, \bar{X}_{l-1})$  such that the corresponding A-paths

$$\begin{array}{c} Y_0, S_1, Y_1, S_2, \dots, S_{l-1}, Y_{l-1}, S_l \\ \bar{Y}_0, \bar{S}_1, \bar{Y}_1, \bar{S}_2, \dots, \bar{S}_{l-1}, \bar{Y}_{l-1}, \bar{S}_l \end{array}$$

fulfill  $Y_{l-1} \neq \bar{Y}_{l-1}$  and have equal results, i.e.,  $[\rho(Y_{l-1})]_{128} \oplus K = R = \bar{R} = [\rho(\bar{Y}_{l-1})]_{128} \oplus \bar{K}$ . By definition of an A-path, this implies

$$[\rho(S_{l-1} \oplus (X_{l-1} \parallel K \parallel 0^{128}))]_{128} \oplus K = [\rho(\bar{S}_{l-1} \oplus (\bar{X}_{l-1} \parallel \bar{K} \parallel 0^{128}))]_{128} \oplus \bar{K}. \quad (2)$$

Since  $\rho$  is a random permutation and  $\mathcal{A}$  only used forward queries (as  $\neg \text{PP}_A$  holds), finding such a collision is unlikely. We assume—to the benefit of the adversary  $\mathcal{A}$ —that it can choose  $S_{l-1}, \bar{S}_{l-1} \in \{0, 1\}^{128}$  freely, i.e., it must not be part of an A-path for some input. The probability of  $\mathcal{A}$  finding  $(S_{l-1}, X_{l-1}, K)$ ,

<sup>40</sup> While not visible in this representation, by definition of A-paths,  $Y_1$  and  $S_1$  (respectively  $Y_{l-1}$  and  $S_{l-1}$ ) differ not only in their first 64 bits but also from bit 65 to 192, where the key is XORed.

$(\bar{S}_{\bar{l}-1}, \bar{X}_{\bar{l}-1}, \bar{K})$ <sup>41</sup> such that  $Y_{l-1} \neq \bar{Y}_{\bar{l}-1}$  and Eq. (2) holds with  $q$  queries, equals the probability of finding a collision in a list of  $q$  uniformly distributed elements. Using Theorem 17 for  $q \leq 2^{127}$ , the latter can be bounded from above by

$$1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right).$$

Next, we turn our attention to deriving an upper bound for  $\Pr[\text{PP}_A]$ . We will see that finding problematic paths is hard, even for a plain sponge without ASCON's blinding mechanisms, which is why we reduce to this setting. We consider a sponge-based hash function  $\mathcal{H}$  obtained from the permutation  $\rho$  with rate 256 for the first round of absorption and rate 196 for all remaining ones. Further, its initial state is given by  $0^{256} \parallel IV$  and the output produced by  $\mathcal{H}$  has length 128. Analogously to A-paths, we also model evaluations of  $\mathcal{H}$  by paths in the directed graph  $G$ . For  $s \geq 1$ , a *PS-path* of length  $s$  is a sequence of  $2s$  nodes

$$Y_0, S_1, Y_1, S_2, \dots, S_{s-1}, Y_{s-1}, S_s$$

with

1.  $\lfloor Y_0 \rfloor_{64} = IV$ ,
2.  $\lfloor Y_i \rfloor_{128} = \lfloor S_i \rfloor_{128}$  for all  $i \in \{1, \dots, s-1\}$ , and
3.  $G$  contains edges from  $Y_{i-1}$  to  $S_i$  for all  $i \in \{1, \dots, s\}$ .

We define the *input* of a PS-path as  $I := (Z_0, \dots, Z_{s-1}) \in \{0, 1\}^{256} \times \{0, 1\}^{192} \times \dots \times \{0, 1\}^{192}$  for  $Z_0 = \lceil Y_0 \rceil_{256}$  and  $Z_i = \lceil S_i \rceil_{192} \oplus \lceil Y_i \rceil_{192}$  for all  $i \in \{1, \dots, s-1\}$ . As a notation for PS-paths that incorporates the input, we write

$$Z_0 | Y_0 \rightarrow \dots \rightarrow S_{s-2} | Z_{s-2} | Y_{s-2} \rightarrow S_{s-1} | Z_{s-1} | Y_{s-1} \rightarrow S_s.$$

The notion of problematic A-paths can be directly transferred to PS-paths. The event that  $\mathcal{A}$  finds a pair of problematic PS-paths is denoted by  $\text{PP}_{\text{ps}}$ .

We next observe that a pair of problematic A-paths, can also be considered as a pair of problematic PS-paths, i.e., in particular the event  $\text{PP}_A$  implies the event  $\text{PP}_{\text{ps}}$ . Let  $(P_A, \bar{P}_A)$  be a pair of problematic A-paths, i.e.,

$$\begin{aligned} P_A &= (K \parallel N) | Y_0 \rightarrow \dots \rightarrow S_{l-2} | X_{l-2} | Y_{l-2} \rightarrow S_{l-1} | X_{l-1} | Y_{l-1} \rightarrow S_l \\ \bar{P}_A &= (\bar{K} \parallel \bar{N}) | \bar{Y}_0 \rightarrow \dots \rightarrow \bar{S}_{\bar{l}-2} | \bar{X}_{\bar{l}-2} | \bar{Y}_{\bar{l}-2} \rightarrow \bar{S}_{\bar{l}-1} | \bar{X}_{\bar{l}-1} | \bar{Y}_{\bar{l}-1} \rightarrow \bar{S}_{\bar{l}}. \end{aligned}$$

By defining

$$\begin{aligned} Z_0 &= K \parallel N \in \{0, 1\}^{256} & \bar{Z}_0 &= \bar{K} \parallel \bar{N} \in \{0, 1\}^{256} \\ Z_1 &= X_1 \parallel K \in \{0, 1\}^{192} & \bar{Z}_1 &= \bar{X}_1 \parallel \bar{K} \in \{0, 1\}^{192} \\ Z_i &= X_i \parallel 0^{128} \in \{0, 1\}^{192} & \bar{Z}_i &= \bar{X}_i \parallel 0^{128} \in \{0, 1\}^{192} \\ Z_{l-1} &= X_{l-1} \parallel K \in \{0, 1\}^{192} & \bar{Z}_{\bar{l}-1} &= \bar{X}_{\bar{l}-1} \parallel \bar{K} \in \{0, 1\}^{192} \end{aligned}$$

<sup>41</sup> Note that,  $\mathcal{A}$  must choose  $(S_{l-1}, X_{l-1}, K) \neq (\bar{S}_{\bar{l}-1}, \bar{X}_{\bar{l}-1}, \bar{K})$  to ensure  $Y_{l-1} \neq \bar{Y}_{\bar{l}-1}$ .

we obtain the following presentation of  $(P_A, \bar{P}_A)$  as PS-paths:

$$\begin{aligned} P_{\text{PS}} &= Z_0|Y_0 \rightarrow \cdots \rightarrow S_{l-2}|Z_{l-2}|Y_{l-2} \rightarrow S_{l-1}|Z_{l-1}|Y_{l-1} \rightarrow S_l \\ \bar{P}_{\text{PS}} &= \bar{Z}_0|\bar{Y}_0 \rightarrow \cdots \rightarrow \bar{S}_{l-2}|\bar{Z}_{l-2}|\bar{Y}_{l-2} \rightarrow \bar{S}_{l-1}|\bar{Z}_{l-1}|\bar{Y}_{l-1} \rightarrow \bar{S}_l. \end{aligned}$$

Visualization for this is provided in Fig. 31. As we neither change  $(Y_{l-1}, \bar{Y}_{l-1})$  nor any of the edges, the paths  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$  form a pair of problematic PS-paths. Thus, we have shown that  $\text{PP}_A$  implies  $\text{PP}_{\text{PS}}$ , hence  $\Pr[\text{PP}_A] \leq \Pr[\text{PP}_{\text{PS}}]$ . This allows us to focus on the plain sponge setting for the rest of the proof. More precisely, we show that it is hard to find a pair of problematic PS-paths, i.e., we derive an upper bound for  $\Pr[\text{PP}_{\text{PS}}]$ .

For this, we define the following auxiliary events:

1. Event  $E_t$  (target hitting query):  
 $\mathcal{A}$  makes a query  $Y$  to  $\rho$  such that  $\lfloor \rho(Y) \rfloor_{64} = IV$  or  $\mathcal{A}$  makes a query  $S$  to  $\rho^{-1}$  such that  $\lfloor \rho^{-1}(S) \rfloor_{64} = IV$ .
2. Event  $E_c$  (colliding queries):  
 $\mathcal{A}$  makes queries  $Y \neq \bar{Y}$  to  $\rho$  such that  $\lfloor \rho(Y) \rfloor_{128} = \lfloor \rho(\bar{Y}) \rfloor_{128}$  or  $\mathcal{A}$  makes queries  $Y$  to  $\rho$  and  $\bar{S}$  to  $\rho^{-1}$  such that  $\lfloor \rho(Y) \rfloor_{128} = \lfloor \rho^{-1}(\bar{S}) \rfloor_{128}$ .

Next, we show that if  $\mathcal{A}$  triggers  $\text{PP}_{\text{PS}}$ , then it triggers one of the events defined above.

For the proof assume that  $\text{PP}_{\text{PS}}$  holds and denote the problematic paths  $\mathcal{A}$  finds by  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$ . We first consider the case that there is at least one backward edge in  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$ . We assume w.l.o.g. that  $P_{\text{PS}}$  contains at least one backward edge. Note that for each PS-path that contains at least one backward edge, we can define a corresponding minimal PS-path containing exactly one backward edge. To do this—starting from the end of the path—all edges are removed, until the last edge of the path is a backward edge and all other edges (if any remain) are forward edges. For sake of simplicity, we write  $P_{\text{PS}}$  also for the minimal path corresponding to  $P_{\text{PS}}$  in the following and write it as

$$P_{\text{PS}} = Z_0|Y_0 \rightarrow \cdots \rightarrow S_{s-2}|Z_{s-2}|Y_{s-2} \rightarrow S_{s-1}|Z_{s-1}|Y_{s-1} \rightarrow S_s.$$

We further distinguish the following two sub-cases:

**Case 1:**  $s = 1$

The path is simply  $Z_0|Y_0 \rightarrow S_1$  and  $\mathcal{A}$  queried  $S_1$  to  $\rho^{-1}$ . By construction, we have  $\lfloor Y_0 \rfloor_{64} = IV$  and  $\rho^{-1}(S_1) = Y_0$ , hence in particular  $\lfloor \rho^{-1}(S_1) \rfloor_{64} = \lfloor Y_0 \rfloor_{64} = IV$ . Thus  $\mathcal{A}$ 's query triggered event  $E_t$ .

**Case 2:**  $s \geq 2$

The path is  $Z_0|Y_0 \rightarrow \cdots \rightarrow S_{s-2}|Z_{s-2}|Y_{s-2} \rightarrow S_{s-1}|Z_{s-1}|Y_{s-1} \rightarrow S_s$ . Except for the last edge, all edges are forward edges. By construction, we have  $\lfloor S_{s-1} \rfloor_{128} = \lfloor Y_{s-1} \rfloor_{128}$ . Furthermore,  $S_{s-1}$  is the result of querying  $Y_{s-2}$  to  $\rho$  (forward edge) and  $Y_{s-1}$  is the result of querying  $S_s$  to  $\rho^{-1}$  (backward edge). This yields that these two queries trigger event  $E_c$ .

We now consider the case that the penultimate states  $Y_{s-1}$  and  $\bar{Y}_{\bar{s}-1}$  are equal, but  $P_{\text{PS}}$  and  $\bar{P}_{\text{PS}}$  contain no backward edges. For such a pair of paths  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$ , we define the corresponding minimal pair of PS-paths by choosing  $s + \bar{s}$  minimal such that  $(Z_0, \dots, Z_{s-1}) \neq (\bar{Z}_0, \dots, \bar{Z}_{\bar{s}-1})$  and  $Y_{s-1} = \bar{Y}_{\bar{s}-1}$  still hold. We consider the minimal pair of paths corresponding to  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$  and—for sake of simplicity—also denote them by  $(P_{\text{PS}}, \bar{P}_{\text{PS}})$ . As before we use the following representation

$$\begin{aligned} P_{\text{PS}} &= Z_0|Y_0 \rightarrow \dots \rightarrow S_{s-2}|Z_{s-2}|Y_{s-2} \rightarrow S_{s-1}|Z_{s-1}|Y_{s-1} \rightarrow S_s \\ \bar{P}_{\text{PS}} &= \bar{Z}_0|\bar{Y}_0 \rightarrow \dots \rightarrow \bar{S}_{\bar{s}-2}|\bar{Z}_{\bar{s}-2}|\bar{Y}_{\bar{s}-2} \rightarrow \bar{S}_{\bar{s}-1}|\bar{Z}_{\bar{s}-1}|\bar{Y}_{\bar{s}-1} \rightarrow \bar{S}_{\bar{s}}. \end{aligned}$$

Without loss of generality, we further assume  $s \leq \bar{s}$  and distinguish between the following three sub-cases:

**Case 1:**  $s = 1 \wedge \bar{s} = 1$

The paths are simply  $Z_0|Y_0 \rightarrow S_1$  and  $\bar{Z}_0|\bar{Y}_0 \rightarrow \bar{S}_1$ . By construction, we have  $[Y_0]_r = Z_0$  and  $[\bar{Y}_0]_r = \bar{Z}_0$  which leads to a contradiction as  $Z_0 \neq \bar{Z}_0$  (recall that problematic paths have different inputs) but  $Y_0 = \bar{Y}_0$ . Thus, this case cannot occur.

**Case 2:**  $s = 1 \wedge \bar{s} \geq 2$

The paths are

$$\begin{aligned} Z_0|Y_0 &\rightarrow S_1 \\ \bar{Z}_0|\bar{Y}_0 &\rightarrow \dots \rightarrow \bar{S}_{\bar{s}-2}|\bar{Z}_{\bar{s}-2}|\bar{Y}_{\bar{s}-2} \rightarrow \bar{S}_{\bar{s}-1}|\bar{Z}_{\bar{s}-1}|\bar{Y}_{\bar{s}-1} \rightarrow \bar{S}_{\bar{s}}. \end{aligned}$$

We have  $[Y_0]_{64} = IV$  by construction and  $Y_0 = \bar{Y}_{\bar{s}-1}$  by assumption. This allows to deduce  $[\bar{S}_{\bar{s}-1}]_{64} = IV$ . Since  $\bar{S}_{\bar{s}-1}$  is the result of querying  $\bar{Y}_{\bar{s}-2}$  to  $\rho$ , this query triggered event  $E_t$ .

**Case 3:**  $s \geq 2 \wedge \bar{s} \geq 2$

The paths are

$$\begin{aligned} Z_0|Y_0 &\rightarrow \dots \rightarrow S_{s-2}|Z_{s-2}|Y_{s-2} \rightarrow S_{s-1}|Z_{s-1}|Y_{s-1} \rightarrow S_s \\ \bar{Z}_0|\bar{Y}_0 &\rightarrow \dots \rightarrow \bar{S}_{\bar{s}-2}|\bar{Z}_{\bar{s}-2}|\bar{Y}_{\bar{s}-2} \rightarrow \bar{S}_{\bar{s}-1}|\bar{Z}_{\bar{s}-1}|\bar{Y}_{\bar{s}-1} \rightarrow \bar{S}_{\bar{s}}. \end{aligned}$$

Consider the penultimate edges of both paths, i.e., the edges from  $Y_{s-2}$  to  $S_{s-1}$  and from  $\bar{Y}_{\bar{s}-2}$  to  $\bar{S}_{\bar{s}-1}$ , which are both forward edges. By assumption,  $Y_{s-1} = \bar{Y}_{\bar{s}-1}$  holds. We have to distinguish two more cases based on the inputs  $Z_{s-1}$  and  $\bar{Z}_{\bar{s}-1}$ :

**Case 3.1:**  $Z_{s-1} = \bar{Z}_{\bar{s}-1}$

From  $Y_{s-1} = \bar{Y}_{\bar{s}-1}$  and  $Z_{s-1} = \bar{Z}_{\bar{s}-1}$ , we can conclude that  $S_{s-1} = \bar{S}_{\bar{s}-1}$ . Then, however, one can obtain a pair of shorter paths by dropping the last edges of both  $P$  and  $\bar{P}$  while maintaining the desired property. Thus this case is impossible as it contradicts the minimality of the paths.

**Case 3.2:**  $Z_{s-1} \neq \bar{Z}_{\bar{s}-1}$

As  $Y_{s-1} = \bar{Y}_{\bar{s}-1}$  and  $Z_{s-1} \neq \bar{Z}_{\bar{s}-1}$ , we can deduce  $S_{s-1} \neq \bar{S}_{\bar{s}-1}$  and  $[S_{s-1}]_{128} = [\bar{S}_{\bar{s}-1}]_{128}$ . Then  $Y_{s-2} \neq \bar{Y}_{\bar{s}-2}$  and the queries on  $Y_{s-2}$  and  $\bar{Y}_{\bar{s}-2}$  triggered event  $E_c$ .

Collecting the above, yields

$$\Pr[\text{PP}_A] \leq \Pr[\mathbf{E}_t \vee \mathbf{E}_c] \leq \Pr[\mathbf{E}_t] + \Pr[\mathbf{E}_c].$$

We apply [Lemma 19](#) for  $c = 2^{64}$  and obtain  $\Pr[\mathbf{E}_t] \leq \frac{q}{2^{63}}$  and once again for  $c = 2^{128}$ , which gives  $\Pr[\mathbf{E}_c] \leq \frac{q(q-1)}{2^{128}}$ . In total, we have shown

$$\begin{aligned} \mathbf{Adv}_{\text{ASCON}}^{\text{CMT}}(\mathcal{A}) &\leq \mathbf{Adv}_{\text{ASCON}}^{\text{TagColl}}(\mathcal{A}) \\ &\leq \Pr[\text{CP} \wedge \neg \text{PP}_A] + \Pr[\text{PP}_A] \\ &\leq \Pr[\text{CP} \wedge \neg \text{PP}_A] + \Pr[\mathbf{E}_t] + \Pr[\mathbf{E}_c] \\ &\leq 1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right) + \frac{q}{2^{63}} + \frac{q(q-1)}{2^{128}}, \end{aligned}$$

which finishes the proof.  $\square$

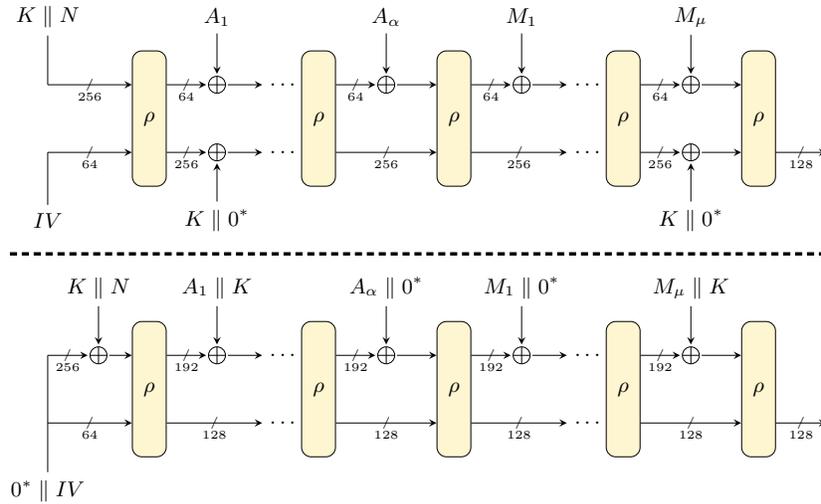


Fig. 31: Illustration of a proof step for ASCON ([Theorem 10](#)).

ASCON.ENC( $K, N, A, M$ )	ENC <sub>M</sub> ( $K, S, M$ )
1: $S \leftarrow \text{ENC}_e(K, N, A)$ 2: $(C, T) \leftarrow \text{ENC}_M(K, S, M)$ 3: <b>return</b> $(C, T)$	13: $M_1, \dots, M_\mu \xleftarrow{r} \text{pad}_{10^*}(M, r)$ 14: $Y \leftarrow S \oplus (M_1 \parallel 0^c)$ 15: $C_1 \leftarrow \lceil Y \rceil_r$ 16: <b>for</b> $i = 2, \dots, \mu$ 17: $S \leftarrow \rho^b(Y)$ 18: $Y \leftarrow S \oplus (M_i \parallel 0^c)$ 19: $C_i \leftarrow \lceil Y \rceil_r$ 20: $C \leftarrow \lceil C_1 \parallel \dots \parallel C_\mu \rceil_{ M }$ 21: $Y \leftarrow Y \oplus (0^r \parallel K \parallel 0^{c-\kappa})$ 22: $S \leftarrow \rho^a(Y)$ 23: $T \leftarrow \lfloor S \rfloor_\tau \oplus K$ 24: <b>return</b> $(C, T)$
ENC <sub>e</sub> ( $K, N, A$ )	
4: $A_1, \dots, A_\alpha \xleftarrow{r} \text{pad}_{10^*}(A, r)$ 5: $Y \leftarrow IV \parallel K \parallel N$ 6: $S \leftarrow \rho^a(Y)$ 7: $S \leftarrow S \oplus 0^{n-\kappa} \parallel K$ 8: <b>for</b> $i = 1, \dots, \alpha$ 9: $Y \leftarrow S \oplus (A_i \parallel 0^c)$ 10: $S \leftarrow \rho^b(Y)$ 11: $S \leftarrow S \oplus 0^{n-1} \parallel 1$ 12: <b>return</b> $S$	

Fig. 32: Pseudocode of ASCON [31] in terms of ENC<sub>e</sub> and ENC<sub>M</sub>.

## B.9 Proof of Theorem 11 (SCHWAEMM)

*Proof.* Firstly, we observe that finding different inputs to  $\text{SCHWAEMM}_{IV}.\text{ENC}$  that result in the same ciphertext is easy. However, breaking CMT security also includes finding colliding tags, which is what we focus on in the following. An adversary  $\mathcal{A}$  that wins the game CMT against  $\text{SCHWAEMM}_{IV}$  its output denoted by  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ , in particular finds a tag collision, i.e., wins the game  $\text{TagColl}$  (see Fig. 16). Hence we can deduce

$$\mathbf{Adv}_{\text{SCHWAEMM}_{IV}}^{\text{CMT}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{SCHWAEMM}_{IV}}^{\text{TagColl}}(\mathcal{A}).$$

As a next step, we pass over to a plain sponge construction. For this, we construct a  $\text{ShiftedColl}_{128}$  adversary  $\mathcal{B}$  against a sponge hash function  $\mathcal{H}$  obtained from the permutation  $\rho$  with rate 256, capacity 128 and output length 128. Further, its initial state is given by  $0^{256} \parallel IV$ . First,  $\mathcal{B}$  runs  $\mathcal{A}$ , which outputs  $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$ . For every query that  $\mathcal{A}$  makes to  $\rho$ , the adversary  $\mathcal{B}$  makes the same query to its own permutation and sends the response back to  $\mathcal{A}$ . Then  $\mathcal{B}$  computes the state  $S_i$  (and respectively  $\overline{S}_i$ ) after the  $i$ -th application of the permutation in  $\text{SCHWAEMM}_{IV}$  evaluated in  $(K, N, A, M)$  (and respectively  $(\overline{K}, \overline{N}, \overline{A}, \overline{M})$ ). Denote by  $S_{i,r}$  and  $S_{i,c}$  (and respectively  $\overline{S}_{i,r}$  and  $\overline{S}_{i,c}$ ) the rate and the capacity part of the state  $S_i$  (and respectively  $\overline{S}_i$ ). The adversary  $\mathcal{B}$  then outputs

$$\begin{aligned} X &= K \parallel N \parallel \\ &\quad (S_{1,r} \oplus \xi_2(S_{1,r}, A_1) \oplus \omega_{c,r}(S_{1,c})) \parallel \dots \parallel \\ &\quad (S_{\alpha,r} \oplus \xi_2(S_{\alpha,r}, A_\alpha) \oplus \omega_{c,r}(S_{\alpha,c})) \parallel \\ &\quad (S_{\alpha+1,r} \oplus \xi_2(S_{\alpha+1,r}, M_1) \oplus \omega_{c,r}(S_{\alpha+1,c})) \parallel \dots \parallel \\ &\quad (S_{\alpha+\mu,r} \oplus \xi_2(S_{\alpha+\mu,r}, M_\mu) \oplus \omega_{c,r}(S_{\alpha+\mu,c})) \\ \overline{X} &= \overline{K} \parallel \overline{N} \parallel \\ &\quad (\overline{S}_{1,r} \oplus \xi_2(\overline{S}_{1,r}, \overline{A}_1) \oplus \omega_{c,r}(\overline{S}_{1,c})) \parallel \dots \parallel \\ &\quad (\overline{S}_{\overline{\alpha},r} \oplus \xi_2(\overline{S}_{\overline{\alpha},r}, \overline{A}_{\overline{\alpha}}) \oplus \omega_{c,r}(\overline{S}_{\overline{\alpha},c})) \parallel \\ &\quad (\overline{S}_{\overline{\alpha}+1,r} \oplus \xi_2(\overline{S}_{\overline{\alpha}+1,r}, \overline{M}_1) \oplus \omega_{c,r}(\overline{S}_{\overline{\alpha}+1,c})) \parallel \dots \parallel \\ &\quad (\overline{S}_{\overline{\alpha}+\overline{\mu},r} \oplus \xi_2(\overline{S}_{\overline{\alpha}+\overline{\mu},r}, \overline{M}_{\overline{\mu}}) \oplus \omega_{c,r}(\overline{S}_{\overline{\alpha}+\overline{\mu},c})), \end{aligned}$$

which guarantees that  $\mathcal{H}(X)$  (and respectively  $\mathcal{H}(\overline{X})$ ) models  $\text{SCHWAEMM}_{IV}$  evaluated on  $(K, N, A, M)$  (and  $(\overline{K}, \overline{N}, \overline{A}, \overline{M})$ , respectively). More precisely, instead of the state-update-function and rate-whitening applied in  $\text{SCHWAEMM}_{IV}$ , for  $\mathcal{H}$  we XOR a suitable value which imitates these operations. A visualization for this is provided in Fig. 33.

Next, we show that if  $\mathcal{A}$  wins the game  $\text{TagColl}$  against  $\text{SCHWAEMM}_{IV}$ , then the constructed adversary  $\mathcal{B}$  wins the game  $\text{ShiftedColl}_{128}$  against  $\mathcal{H}$ . First observe that  $\mathcal{A}$  winning implies that  $(K, N, A, M) \neq (\overline{K}, \overline{N}, \overline{A}, \overline{M})$  and the corresponding tags  $T, \overline{T}$ —computed with  $\text{SCHWAEMM}_{IV}$ —agree. Note that the latter implies that  $S_{\alpha+\mu+1,c} \oplus K = \overline{S}_{\overline{\alpha}+\overline{\mu}+1,c} \oplus \overline{K}$ , hence by choice of  $X$  and  $\overline{X}$  it holds that  $\mathcal{H}(X) \oplus [X]_{128} = \mathcal{H}(\overline{X}) \oplus [\overline{X}]_{128}$ . Further, the fact that

$(K, N, A) \neq (\overline{K}, \overline{N}, \overline{A})$ , implies that  $X \neq \overline{X}$ : for  $(K, N) \neq (\overline{K}, \overline{N})$  this is obvious while for  $(K, N) = (\overline{K}, \overline{N})$  and  $A \neq \overline{A}$ , a simple analysis shows that  $X$  and  $\overline{X}$  differ at the point where the associated data blocks differ for the first time. This implies that  $\mathcal{B}$  wins the game  $\text{ShiftedColl}_{128}$ .

Using the indifferentiability of sponges (cf. [Theorem 16](#)), we can replace  $\mathcal{H}$  by a random function  $F$ , as there exists an efficient simulator for the underlying permutation such that  $\mathcal{A}$  cannot distinguish between  $\mathcal{H}$  and  $F$ . This yields

$$\mathbf{Adv}_{\mathcal{H}}^{\text{ShiftedColl}_{128}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{ShiftedColl}_{128}}(\mathcal{A}) + \epsilon,$$

for  $\epsilon > \frac{(1-2^{-256})q^2 + (1+2^{-256})q}{2^{129}}$ , which results from the application of [Theorem 16](#).

As a last step, we observe that for a random oracle  $F : \{0, 1\}^{\geq 128} \rightarrow \{0, 1\}^{128}$ , it is unlikely that  $\mathcal{B}$  wins the game  $\text{ShiftedColl}_{128}$ . For this, note that an adversary that wins the game  $\text{ShiftedColl}_{128}$  against  $F$  with  $q$  queries, finds a collision in the following list of uniformly distributed elements

$$L = \{F(X_1) \oplus [X_1]_{128}, F(X_2) \oplus [X_2]_{128}, \dots, F(X_q) \oplus [X_q]_{128}\},$$

for  $X_i \in \{0, 1\}^{\geq 128}$  being the inputs  $\mathcal{B}$  queries to  $F$ . By [Theorem 17](#), the probability for this is bounded above by

$$1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right),$$

for  $q \leq 2^{127}$ . In total, we obtain

$$\begin{aligned} \mathbf{Adv}_{\mathcal{H}}^{\text{ShiftedColl}_{128}}(\mathcal{A}) &\leq \mathbf{Adv}_F^{\text{ShiftedColl}_{128}}(\mathcal{A}) + \epsilon \\ &\leq 1 - \exp\left(\frac{-q(q-1)}{2^{128}}\right) + \epsilon, \end{aligned}$$

which finishes the proof of the theorem. □

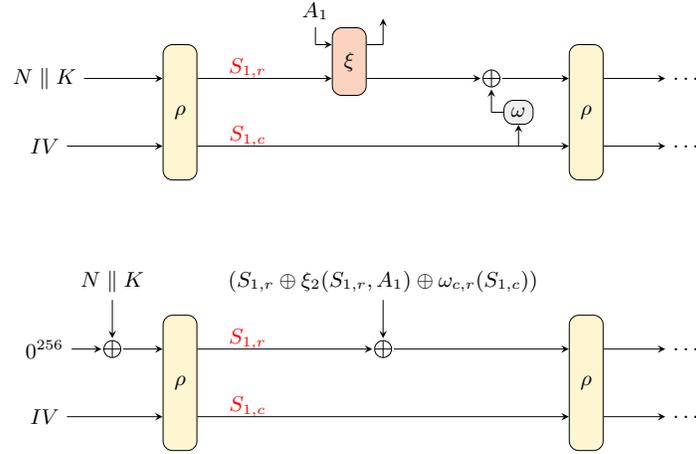


Fig. 33: Illustration of a proof step for  $\text{SCHWAEMM}_{IV}$  (Theorem 11).  $\text{SCHWAEMM}_{IV}$  is represented as a plain sponge as is shown for the first XOR in the above figure.

Game  $\text{ShiftedColl}_\kappa$

- 
- 1:  $X, \bar{X} \leftarrow \mathcal{A}()$
  - 2: **if**  $X = \bar{X}$
  - 3:     **return** 0
  - 4: **return**  $(F(X) \oplus F(\bar{X})) = [X]_\kappa \oplus [\bar{X}]_\kappa$

Fig. 34: Security game  $\text{ShiftedColl}_\kappa$  defined for a function  $F: \{0, 1\}^{\geq \kappa} \rightarrow \{0, 1\}^\kappa$  and used in the proof for  $\text{SCHWAEMM}_{IV}$  (Theorem 11).

SCHWAEMM.ENC( $K, N, A, M$ )	ENC $_M(K, S, M)$
<pre> 1 : <math>S \leftarrow \text{ENC}_e(K, N, A)</math> 2 : <math>(C, T) \leftarrow \text{ENC}_M(K, S, M)</math> 3 : <b>return</b> <math>(C, T)</math> </pre>	<pre> 16 : <math>M_1, \dots, M_\mu \xleftarrow{r} \text{pad}_{10^*}(M, r)</math> 17 : <b>for</b> <math>i = 1, \dots, \mu - 1</math> 18 :   <math>(X, C_i) \leftarrow \xi(\lceil S \rceil_r, M_i)</math> 19 :   <math>Y \leftarrow (X \oplus \omega(\lfloor S \rfloor_c)) \parallel \lfloor S \rfloor_c</math> 20 :   <math>S \leftarrow \rho^b(Y)</math> 21 : <math>(X, C_\mu) \leftarrow \xi(\lceil S \rceil_r, M_\mu)</math> 22 : <math>Y \leftarrow X \parallel (\lfloor S \rfloor_c \oplus \iota_M)</math> 23 : <math>Y \leftarrow (\lceil Y \rceil_r \oplus \omega(\lfloor Y \rfloor_c)) \parallel \lfloor Y \rfloor_c</math> 24 : <math>C \leftarrow \lceil C_1 \rceil \parallel \dots \parallel C_\mu \lfloor_{ M }</math> 25 : <math>S \leftarrow \rho^a(Y)</math> 26 : <math>T \leftarrow \lceil S \rceil_r \oplus K</math> 27 : <b>return</b> <math>(C, T)</math> </pre>
<pre> ENC<math>_e(K, N, A)</math> 4 : <math>A_1, \dots, A_\alpha \xleftarrow{r} \text{pad}_{10^*}(A, r)</math> 5 : <math>Y \leftarrow K \parallel N</math> 6 : <math>S \leftarrow \rho^a(Y)</math> 7 : <b>for</b> <math>i = 1, \dots, \alpha - 1</math> 8 :   <math>(X, \cdot) \leftarrow \xi(\lceil S \rceil_r, A_i)</math> 9 :   <math>Y \leftarrow (X \oplus \omega(\lfloor S \rfloor_c)) \parallel \lfloor S \rfloor_c</math> 10 :   <math>S \leftarrow \rho^b(Y)</math> 11 : <math>(X, \cdot) \leftarrow \xi(\lceil S \rceil_r, A_\alpha)</math> 12 : <math>Y \leftarrow X \parallel (\lfloor S \rfloor_c \oplus \iota_A)</math> 13 : <math>Y \leftarrow (\lceil Y \rceil_r \oplus \omega(\lfloor Y \rfloor_c)) \parallel \lfloor Y \rfloor_c</math> 14 : <math>S \leftarrow \rho^a(Y)</math> 15 : <b>return</b> <math>S</math> </pre>	<pre> ENC<math>_e(K, N, A)</math> 28 : <math>Y \leftarrow \text{FeistelSwap}(S) \oplus I</math> 29 : <math>O \leftarrow S \oplus I</math> 30 : <b>return</b> <math>(Y, O)</math> </pre>

Fig. 35: Pseudocode of SCHWAEMM [7] in terms of ENC $_e$  and ENC $_M$ .