

Committing AE from Sponges

Security Analysis of the NIST LWC Finalists

Juliane Krämer¹, Patrick Struck² and Maximiliane Weishäupl¹

¹ Universität Regensburg, Regensburg, Germany

juliane.kraemer@ur.de maximiliane.weishaeupl@ur.de

² Universität Konstanz, 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 entail 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 are well-suited for building committing AE schemes. Furthermore, we show several negative results when applying the zero-padding method to the NIST finalists.

Keywords: Authenticated Encryption · Committing Security · NIST LWC Finalists

*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 Universität 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).

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	7
2.3	NIST LWC Finalists	8
3	Committing Security Analysis	10
3.1	Romulus	10
3.2	TinyJAMBU	14
3.3	Ascon	17
4	Zero-Padding the NIST Finalists	22
4.1	Motivation and Description of the Zero-Padding Approach	22
4.2	Zero-Padding and Context-Pre-Processing Schemes	23
4.3	Zero-Padding and Elephant	24
4.4	Zero-Padding and ISAP	25
5	Conclusion	25
	References	25
A	Additional Preliminaries	30
A.1	The NIST LWC Finalists	30
A.2	Paddings and Security Notions	31
A.3	(Tweakable) Block-Ciphers	33
A.4	Sponges	33
A.5	Existing Results	34
B	Additional Committing Security Analysis	36
B.1	Elephant	36
B.2	GIFT-COFB	39
B.3	PHOTON-Beetle	43
B.4	Xoodyak	45
B.5	ISAP	47
B.6	Schwaemm	52
C	Deferred Proofs	56
C.1	Proof of Lemma 2	56
C.2	Proof of Lemma 3	56
C.3	Proof of Lemma 4	57
C.4	Proof of Theorem 3	58
C.5	Proof of Theorem 5	59
C.6	Proof of Theorem 7	59
C.7	Proof of Theorem 8	59

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 [Vau02], 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 [NIST15, Ber14]. 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, while 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 [Res18]. The CAESAR competition for authenticated encryption [Ber14] and the recent NIST lightweight cryptography (LWC) standardization process [NIST15], both called specifically for AE schemes which are deemed secure if they provide both confidentiality and authenticity.

However, a series of recent attacks [LGR21, DGRW18, ADG⁺22] has shown that there still might be attack scenarios not covered by the standard notions. The *Facebook message franking attack* [DGRW18] 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* [ADG⁺22] and the *partitioning oracle attack* [LGR21]. The latter allows for more efficient key recovery: The adversary crafts a ciphertext that decrypts validly under multiple keys (for instance a list of leaked keys) and sends it to the recipient; if the recipient rejects the ciphertext, 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, bearing the need for an additional security notion. To this end, committing security [BH22] was defined by requiring each ciphertext to be a commitment to the key (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, and message, and wins if their contexts differ, i.e., $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$, and $\text{AE.ENC}(K, N, A, M) = \text{AE.ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$ holds.

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. However, this approach is ill-advised as it requires a separate analysis for each protocol and puts the burden on the protocol designers. A better approach is to prove authenticated encryption schemes to be committing as those 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 [BH22, MLGR23]. A majority of them were 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 [TMC⁺23], there is barely any research with respect to their committing security [NSS23, DGL23].

While not all of the existing committing attacks are immediately critical in the lightweight setting, the partitioning oracle attack is applicable and can facilitate key recovery on embedded devices, as is classically done via side-channel attacks (described in [SW24]). This highlights the importance of committing security for lightweight AE schemes—in particular if one takes into consideration that more committing attacks that are relevant for such schemes might arise in the future.

1.1 Contribution

In this paper we analyze the committing security of the NIST LWC finalists that are based on (tweakable) block-ciphers or permutations.¹ 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 [MLGR23], 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.

We divide the NIST LWC finalists into two groups: Firstly, ELEPHANT [BCDM21] and ISAP [DEM⁺21] follow the *Encrypt-then-MAC* (EtM) paradigm, as they start by encrypting the message and then authenticate the resulting ciphertext alongside the context. Secondly, ROMULUS [IKM⁺21], GIFT-COFB [BCI⁺21], PHOTON-BEETLE [BCD⁺21], XOODYAK [DHP⁺21], TINYJAMBU [WH21], ASCON [DEMS21], and SCHWAEMM [BBdS⁺21] share a common structure, in the sense that they first process the context and then the message. We refer to schemes of this 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 queries² to attacks that are costlier—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 schemes³ for which we show committing security.

Several of our attacks share the same idea. This is the case for ROMULUS and GIFT-COFB, which are both block-cipher-based AE schemes and 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, as we need to invert the state-update-function for each ciphertext block and the attack only works if *all* are invertible. 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 permutation-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

¹This covers all finalists except GRAIN-128AEAD [HJM⁺21] which comes with a dedicated design.

²More precisely, these attacks need only the minimal cost of computing the respective encryption algorithm twice (once for each of the output tuples).

³Note that we consider a slightly modified version of SCHWAEMM.

the ciphertext correctly—in this case the decryption of ELEPHANT will never return \perp .⁴ 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, as it is 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 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 considered schemes⁵), hence colliding tags can be found 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 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. 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.

We can make the following observations for the committing security of a scheme. A too short tag enables efficient birthday attacks that might extend to committing attacks, as—depending on the scheme—colliding ciphertexts might be easy to find. Further, our analysis reveals that the size of the state that cannot directly be influenced by the inputs (for sponge-based schemes, this is the capacity) plays a crucial role regarding the committing security of a scheme. If the size of this “input-independent” state is too small (even if that happens only once in the scheme), a CMT adversary can produce a full-state collision with only a few tries, which in almost all cases we considered, enables an efficient attack. In fact, a variation of this concept is used in all of our attacks except for the one on TINYJAMBU, which exploits its short tag length. This weakness can also be present in the initial state, i.e., if it contains only a small fixed part, while the rest is filled with inputs provided by the CMT adversary. Note that for the NIST finalists the aforementioned weakness is only present in some of the sponge-based schemes. From this, we can derive some explicit criteria. The schemes we prove committing secure fulfill all of these criteria, while the ones proven insecure lack at least one of them. This suggests that these properties might be sufficient to achieve CMT security—however, we want to emphasize that the criteria should only be seen as a heuristic, as they are observations from the analysis of the schemes, not proven statements. Taking into account our convention of 64-bit committing security, the three criteria are: (I) the tag length is at least 128 bits, (II) at least 128 bits of the state are unaffected by the inputs, and—solely for sponges—(III) the initial state contains at least 64 bits that are independent of the inputs. While we observe that the absence of these properties often enables committing attacks, this is not necessarily the case as the example of SCHWAEMM shows: The unmodified version does not fulfill (III), as the initial

⁴Menda et al. [MLGR23] coin this property as *NoFailDecrypt*.

⁵ELEPHANT uses 64-bit tags as well, but also gives a parameter set with 128-bit tags.

Table 1: Overview of results: a **✗** indicates a committing (CMT) attack with essentially no queries; a **◆** indicates a CMT attack with significantly less than 2^{64} queries; and a **✓** indicates about 64-bit CMT security. Further, the table depicts three properties (I), (II), and (III), that are relevant for the committing analysis: (I) the tag length is at least 128 bits, (II) at least 128 bits of the state are unaffected by the inputs, and (III) at least 64 bits of the initial state are independent of the inputs. Note that we consider property (III) only for sponge-based schemes. The letters **y** and **n** indicate that the property is and is not, respectively, present in the scheme. If the letter is red, we exploit the respective property in our committing attack.

Scheme	CMT	Properties			Sponge	Section
		(I)	(II)	(III)		
ROMULUS	✗	y	n	-	n	Section 3.1
ELEPHANT	✗	n ^a	n	-	n	Appendix B.1
GIFT-COFB	✗	y	n	-	n	Appendix B.2
PHOTON-BEETLE	✗	y	y	n	y	Appendix B.3
TINYJAMBU	◆	n	n	y	y ^b	Section 3.2
XOODYAK	◆	y	n	y	y	Appendix B.4
ASCON	✓	y	y	y	y	Section 3.3
ISAP	✓	y	y	y	y	Appendix B.5
SCHWAEMM	✓	y	y	y ^c	y	Appendix B.6

^a This is only the case for two out of three parameter sets, including the main one.

^b TINYJAMBU can be viewed as a sponge construction based on a block-cipher.

^c This holds for our slightly modified version SCHWAEMM_{IV}, but not the original one.

state is completely determined by key and nonce—however, the output-blinding deployed in SCHWAEMM prevents a CMT attack.

Given the many negative results, we continue with the question whether the NIST finalists can be patched easily to achieve committing security. A very simple method—especially to preserve the lightweight aspect—is the zero-padding approach [ADG⁺22]. However, we show several scheme-specific negative results regarding the NIST finalists when applying zero-padding. This implies that some of the finalists require the application of costlier transforms to achieve committing security.

In conclusion, our analysis shows that most NIST LWC finalists do not achieve committing security (some even after being zero-padded) and those that do, are all sponge-based.

1.2 Related Work

Committing security can be traced back to [ABN10, FLPQ13] where the focus was on public-key encryption. In [FOR17]—using the name *key-robustness*—Farshim et al. gave first definitions of committing security for symmetric encryption. Recently, Bellare and Hoang [BH22] 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. [MLGR23] developed a framework for fine-grained committing security notions. Instead of just having a ciphertext being a commitment 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 [MLGR23], Chan and Rogaway [CR22] also developed

a 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.

In concurrent and independent work, Naito et al. [NSS23] study the committing security of ASCON. While we aim at analyzing all NIST finalists, the focal point of their work is on giving an exhaustive analysis of ASCON. They study both the mode and the underlying permutation of ASCON, whereas we focus our analysis on the mode. Further, they analyze how committing security can be increased by zero-padding the message. The results on the committing security of the unmodified ASCON mode, treated in both works, agree.

2 Authenticated Encryption and the NIST LWC Finalists

In this section, we introduce the notation and recall important definitions. 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 written 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 resp. bit shift of x by b bits to the left is written as $x \lll b$ resp. $x \ll b$ (\ggg resp. \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 ι 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 definitions of authenticated encryption 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 \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* [NRS14], i.e., $M = \text{DEC}(K, N, A, C)$ implies that $C = \text{ENC}(K, N, A, M)$. Following [MLGR23], 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 in [Appendix A](#).

Game CMT (CMT-3 in [BH22])

```

1:  $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 
2: if  $(K, N, A) = (\bar{K}, \bar{N}, \bar{A})$ 
3:   return 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: return  $((C, T) = (\bar{C}, \bar{T}))$ 

```

Figure 1: Security game CMT.

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 [NIST15] 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 [DGRW18] that exploited the absence 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.⁶

The AE schemes that we study in this work are the NIST LWC finalists that are based on (tweakable) block-ciphers or permutations.⁷ Details on the classification and parameters of the schemes can be found in Appendix A.1—note that for each candidate we focus on the main parameter set.

2.3.1 Classes of AE Schemes.

The considered schemes can be divided into two classes. The first class encompasses AE schemes that follow the *Encrypt-then-MAC* (EtM) paradigm [BN00]. 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 $\text{ENC}_{\mathcal{C}}$. The result is then processed together with M , and optionally K and N , yielding the ciphertext (C, T) via a function $\text{ENC}_{\mathcal{M}}$. Both classes are illustrated in Fig. 2. 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 Schemes. For the EtM schemes, we can focus on the underlying MAC. Once we have two contexts $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$ that verify the same ciphertext (C, T) , we can immediately derive a committing attack. This is the case because

⁶A recent NIST workshop [NIST23] mentions committing security as a desirable property.

⁷This covers all finalists except GRAIN-128AEAD [HJM+21].




Figure 2: Illustration of Encrypt-then-MAC schemes (left) and Context-pre-Processing schemes (right). Both ELEPHANT and ISAP are Encrypt-then-MAC schemes. The other schemes are Context-pre-Processing 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 arrow; and ROMULUS has both arrows.

for the described contexts, the decryption algorithm will return some messages $M, \bar{M} \neq \perp$.⁸ Using the tidyness property, we get $\text{ENC}(K, N, A, M) = (C, T) = \text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M})$, hence winning the game CMT.

Attacking Context-pre-Processing Schemes. For the CpP schemes, we focus on the state S that is outputted by ENC_e and then fed into ENC_M . The general idea is to generate the first context (K, N, A) and a message M at random, and compute the corresponding ciphertext (C, T) . Then, we invert ENC_M for the same ciphertext (C, T) and—depending on the scheme—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 ENC_e with input $(\bar{K}, \bar{N}, \bar{A})$ results in \bar{S} , which ultimately yields a committing attack as $\text{ENC}(\bar{K}, \bar{N}, \bar{A}, \bar{M}) = \text{ENC}_M(\bar{K}, \bar{N}, \text{ENC}_e(\bar{K}, \bar{N}, \bar{A}), \bar{M}) = \text{ENC}_M(\bar{K}, \bar{N}, S, \bar{M}) = (C, T)$. The step of finding \bar{A} is essentially what was recently coined a *context discovery attack* [MLGR23]. This is a stronger attack that easily translates to a committing attack as shown in [MLGR23]. Indeed, our 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.

2.3.2 State-Update-Function.

Out of the nine finalists, ROMULUS, GIFT-COFB, PHOTON-BEETLE, and SCHWAEMM deploy a so-called *state-update-function* (this name is adopted from [IKM⁺21]). This function—which we will generally denote by ξ —takes as input a state S and some additional input data I , and outputs a new state Y and additional output data O . The state-update-functions works 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. Details are given in Appendix A.1.

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 ξ . 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.

⁸For both ELEPHANT and ISAP, the underlying decryption algorithm never returns \perp , thus the AE scheme returns \perp iff the verification of the tag fails.

Our attack against PHOTON-BEETLE is independent of the used state-update-function and for SCHWAEMM the state-update-function is incorporated into our security proof.

2.3.3 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 notions [MLGR23]. 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.⁹

3 Committing Security Analysis

Here, we analyze the CMT security of the NIST LWC finalists. For ROMULUS (cf. Section 3.1) we give an attack that breaks committing security with essentially no cost—requiring the bare minimum of two encryptions. For TINYJAMBU (cf. Section 3.2), we provide a committing attack requiring about 2^{33} queries. For ASCON (cf. Section 3.3) we give a formal proof showing that the scheme achieves committing security of about 64-bit. These three schemes represent the different kind of results we have for the NIST LWC finalists and illustrate the core ideas. At the end of each section, we briefly discuss the remaining schemes with similar results and a formal analysis is given in Appendix B.

Our analysis covers all of the NIST LWC finalists except GRAIN-128AEAD, which we exclude due to our focus on schemes based on (tweakable) block-ciphers and permutations. While it is easy to see that the committing security of GRAIN-128AEAD is upper-bound at 32-bit (as there is only one parameter set featuring a 64-bit tag), developing more efficient attacks requires a dedicated analysis. In particular, the so-called accumulator used in GRAIN-128AEAD for tag generation has to be analyzed in detail: Since GRAIN-128AEAD is a stream cipher, achieving the same ciphertexts can be easily done, while the more challenging part seems to be forcing same tags. The accumulator gathers message, associated data, and a pseudorandom bit string into the final tag; it is however no ideal primitive, i.e., a concrete analysis is necessary to find better committing attacks (or exclude the existence of such).

3.1 ROMULUS

ROMULUS [IKM⁺21, IKMP20] is an authenticated encryption scheme based on tweakable block-ciphers. For the concrete instantiation of the TBC, they use SKINNY [BJK⁺16] and the authenticated encryption mode bears similarities with COFB [CIMN17]. 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.

⁹Very recently, Struck and Weishäupl [SW24] developed a generic transformation that turns an AE scheme into one that is both committing and leakage-resilient.

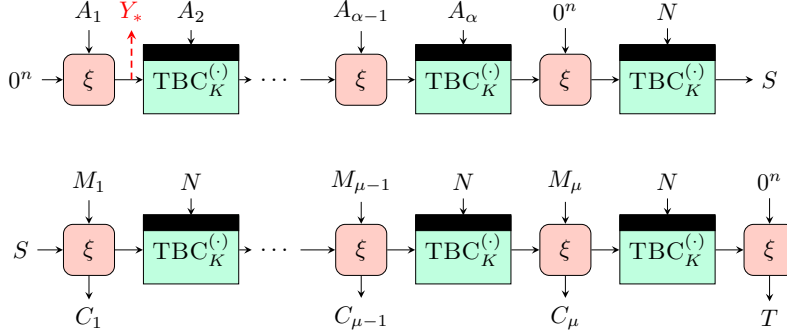


Figure 3: Illustration of ROMULUS (for α an even number) in terms of ENC_e (top) and ENC_M (bottom). The values that are input from the top into $\text{TBC}_K^{(\cdot)}$ are used as tweaks (for simplicity, we drop the counters making the tweaks unique). The state Y_* , marked in red, is used in our CMT attack.

3.1.1 Description of ROMULUS

The pseudocode of ROMULUS is given in Fig. 4. The scheme is further illustrated in Fig. 3. 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 ENC_e and ENC_M apply the tweakable block-cipher and the state-update-function ξ in an alternating manner. The latter, i.e.,

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

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

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

3.1.2 Committing Attack Against ROMULUS

We show that ROMULUS does not achieve CMT security. The attack is stated in the following theorem.

Theorem 1. *Consider ROMULUS which is illustrated and described in Fig. 3 and Fig. 4, respectively. 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, μ is the number of message blocks, α is the number of associated data blocks of the first tuple, and $\bar{\alpha}$ is the number of blocks for the second tuple that \mathcal{A} outputs.

For the proof of Theorem 1, we formulate and prove three lemmas. Firstly, we show that the state-update-function ξ is invertible (Lemma 1). Secondly, we prove that we can invert both ENC_e and ENC_M (Lemma 2 and Lemma 3), where, for the latter, we make use of the invertibility of ξ .

ROMULUS.ENC(K, N, A, M)	ENC \mathcal{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: $(C, T) \leftarrow \text{ENC}_{\mathcal{M}}(K, N, S, M)$	16: for $i = 1, \dots, \mu - 1$
3: return (C, T)	17: $(Y, C_i) \leftarrow \xi(S, M_i)$
	18: $S \leftarrow \text{TBC}^N(K, Y)$
	19: $(Y, C_\mu) \leftarrow \xi(S, M_\mu)$
	20: $S \leftarrow \text{TBC}^N(K, Y)$
	21: $(\cdot, O) \leftarrow \xi(S, 0^n)$
	22: $T \leftarrow [O]_\tau$
	23: $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$
	24: return (C, T)
	$\xi(S, I)$
	<hr style="width: 50%; margin: 0 auto;"/>
	25: $Y \leftarrow S \oplus I$
	26: $O \leftarrow G(S) \oplus I$
	27: return (Y, O)
ENC $_e$ (K, N, A)	
4: $A_1, \dots, A_\alpha \xleftarrow{n} \text{pad}_L(A, n)$	
5: $S \leftarrow 0^n$	
6: for $i = 1, \dots, \lfloor \frac{\alpha}{2} \rfloor$	
7: $(Y, \cdot) \leftarrow \xi(S, A_{2i-1})$	
8: $S \leftarrow \text{TBC}^{A_{2i}}(K, Y)$	
9: $V \leftarrow 0^n$	
10: if $\alpha \bmod 2 \neq 0$	
11: $V \leftarrow A_\alpha$	
12: $(Y, \cdot) \leftarrow \xi(S, V)$	
13: $S \leftarrow \text{TBC}^N(K, Y)$	
14: return S	

Figure 4: Pseudocode of ROMULUS [IKM⁺21] in terms of ENC $_e$ and ENC \mathcal{M} . For sake of simplicity, we drop the counter that is part of the tweak.

Recall that the state-update-function ξ 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 ξ 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 ξ is discarded in ENC $_e$ —a fact that will be exploited later. For our attack against ENC \mathcal{M} , this no longer works, as we have to invert the output of ξ while maintaining equal ciphertexts. The following lemma shows that we can invert ξ with respect to *both* its output and state. We write M for the input and C for the output of ξ (instead of I and O), which is the case for our scenario.

Lemma 1. *The state-update-function of ROMULUS is invertible.*

Proof. Note that ξ can be expressed as the block-matrix $\begin{pmatrix} \text{id} & | & \text{id} \\ G & | & \text{id} \end{pmatrix}$ and hence in order to show the claim, we need to find a right-sided inverse to this matrix. For $+$ denoting component-wise addition mod 2, consider $\begin{pmatrix} F & | & F \\ F + \text{id} & | & F \end{pmatrix}$,

$$\text{with } F = \begin{pmatrix} F_s & 0 & 0 & \dots & 0 \\ 0 & F_s & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \dots & 0 & F_s & 0 \\ 0 & \dots & 0 & 0 & F_s \end{pmatrix} \text{ and } F_s = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix}.$$

Observe that

$$\begin{pmatrix} \text{id} & | & \text{id} \\ G & | & \text{id} \end{pmatrix} \cdot \begin{pmatrix} F & | & F \\ F + \text{id} & | & F \end{pmatrix} = \begin{pmatrix} F + F + \text{id} & | & F + F \\ GF + F + \text{id} & | & GF + F \end{pmatrix} = \begin{pmatrix} \text{id} & | & 0 \\ 0 & | & \text{id} \end{pmatrix},$$

where the last equality follows from a direct computation which shows that $G_s F_s + F_s = F_s(G_s + \text{id}) = \text{id}$ and thus $GF + F = F_s(G + \text{id}) = \text{id}$. \square

Next, we give an adversary that inverts ENC_e , i.e., for a given output S of ENC_e and a partial context (K, N) , it finds matching associated data. We exploit the fact that the associated data blocks are XORed to the full state in ENC_e . The detailed proof can be found in [Appendix C.1](#).

Lemma 2. *Consider ROMULUS which is illustrated and described in [Fig. 3](#) and [Fig. 4](#), respectively. 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 α being the number of associated data blocks that \mathcal{A}_e outputs.

We now give an adversary that inverts 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 ξ as shown in [Lemma 1](#) and is proven in [Appendix C.2](#).

Lemma 3. *Consider ROMULUS which is illustrated and described in [Fig. 3](#) and [Fig. 4](#), respectively. There exists an adversary \mathcal{A}_M , making q queries to $\tilde{\text{E}}$ such that for any $(K, N, (C, T)) \in \mathcal{K} \times \mathcal{N} \times \mathcal{C}$, it holds that*

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

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

Proof (of [Theorem 1](#)). We construct the following adversary \mathcal{A} against ROMULUS as shown in [Fig. 5](#). 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 $(\bar{S}, \bar{M}) \leftarrow \mathcal{A}_M(\bar{K}, \bar{N}, (C, T))$, 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 2](#) and [Lemma 3](#), we obtain

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

As for the number of queries to the ideal tweakable cipher $\tilde{\text{E}}$, \mathcal{A} makes $\mu + \lfloor \frac{\alpha}{2} \rfloor + 1$ while computing the ciphertext (C, T) for the first tuple and additionally μ and $\lfloor \frac{\alpha}{2} \rfloor + 1$ queries while running \mathcal{A}_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 gist of the attack is finding a different \bar{A} which yields the target ciphertext. The attack easily extends to a context discovery attack (CDY_A^*) [[MLGR23](#)]. Hence, we can conclude that ROMULUS is also vulnerable with respect to the weaker security notions CMT_K and CMT_N by using [[MLGR23](#), [Corollary 3](#)]. Furthermore, the attack can be translated to one against CMT_A 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_A^* by choosing the second key-nonce pair (\bar{K}, \bar{N}) not at random but equal to the first pair (K, N) .

Similar Results. Just as ROMULUS, we can attack ELEPHANT (cf. [Appendix B.1](#)), GIFT-COFB (cf. [Appendix B.2](#)), and PHOTON-BEETLE (cf. [Appendix B.3](#)) with a minimum number of queries. The attack against GIFT-COFB is very similar to the one given here, the core difference is the state-update-function. The attack against ELEPHANT is even simpler as the scheme does not use a state-update-function. For PHOTON-BEETLE, our attack exploits the choice of its initial state.

ROMULUS adversary $\mathcal{A}()$	ENC_e adversary $\mathcal{A}_e(K, N, S)$
1 : $(K, N, A, M) \leftarrow_{\$} \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$	14 : $\alpha \leftarrow_{\$} 2\mathbb{N}$
2 : $(C, T) \leftarrow \text{ENC}(K, N, A, M)$	15 : $A_2, \dots, A_\alpha \leftarrow_{\$} \{0, 1\}^n$
3 : $(\overline{K}, \overline{N}) \leftarrow_{\$} \mathcal{K} \times \mathcal{N}$	16 : $A_{\alpha+1} \leftarrow 0^n$
4 : $(\overline{S}, \overline{M}) \leftarrow \mathcal{A}_{\mathcal{M}}(\overline{K}, \overline{N}, (C, T))$	17 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, N, S)$
5 : $\overline{A} \leftarrow \mathcal{A}_e(\overline{K}, \overline{N}, \overline{S})$	18 : for $i = \frac{\alpha}{2}, \dots, 1$ do
6 : return $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, \overline{M})$	19 : $S \leftarrow Y \oplus A_{2i+1}$
	20 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, A_{2i}, S)$
$\text{ENC}_{\mathcal{M}}$ adversary $\mathcal{A}_{\mathcal{M}}(K, N, (C, T))$	21 : $A_1 \leftarrow Y$
7 : $C_1, \dots, C_\gamma \xleftarrow{n} C$	22 : $A \leftarrow A_1 \parallel \dots \parallel A_\alpha$
8 : $S \leftarrow G^{-1}(T)$	23 : return A
9 : for $i = \mu, \dots, 1$ do	
10 : $Y \leftarrow \tilde{\text{E}}^{-1}(K, N, S)$	
11 : $(S, M_i) \leftarrow \xi^{-1}(Y, C_i)$	
12 : $M \leftarrow M_1 \parallel \dots \parallel M_\mu$	
13 : return (S, M)	

Figure 5: ROMULUS adversary \mathcal{A} from Theorem 1 which uses the inverse state-update-function ξ^{-1} from Lemma 1.

3.2 TINYJAMBU

TINYJAMBU [WH21] is a block-cipher-based authenticated encryption scheme. The specification introduces the TINYJAMBU mode, which is a lightweight variant of the JAMBU mode [WH16]. The latter was part of the CAESAR competition [Ber14]. For the permutation underlying TINYJAMBU, a keyed permutation based on non-linear feedback shift registers is defined.

3.2.1 Description of TINYJAMBU

The pseudocode of TINYJAMBU is given in Fig. 7 and an illustration of the scheme can be found in Fig. 6. TINYJAMBU follows the CpP-approach, i.e., it first processes the context (K, N, A) via the function ENC_e and then passes the output on to $\text{ENC}_{\mathcal{M}}$, where it is processed together with the message. TINYJAMBU uses two keyed permutations BC_1 and BC_2 , both based on the same keyed permutation that is applied 640 and 1024 times for BC_1 and BC_2 , respectively.

3.2.2 Committing Attack Against TINYJAMBU

In this section, we show that TINYJAMBU does not achieve CMT security.¹⁰ 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 [WH21, Section 6]), this setting is modeled with only one permutation ρ . We adopt the same for the TINYJAMBU attack given in the following.

Theorem 2. *Consider TINYJAMBU which is illustrated and described in Fig. 6 and Fig. 7, respectively. Let BC_1 and BC_2 be modeled as an ideal cipher E . Then there exists an*

¹⁰Dunkelmann et al. [DGL23] show that TINYJAMBU is not $\text{CMT}_{\mathcal{K}}$ -secure. Their committing attack is based on a related-key forgery attack on the underlying block-cipher and only applies to the larger parameter sets but not the main one. In contrast, our attack is based on the mode and works for all parameter sets.

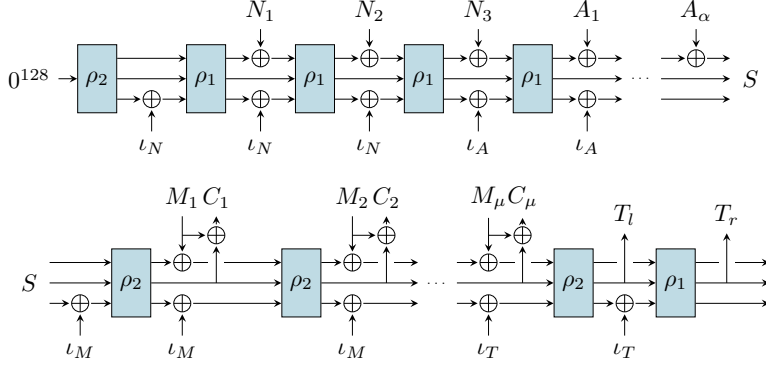


Figure 6: Illustration of TINYJAMBU in terms of ENC_E (top) and ENC_M (bottom), where $\rho_1 = \text{BC}_1(K, \cdot)$ and $\rho_2 = \text{BC}_2(K, \cdot)$.

adversary \mathcal{A} that makes q queries to \mathbf{E} such that

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

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

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 distinct (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 distinct (\bar{N}_i, \bar{A}_i) ¹¹ 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 ρ .

For the context (K, N_i, A_i) , denote the states before the second to last permutation application (see Fig. 6) 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 $r = 32$ and capacity $c = 96$. Therefore, the event $S_i = S_j$ for $i \neq j$ (and analogously $\bar{S}_i = \bar{S}_j$ for $i \neq j$), constitutes an inner collision, which is—for a sponge with capacity 96—highly unlikely¹². As we model both BC_1 and BC_2 by an ideal cipher \mathbf{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 5 (for $l_1 = l_2 = 2^{32} + 1$ and $\tau = 64$), hence we obtain the following lower bound for finding a collision $T_i = \bar{T}_j$:

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

¹¹For sake of simplicity, we assume that \mathcal{A} chooses the \bar{A}_i with same block length.

¹²More precisely, the probability is $\frac{q(q+1)}{2^{97}} - \frac{q(q-1)}{2^{129}}$ [BDPVA07].

TINYJAMBU.ENC(K, N, A, M)	ENC $_M(K, S, M)$
1 : $S \leftarrow \text{ENC}_E(K, N, A)$	16 : $M_1, \dots, M_\mu \xleftarrow{32} \text{pad}_{0^*}(M, 32)$
2 : $(C, T) \leftarrow \text{ENC}_M(K, S, M)$	17 : for $i = 1, \dots, \mu$
3 : return (C, T)	18 : $S \leftarrow S \oplus (0^{64} \parallel \iota_M)$
ENC $_E(K, N, A)$	19 : $S \leftarrow \text{BC}_2(K, S)$
4 : $N_1, N_2, N_3 \xleftarrow{32} N$	20 : $S \leftarrow ([S]_{32} \oplus M_i) \parallel [S]_{96}$
5 : $A_1, \dots, A_\alpha \xleftarrow{32} \text{pad}_{0^*}(A, 32)$	21 : $C_i \leftarrow [S]_{33}^{64}$
6 : $S \leftarrow \text{BC}_2(K, 0^{128})$	22 : $S \leftarrow S \oplus (0^{64} \parallel \iota_T)$
7 : for $i = 1, \dots, 3$	23 : $S \leftarrow \text{BC}_2(K, S)$
8 : $S \leftarrow S \oplus (0^{64} \parallel \iota_N)$	24 : $T_l \leftarrow [S]_{33}^{64}$
9 : $S \leftarrow \text{BC}_1(K, S)$	25 : $S \leftarrow S \oplus (0^{64} \parallel \iota_T)$
10 : $S \leftarrow ([S]_{32} \oplus N_i) \parallel [S]_{96}$	26 : $S \leftarrow \text{BC}_1(K, S)$
11 : for $i = 1, \dots, \alpha$	27 : $T_r \leftarrow [S]_{33}^{64}$
12 : $S \leftarrow S \oplus (0^{64} \parallel \iota_A)$	28 : $C \leftarrow [C_1 \parallel \dots \parallel C_\mu]_{ M }$
13 : $S \leftarrow \text{BC}_1(K, S)$	29 : $T \leftarrow T_l \parallel T_r$
14 : $S \leftarrow ([S]_{32} \oplus A_i) \parallel [S]_{96}$	30 : return (C, T)
15 : return S	

Figure 7: Pseudocode of TINYJAMBU [WH21] in terms of ENC $_E$ and 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 ι_A and ι_M).

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

The attack exploits the fact that TINYJAMBU uses a very short tag (64 bits) compared to the other schemes—the only other scheme considered in this work 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 be absorbed bit-wise.

By construction, the above attack is a 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 CMT $_N$ and a CMT $_A$ attack.

Similar Results. The sponge-based AE scheme XOODYAK (cf. Appendix B.4) can also be efficiently attacked using a birthday attack. We target the full-state sponge part of XOODYAK when the associated data is processed. Due to the padding of XOODYAK, the adversary cannot control the entire state—32 bits are reserved for the padding. Hence, a birthday attack on these 32 bits is required which then extends to a CMT attack.

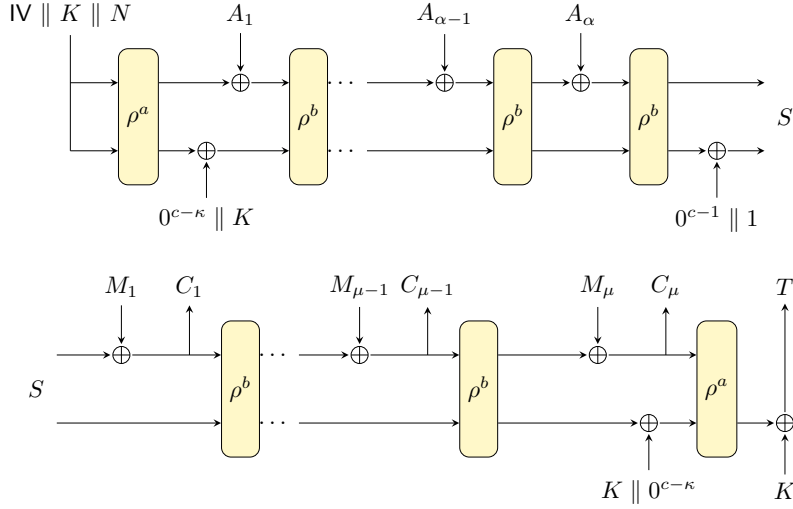


Figure 8: Illustration of ASCON in terms of ENC_e (top) and ENC_M (bottom).

3.3 ASCON

ASCON [DEMS21] is a sponge-based AE 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 [LM24] gave the first security proof for ASCON.¹³

3.3.1 Description of ASCON

The pseudocode of ASCON is given in Fig. 9 and further illustration is provided in Fig. 8. Similar to the other schemes, ASCON can be viewed as a CpP scheme which first processes the context using ENC_e before the message is processed using ENC_M . A core feature of ASCON is that at the very start (first permutation of ENC_e) and the very end (last permutation of ENC_M), it uses more rounds of the underlying permutation for security (ρ^a and ρ^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.

3.3.2 Committing Security of ASCON

We show that ASCON achieves CMT security. We model the two permutations ρ^a and ρ^b by one ideal permutation ρ , essentially adopting the approach of ISAP [DEM⁺20], where the different permutations for the re-keying function are also modeled by one random permutation. 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.¹⁴ Note that these changes are purely cosmetic. They do not influence the overall security of ASCON but simplify our proof, as we can capture the state-blinding by considering a larger rate.

¹³An earlier work [JLM14] showed security for a simplified version of ASCON.

¹⁴This only affects the first state-blinding; the second one is already of that form.

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

Figure 9: Pseudocode of ASCON [DEMS21] in terms of ENC_e and $\text{ENC}_{\mathcal{M}}$.

Lastly, for sake of simplicity, we drop the domain separation of ASCON in the proof. It can, however, be easily incorporated by XORing $1 \parallel 0^{c-1}$ instead of $0^{c-1} \parallel 1$, i.e., moving the domain separation to the first bit of the capacity as opposed to the last bit. This neither influences the security of ASCON nor interferes with the purpose of the domain separation but it allows us to incorporate the domain separation into the rate.¹⁵

We show committing security of ASCON by arguing about the collision resistance of plain sponges—an alternative approach is to make use of the indistinguishability, as was done in concurrent and independent work [NSS23]. Theorem 9 provides a bound for the collision resistance of plain sponges, however, applying it to (a plain sponge version of) ASCON does not yield a good bound. The reason for this is that the capacity of ASCON’s initial state (64 bits) is small compared to the one of the remaining states (≥ 128 bits). As part of our argument, we prove that Theorem 9 also holds if the initial state has a smaller capacity than the rest of the sponge construction.¹⁶

This generalized version of [BS23, Theorem 8.6] uses the same graph modelling as the original proof. We consider a sponge-based hash function \mathcal{H} , obtained from a permutation $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$, with capacity c , rate r , initial capacity c^* , and output length $w \leq r$. We model a CR adversary against \mathcal{H} as follows: We build a directed graph G from the ideal permutation queries the adversary makes. The nodes in G are the 2^n bit strings of length n 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 ρ queries are called *forward edges* and the ones resulting from ρ^{-1} queries are referred to as *backward edges*. 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. $\lfloor Y_0 \rfloor_{c^*} = \text{IV}$,
2. $\lfloor Y_i \rfloor_c = \lfloor S_i \rfloor_c$ for all $i \in \{1, \dots, s-1\}$, and

¹⁵The same argument appears in [DJS19] (AE scheme SLAE) and [DHP⁺21] (XOODYAK).

¹⁶Naito and Ohta [NO14] showed that a smaller capacity in the initial state can also be tolerated when considering indistinguishability, which the initial result [BDPV08] did not.

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 $I := (Z_0, \dots, Z_{s-1}) \in \{0, 1\}^{n-c^*} \times \{0, 1\}^r \times \dots \times \{0, 1\}^r$ for $Z_0 = [Y_0]_{n-c^*}$ and $Z_i = [S_i]_r \oplus [Y_i]_r$ for all $i \in \{1, \dots, s-1\}$. Further, we define the *result* of a PS-path as $[S_s]_w$ which corresponds to the output of the hash function. As a notation for PS-paths that incorporates the input, we write

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

Next, we define two properties a pair (P, \bar{P}) of PS-paths can have. For this, denote the nodes in \bar{P} by $\bar{Y}_0, \bar{S}_1, \bar{Y}_1, \dots, \bar{Y}_{l-1}, \bar{S}_l$. Firstly, the paths P and \bar{P} are *colliding* if their inputs differ but their results agree. Secondly, the paths P and \bar{P} 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 can show the following lemma regarding the probability of finding a pair of problematic paths. The proof can be found in Appendix C.3.

Lemma 4. *Consider the graph modelling as described above. Let PP be the event that \mathcal{A} finds a pair of problematic paths and q the number of queries that \mathcal{A} makes to ρ , then*

$$\Pr[\text{PP}] \leq \frac{q}{2^{c^*-1}} + \frac{q(q-1)}{2^c}.$$

Using this result yields a generalization of [BS23, Theorem 8.6] to the case that the initial state has a smaller capacity. This is proven in Appendix C.4.

Theorem 3 (Generalized version from [BS23, Theorem 8.6]). *For \mathcal{H} a hash function obtained from a permutation $\rho: \{0, 1\}^n \rightarrow \{0, 1\}^n$, with capacity c , rate r (so $n = r + c$), initial capacity c^* , and output length $w \leq r$, it holds that*

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

The theorem below establishes the committing security of ASCON.

Theorem 4. *Consider ASCON which is illustrated and described in Fig. 8 and Fig. 9, respectively. Let ρ^a and ρ^b be modeled as a random permutation ρ . Then for any adversary \mathcal{A} making $q \leq 2^{127}$ queries to ρ , it holds that*

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

Proof. 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 TagColl (see Fig. 15). Hence we obtain

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

We consider the directed graph described above for $n = 320$. We assume \mathcal{A} to make queries to ρ that correspond to its output, i.e., querying all states that occur during the evaluation

of ASCON for the output tuples of \mathcal{A} . This assumption is without loss of generality, as we can easily transform any adversary \mathcal{A} 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 ρ corresponding to (K, N, A, M) and $(\overline{K}, \overline{N}, \overline{A}, \overline{M})$. 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 ρ queries are made on Y and no further ρ^{-1} queries are made on S .

We define a special kind of path, the *A-path*, which models ASCON’s tag generation—recall that we passed over to the game **TagColl** at the beginning of the proof. An A-path P_A of length l^{17} is a sequence of $2l$ nodes

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

with the following four properties:

1. $Y_0 = K \parallel N \parallel IV$,
2. $[Y_1]_{256} = [S_1]_{256} \oplus (K \parallel 0^{128})$ and $[Y_{l-1}]_{256} = [S_{l-1}]_{256} \oplus (K \parallel 0^{128})$,
3. $[Y_i]_{256} = [S_i]_{256}$ 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 $K \parallel N := [Y_0]_{256}$ 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: The sponge processes the input consisting of key K , nonce N , and the tuple of associated data and message $(A, M) = (X_1, \dots, X_{l-1})$ and the result is the tag. As a notation for A-paths that incorporates the input, we write

$$(K \parallel N) | Y_0 \rightarrow \dots \rightarrow S_{l-2} | X_{l-2} | Y_{l-2} \rightarrow S_{l-1} | X_{l-1} | Y_{l-1} \rightarrow S_l.^{18}$$

Next, we define two properties a pair (P_A, \overline{P}_A) of A-paths can have. For this, denote the nodes in \overline{P}_A by $\overline{Y}_0, \overline{S}_1, \overline{Y}_1, \dots, \overline{Y}_{l-1}, \overline{S}_l$. Firstly, the paths P_A and \overline{P}_A are *colliding* if their inputs differ but their results agree. Secondly, the paths P_A and \overline{P}_A are *problematic* if their inputs differ and (1) $Y_{l-1} = \overline{Y}_{l-1}$ or (2) at least one of the edges in P or \overline{P} is a backward edge. We are interested in the event CP_A that \mathcal{A} finds a pair of colliding A-paths. Note that finding such paths means that \mathcal{A} wins the game **TagColl**. In order to compute the probability of CP_A , we define the auxiliary event PP_A that \mathcal{A} finds a pair of problematic A-paths. Using this, we obtain

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

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 $\text{CP}_A \wedge \neg \text{PP}_A$ 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 $\overline{I} = (\overline{K}, \overline{N}, \overline{X}_1, \dots, \overline{X}_{l-1})$ such that the corresponding A-paths

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

¹⁷Note that $l \geq 3$, as ASCON involves at least three applications of ρ (for $\mu = \alpha = 1$).

¹⁸While 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.

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

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

Since ρ 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 $(S_{l-1}, X_{l-1}, K), (\bar{S}_{l-1}, \bar{X}_{l-1}, \bar{K})$ ¹⁹ such that $Y_{l-1} \neq \bar{Y}_{l-1}$ and Eq. (2) holds with q queries, equals the probability of finding a collision in a list of q uniformly distributed elements. Using Theorem 11 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 $\Pr[\text{PP}_A]$. Lemma 4 shows 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 ρ with rate 256 for the first round of absorption and rate 196 for all remaining ones. Further, its initial state is given by $0^{256} \| \text{IV}$ and the output produced by \mathcal{H} has length 128. 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 PP_A implies the event PP. Let (P_A, \bar{P}_A) be a pair of problematic A-paths, i.e.,

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

By defining

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

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

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

Visualization for this is provided in Fig. 10. As we neither change (Y_{l-1}, \bar{Y}_{l-1}) nor any of the edges, the paths (P, \bar{P}) form a pair of problematic PS-paths. Thus, we have shown that PP_A implies PP, hence $\Pr[\text{PP}_A] \leq \Pr[\text{PP}]$. This allows us to focus on the plain sponge setting for the rest of the proof. Applying Lemma 4 for $n = 320$, $c^* = 64$, and $c = 128$, then yields the following bound for finding a pair of problematic plain sponge paths

$$\Pr[\text{PP}] \leq \frac{q}{2^{63}} + \frac{q(q-1)}{2^{128}}.$$

In total, we have shown

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

¹⁹Note that \mathcal{A} must choose $(S_{l-1}, X_{l-1}, K) \neq (\bar{S}_{l-1}, \bar{X}_{l-1}, \bar{K})$ to ensure $Y_{l-1} \neq \bar{Y}_{l-1}$.




Figure 10: Illustration of a proof step for ASCON (Theorem 4). ASCON is represented as a plain sponge with a larger rate.

which finishes the proof. \square

Similar Results. Similar to ASCON, we can show committing security for ISAP (cf. Appendix B.5) and SCHWAEMM (cf. Appendix B.6). For ISAP, the overall idea is similar to the one used in the ASCON proof. While ISAP’s large IV allows a direct application of Theorem 9, extra care is necessary to handle its re-keying mechanism. For SCHWAEMM, we consider a slightly modified version $\text{SCHWAEMM}_{\text{IV}}$, where we introduce a fixed IV in the initial state as ASCON and ISAP have. The proof then makes use of the indistinguishability of (the plain sponge version) of $\text{SCHWAEMM}_{\text{IV}}$ from a random function.

4 Zero-Padding the NIST Finalists

In this section we consider the zero-padding approach for the NIST finalists, which is described in Section 4.1. In Section 4.2, we refine the class of CpP schemes and show for a subclass—containing PHOTON-BEETLE and XOODYAK—that CMT_{K} attacks reduce to finding key collisions in the underlying function ENC_{e} . In Section 4.3, we show that ZP-ELEPHANT does not achieve CMT_{K} if the number of padded zeros is less or equal to the block-size of ELEPHANT. In Section 4.4, we show that ZP-ISAP can be attacked with about $2^{\frac{z}{2}}$ queries regardless of the number of padded zeros.

4.1 Motivation and Description of the Zero-Padding Approach

As discussed in Section 2, there are several transformations that turn arbitrary AE schemes into committing AE schemes but these might not work in the context of the NIST LWC finalists. A much simpler and less invasive method is the so-called zero-padding suggested by Albertini et al. [ADG⁺22]. Here, a number of zeros is prepended to the message before encrypting it, i.e.,

$$\text{ZP-AE.ENC}(K, N, A, M) := \text{AE.ENC}(K, N, A, 0^z \parallel M).$$

Decryption uses the decryption algorithm of the AE scheme and additionally checks whether the resulting message has z leading zeros. If it does, the zeros are discarded and

the message is returned, otherwise, the ciphertext is rejected and \perp is outputted. For any scheme AE , we write its corresponding zero-padding scheme with a prefix ZP- , i.e., ZP-AE ; the number of zeros that are padded is denoted by z . Zero-padding is particularly interesting as it does not require any changes to the scheme: prepending zeros to the message and verifying that a decrypted message has leading zeros can be done outside of the AE scheme.

Albertini et al. [ADG⁺22] suggested zero-padding as a method to obtain CMT_K security and Menda et al. [MLGR23] showed that in general it does not achieve CMT security as CMT_A^* attacks are still possible. Since all our attacks are also CMT_A^* attacks—or can be easily modified to be—we cannot achieve CMT security via the zero-padding. Nonetheless, zero-padding as a method to obtain CMT_K security is still interesting, as the latter seems to be the more relevant notion for practical attacks [ADG⁺22, DGRW18, LGR21].

This raises the question whether the broken NIST LWC finalists can be patched by zero-padding to achieve CMT_K security. We provide several negative results for this: The schemes ZP-PHOTON-BEETLE and ZP-XOODYAK remain insecure, irrespectively of the value of z , while ZP-ELEPHANT is vulnerable if $z \leq n$.

Besides achieving CMT_K security, Naito et al. [NSS23] showed that zero-padding can increase CMT security (up to certain level that is) for ASCON . More precisely, they show that padding zeros can compensate for smaller tags. This raises the question whether it also improves the security of other schemes. We provide yet another negative result, showing that ISAP does not benefit from zero-padding regarding its CMT security.

4.2 Zero-Padding and Context-Pre-Processing Schemes

In Section 2, we introduced CpP schemes and argued that most of the NIST finalists fall into this category. Recall that CpP schemes first process the context (K, N, A) via ENC_e and afterwards process the result together with the message (and optionally the key K and nonce N) via ENC_M yielding the ciphertext. The classification of CpP schemes can be refined further, depending on whether ENC_M takes the key K and nonce N as input. We call schemes where ENC_M solely takes the output of ENC_e and the message M as input *full-CpP schemes*. Out of the NIST LWC finalists, PHOTON-BEETLE and XOODYAK are full-CpP schemes, whereas for the others the key K is also input to ENC_M (and for ROMULUS additionally also the nonce N).²⁰ For CpP schemes, we define a game KeyColl , which asks the adversary to find a key collision for ENC_e , i.e., two contexts that differ (at least) in the keys but result in the same output of ENC_e .

Definition 3. Let AE be a CpP scheme, with encryption composed of ENC_e and ENC_M , and let the game KeyColl be defined as in Fig. 11. For any adversary \mathcal{A} , its KeyColl advantage is defined as

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

For full-CpP schemes, key committing security can be lower bounded by finding a key collision for the pre-processing component ENC_e . Any key collision attack against ENC_e translates to a key committing attack against the AE scheme—even when applying the zero-padding, regardless of how many zeroes are prepended to the message. This is formally stated in the theorem below, which is proven in Appendix C.5.

Theorem 5. *Let AE be a full-CpP scheme, where encryption is decomposed into ENC_e and ENC_M . Then, for any adversary \mathcal{A} against ENC_e , there exists an adversary \mathcal{B} against ZP-AE such that*

$$\text{Adv}_{\text{ZP-AE}}^{\text{CMT}_K}(\mathcal{B}) \geq \text{Adv}_{\text{ENC}_e}^{\text{KeyColl}}(\mathcal{A}).$$

²⁰There is no scheme for which the nonce is input to ENC_M while the key is not.

```

Game KeyColl
-----
1 : (K, N, A), (K̄, N̄, Ā) ← A()
2 : if K = K̄
3 :   return 0
4 : S ← ENCe(K, N, A)
5 : S̄ ← ENCe(K̄, N̄, Ā)
6 : return (S = S̄)

```

Figure 11: Security game KeyColl for CpP schemes.

Having established [Theorem 5](#) we show that PHOTON-BEETLE and XOODYAK both allow for key collision attacks against the function ENC_e as stated in the following theorem.

Theorem 6. *Let $\text{AE} \in \{\text{PHOTON-BEETLE}, \text{XOODYAK}\}$ with underlying components ENC_e and ENC_M . Let further ρ be modeled as an ideal permutation. Then there exists an adversary \mathcal{A} , making q queries to ρ , such that*

$$\text{Adv}_{\text{ENC}_e}^{\text{KeyColl}}(\mathcal{A}) \geq \frac{1}{2}.$$

For PHOTON-BEETLE, we have $q = \alpha + \bar{\alpha}$. For XOODYAK, we have $q = 2^{17} + 1$.

Proof. The proof follows from the proofs of [Theorem 14](#) (PHOTON-BEETLE) and [Theorem 15](#) (XOODYAK). \square

In combination, the above two theorems show that neither PHOTON-BEETLE nor XOODYAK are CMT_K secure even when using the zero-padding approach.

4.3 Zero-Padding and ELEPHANT

The following theorem shows that ELEPHANT with zero-padding does not achieve CMT_K security provided that the number of padded zeros is smaller than the block size of ELEPHANT. The proof can be found in [Appendix C.6](#).

Theorem 7. *Consider ELEPHANT which is illustrated and described in [Fig. 18](#) and [Fig. 19](#), respectively. Let the tweakable block-cipher TEM be modeled as an ideal tweakable cipher $\tilde{\text{E}}$ and let $z \leq n$. There exists an adversary \mathcal{A} , making q queries to $\tilde{\text{E}}$, such that*

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

where $q = 2\mu + 2\gamma + \alpha + \bar{\alpha}$. Here, μ is the number of message blocks while computing ENC_M (note that the block length is based on $0^z \parallel M$ and not just M) and γ is the number of ciphertext blocks while computing ENC_T . Furthermore, α and $\bar{\alpha}$ are the number of associated data blocks for the two tuples that \mathcal{A} outputs.

[Theorem 7](#) shows that zero-padding is pointless for ELEPHANT if the number of prepended zeros is less than or equal to the block size (160/176/200 for the three parameter sets JUMBO/DUMBO/DELIRIUM). To ensure a valid zero-padding for both tuples, the adversary needs to find (K, N) and (\bar{K}, \bar{N}) such that $\tilde{\text{E}}(K, (0, 1), N) = \tilde{\text{E}}(\bar{K}, (0, 1), \bar{N})$ holds, which is trivial as $\tilde{\text{E}}$ is an ideal block-cipher for a fixed tweak—in this case $(0, 1)$. The attack, however, does not apply anymore if the zero-padding affects more than one block. Assume that $z = 2n$, then \mathcal{A} needs to find (K, N) and (\bar{K}, \bar{N}) such that both $\tilde{\text{E}}(K, (0, 1), N) = \tilde{\text{E}}(\bar{K}, (0, 1), \bar{N})$ and $\tilde{\text{E}}(K, (0, 2), N) = \tilde{\text{E}}(\bar{K}, (0, 2), \bar{N})$. By the same means as above, one can find (K, N) and (\bar{K}, \bar{N}) such that any one of the two equations holds

but not necessarily both of them simultaneously. More generally, for $z = n + x$, the probability is $\frac{1}{2^z}$, i.e., the security increases with the number of padded zeros but only if the zero-padding exceeds the first block.

4.4 Zero-Padding and ISAP

The following theorem (proven in detail in [Appendix C.7](#)) shows that ISAP does not benefit from the zero-padding approach to increase its committing security. This is in sharp contrast to ASCON, for which Naito et al. [[NSS23](#)] showed that zero-padding indeed affects the committing security positively.

Theorem 8. *Consider ISAP which is illustrated and described in [Fig. 30](#) and [Fig. 31](#), respectively. Let ρ_B and ρ_K be modeled by an ideal permutation ρ_1 . Let further ρ_H and ρ_E be modeled by ideal permutations ρ_2 and ρ_3 , respectively. Then there exists an adversary \mathcal{A} , making q_1 , q_2 , and q_3 queries to ρ_1 , ρ_2 , and ρ_3 , respectively, such that*

$$\mathbf{Adv}_{\text{ZP-ISAP}}^{\text{CMT}}(\mathcal{A}) \geq \frac{1}{2},$$

where $q_1 = (2^{\frac{\kappa}{2}+1} + 1)(\kappa + 1) + \nu$, $q_2 = (2^{\frac{\alpha}{2}+1} + 1)(\alpha + \gamma + 2)$, and $q_3 = \mu$.

5 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. From the analysis, we identified three criteria that are related to committing security: (I) requires a tag length of at least 128 bits, (II) requires 128 bits of the state to be unaffected by the inputs, and (III) requires the initial state (for sponge-based AE schemes) to have 64 bits not affected by the inputs. We showed that the absence of these criteria often yields CMT attacks.

Overall, our analysis reveals that sponge constructions built on top of large permutations are favorable for committing security, as having 128 bits unaffected by the inputs (II) is then more easily achievable. Note that for sponges this should also hold for the initial state (III). Constructions built from block-ciphers often fail to achieve (II), as the entire state typically consists only of 128 bits—though the NIST workshop on accordion cipher modes [[NIST24](#)] also considers block-ciphers with 256-bit states which could solve the issue. On the other hand, a tag of at least 128 bits (I) is typically not a problem and can also be achieved by block-ciphers and permutations with smaller state sizes.

Finally, we show that for several of the NIST finalists, the problems that lead to committing attacks persist even when deploying the zero-padding method. Furthermore, we show that the interesting observation from [[NSS23](#)], i.e., zero-padding improves the committing security of ASCON, does not apply to ISAP. Overall, this shows that the benefit of the zero-padding is rather restricted.

References

- [ABN10] 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.
- [ADG⁺22] Ange Albertini, Thai Duong, Shay Gueron, Stefan Kölbl, Atul Luykx, and Sophie Schmieg. How to abuse and fix authenticated encryption without key commitment. In Kevin R. B. Butler and Kurt Thomas, editors, *USENIX Security 2022*, pages 3291–3308. USENIX Association, August 2022.

- [BBC⁺20] 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.
- [BBdS⁺21] 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>.
- [BCD⁺21] 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>.
- [BCDM20] 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.
- [BCDM21] Tim Beyne, Yu Long Chen, Christoph Dobraunig, and Bart Mennink. Elephant. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
- [BCI⁺21] 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>.
- [BDP⁺16] 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>.
- [BDPV08] 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.
- [BDPVA07] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sponge functions. In *ECRYPT Hash Workshop, 2007*.
- [Ber14] Daniel J. Bernstein. CAESAR: Competition for authenticated encryption: Security, applicability, and robustness. <https://competitions.cr.yp.to/caesar.html>, 2014.
- [BH22] Mihir Bellare and Viet Tung Hoang. Efficient schemes for committing authenticated encryption. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 845–875. Springer, Heidelberg, May / June 2022.
- [BJK⁺16] 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.

- [BN00] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, Heidelberg, December 2000.
- [BPP⁺17] 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.
- [BS23] Dan Boneh and Victor Shoup. *A Graduate Course in Applied Cryptography*. 2023. Draft 0.6, <http://toc.cryptobook.us/>.
- [CDNY18] 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>.
- [CIMN17] 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.
- [CR22] John Chan and Phillip Rogaway. On committing authenticated-encryption. In Vijayalakshmi Atluri, Roberto Di Pietro, Christian Damsgaard Jensen, and Weizhi Meng, editors, *ESORICS 2022, Part II*, volume 13555 of *LNCS*, pages 275–294. Springer, Heidelberg, September 2022.
- [DEM⁺17] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, and Thomas Unterluggauer. ISAP – towards side-channel secure authenticated encryption. *IACR Trans. Symm. Cryptol.*, 2017(1):80–105, 2017.
- [DEM⁺20] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, Bart Mennink, Robert Primas, and Thomas Unterluggauer. ISAP v2.0. *IACR Trans. Symm. Cryptol.*, 2020(S1):390–416, 2020.
- [DEM⁺21] 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>.
- [DEMS21] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer. Ascon. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
- [DGL23] Orr Dunkelman, Shibam Ghosh, and Eran Lambooj. Practical related-key forgery attacks on full-round TinyJAMBU-192/256. In *ToSC 2023 (2)*, 2023.
- [DGRW18] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast message franking: From invisible salamanders to encryptment. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 155–186. Springer, Heidelberg, August 2018.

- [DHAK18] 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.
- [DHP⁺21] 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>.
- [DJS19] Jean Paul Degabriele, Christian Janson, and Patrick Struck. Sponges resist leakage: The case of authenticated encryption. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 209–240. Springer, Heidelberg, December 2019.
- [DM19] Christoph Dobraunig and Bart Mennink. Leakage resilience of the duplex construction. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 225–255. Springer, Heidelberg, December 2019.
- [DPU⁺16] 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.
- [FLPQ13] 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.
- [FOR17] 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.
- [GJMN16] 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.
- [GPP11] 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.
- [HJM⁺21] 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>.
- [IKM⁺21] Tetsu Iwata, Mustafa Khairallah, Kazuhiko Minematsu, Thomas Peyrin, and Chun Guo. Romulus. Technical report, National Institute of Standards and Technology, 2021. Available at <https://csrc.nist.gov/projects/lightweight-cryptography/finalists>.
- [IKMP20] 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.

- [JLM14] 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.
- [LGR21] Julia Len, Paul Grubbs, and Thomas Ristenpart. Partitioning oracle attacks. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021*, pages 195–212. USENIX Association, August 2021.
- [LM24] Charlotte Lefevre and Bart Mennink. Generic security of the Ascon mode: On the power of key blinding. In *SAC 2024*, 2024.
- [LRW02] 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.
- [MLGR23] Sanketh Menda, Julia Len, Paul Grubbs, and Thomas Ristenpart. Context discovery and commitment attacks - how to break CCM, EAX, SIV, and more. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part IV*, volume 14007 of *LNCS*, pages 379–407. Springer, Heidelberg, April 2023.
- [NIST15] National Institute of Standards and Technology. Lightweight cryptography standardization process. <https://csrc.nist.gov/projects/lightweight-cryptography>, 2015.
- [NIST23] National Institute of Standards and Technology. The third NIST workshop on block cipher modes of operation 2023. <https://csrc.nist.gov/projects/lightweight-cryptography>, 2023. Accessed: 2024-02-23.
- [NIST24] National Institute of Standards and Technology. NIST workshop on the requirements for an accordion cipher mode 2024. <https://csrc.nist.gov/Events/2024/accordion-cipher-mode-workshop-2024>, 2024. Accessed: 2024-04-04.
- [NO14] Yusuke Naito and Kazuo Ohta. Improved indifferentiable security analysis of PHOTON. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 340–357. Springer, Heidelberg, September 2014.
- [NRS14] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering generic composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, Heidelberg, May 2014.
- [NSS23] Yusuke Naito, Yu Sasaki, and Takeshi Sugawara. Committing security of Ascon: Cryptanalysis on primitive and proof on mode. In *ToSC 2023 (4)*, 2023.
- [Res18] Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.
- [SW24] Patrick Struck and Maximiliane Weishäupl. Constructing committing and leakage-resilient authenticated encryption. In *ToSC 2024 (1)*, 2024.
- [TMC⁺23] 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.

- [Vau02] 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.
- [Wag02] David Wagner. A generalized birthday problem. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 288–303. Springer, Heidelberg, August 2002.
- [WH16] 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>.
- [WH21] 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 give additional details about the NIST LWC finalists and in [Appendix A.2](#) we describe the used paddings and define more security notions. We give some background on (tweakable) block-ciphers and sponges in [Appendix A.3](#) and [Appendix A.4](#), respectively. Finally, in [Appendix A.5](#), we provide some results that are relevant for our attacks and proofs.

A.1 The NIST LWC Finalists

We provide two tables and a figure with additional information about the NIST LWC finalists regarding their similarities and differences. [Table 2](#) gives an overview and [Table 3](#) describes the parameter sets. [Fig. 12](#) shows the state-update-function for the different schemes.

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

Scheme	Class of scheme	Underlying Primitive	Sponge	State-update-function
ROMULUS	CpP	Block-cipher ^a	No	Yes
ELEPHANT	EtM	Permutation	No	No
GIFT-COFB	CpP	Block-cipher	No	Yes
PHOTON-BEETLE	CpP	Permutation	Yes	Yes
TINYJAMBU	CpP	Block-cipher	Yes	No
XOODYAK	CpP	Permutation	Yes	No
ASCON	CpP	Permutation	Yes	No
ISAP	EtM	Permutation	Yes	No
SCHWAEMM	CpP	Permutation	Yes	Yes

^a Note that ROMULUS actually uses a tweakable block-cipher for its underlying primitive.

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 ASCON-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 ENC_c component in XOODYAK is a full-state sponge and in ISAP’s re-keying mechanism a minimal rate of 1 is used.

Scheme	Key κ	Nonce ν	Tag τ	State n	Rate r	Capacity c
ROMULUS	128	128	128	128	-	-
ELEPHANT	128	96	64	160	-	-
GIFT-COFB	128	128	128	128	-	-
PHOTON-BEETLE	128	128	128	256	128	128
TINYJAMBU	128	96	64	128	32	96
XOODYAK	128	128	128	384	192	192
ASCON	128	128	128	320	64	256
ISAP	128	128	128	400	144	256
SCHWAEMM	128	256	128	384	256	128

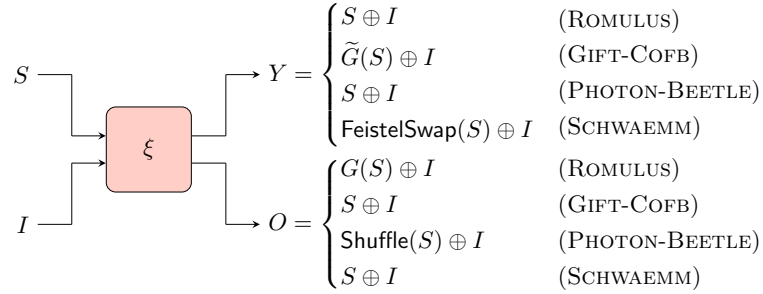


Figure 12: Illustration of the state-update-function ξ for the different schemes. The components G , \tilde{G} , Shuffle , and FeistelSwap —if relevant for our results—are described along with the schemes in the respective sections.

A.2 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 length is a multiple of r . Simply padding with 0s to a length of a multiple of 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 such that the overall length is a multiple of r .

Below we define collision resistance of a hash function.

Definition 4. Let $\mathcal{H}: \{0, 1\}^* \rightarrow \{0, 1\}^w$ be a hash function outputting bit strings of 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. [MLGR23] 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.²¹

Definition 5. Let $\text{AE} = (\text{ENC}, \text{DEC})$ be an authenticated encryption scheme and the games CMT_X and CMT_X^* for $X \in \{K, N, A\}$ be defined as in Fig. 1. For any adversary \mathcal{A} ,

²¹Note that CMT_K originates from [BH22].

<p>Game CMT_K</p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: if $K = \bar{K}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$	<p>Game CMT_K^*</p> <hr/> 1: $((K, \bar{K}), N, A, (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: if $K = \bar{K}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$
<p>Game CMT_N</p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: if $N = \bar{N}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$	<p>Game CMT_N^*</p> <hr/> 1: $(K, (N, \bar{N}), A, (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: if $N = \bar{N}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$
<p>Game CMT_A</p> <hr/> 1: $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M}) \leftarrow \mathcal{A}()$ 2: if $A = \bar{A}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$	<p>Game CMT_A^*</p> <hr/> 1: $(K, N, (A, \bar{A}), (M, \bar{M})) \leftarrow \mathcal{A}()$ 2: if $A = \bar{A}$ 3: return 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: return $((C, T) = (\bar{C}, \bar{T}))$

Figure 13: Security games CMT_K , CMT_N , CMT_A , CMT_K^* , CMT_N^* , and 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.

its CMT_X and CMT_X^* advantages are defined as

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

respectively.

Furthermore, Menda et al. [MLGR23] introduce so-called *context-discovery* notions. In this work, we are only interested in the notion CDY_A^* . Roughly speaking, this notion provides an adversary with a ciphertext C and partial context (consisting of the key K and nonce N used to generate the ciphertext), and the adversary is challenged to find associated data A which validly decrypts the ciphertext. Below we recall this security notion.

Definition 6. Let $\text{AE} = (\text{ENC}, \text{DEC})$ be an authenticated encryption scheme and the game CDY_A^* be defined as in Fig. 14. For any adversary \mathcal{A} , its CDY_A^* advantage is defined as

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

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.


```

Game CDYA*
-----
1 : (K, N, (C, T)) ←s K × N × C
2 : A ← A(K, N, (C, T))
3 : M ← DEC(K, N, A, (C, T))
4 : if M = ⊥
5 :   return 0
6 : return 1

```

Figure 14: Security game CDY_A^{*} for authenticated encryption schemes.

```

Game TagColl
-----
1 : (K, N, A, M), (K̄, N̄, Ā, M̄) ← A()
2 : if (K, N, A) = (K̄, N̄, Ā)
3 :   return 0
4 : (C, T) ← ENC(K, N, A, M)
5 : (C̄, T̄) ← ENC(K̄, N̄, Ā, M̄)
6 : return (T = T̄)

```

Figure 15: Security game TagColl for authenticated encryption schemes used in the proof of [Theorem 4](#).

Definition 7. Let $\text{AE} = (\text{ENC}, \text{DEC})$ be an authenticated encryption scheme and the game TagColl be defined as in Fig. 15. For any adversary \mathcal{A} , its TagColl advantage is defined as

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

A.3 (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 κ 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 [LRW02] $\text{TBC}: \{0, 1\}^\kappa \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ takes as input a key K of length κ , a tweak W (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 $(K, W) \in \{0, 1\}^\kappa \times \mathcal{T}$, $\text{TBC}(K, W, \cdot)$ is a permutation over $\{0, 1\}^n$. We also use $\text{TBC}_K^W(\cdot)$ as an alternative notation for $\text{TBC}(K, W, \cdot)$.

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

A.4 Sponges

Sponges [BDPVA07] are a versatile tool for cryptographic primitives. Rather than just being relevant for cryptographic hash functions—as was their initial 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 ρ . 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



Figure 16: Block-cipher (left) and tweakable block-cipher (right). For tweakable block-ciphers, the black bar indicates that the incoming arrow (W) is the tweak.

state. The higher the rate the faster the sponge as fewer rounds, hence fewer invocations of ρ , 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. For several of the NIST LWC finalists, the capacity of the initial state differs from the one in the rest of the sponge, which is why we introduce c^* as a notation for this initial capacity.

We refer to sponges of the form described above by *plain sponges* and provide an illustration in Fig. 17. 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 9 ([BS23, 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*

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

Theorem 10 ([BDPV08, 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 ρ , \mathcal{H} is indistinguishable from a random oracle F , except with probability at most $\frac{(1-2^{-256})q^2 + (1+2^{-256})q}{2^{129}}$.*

An improved indistinguishability bound is given by Naito et al. [NO14]. This results allow for a smaller capacity in both the initial state of the sponge ($c^* = \frac{c}{2}$ instead of $c^* = c$) and the squeezing phase ($c^* = \frac{c}{2} + \log_2 c$ instead of $c^* = c$) while maintaining the security bound. Using this improved indistinguishability one can also obtain committing security of ASCON (in contrast to our analysis which uses collision resistance of plain sponges). For ISAP and (our modified version of) SCHWAEMM, the older indistinguishability bound is sufficient as they both feature a larger IV.

A.5 Existing Results

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

Theorem 11 ([BS23, 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*




Figure 17: Illustration of a plain sponge construction with four rounds of absorbing and two rounds of squeezing.

$i, j \in \{1, \dots, k\}$ we have that $X_i = X_j$. Then

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

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

We use the formulation of the birthday problem presented in [Wag02], but provide a more formal description using [Theorem 11](#).

Lemma 5. 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 \mathbf{C} the event that $x_k = x_j$ holds for some $k \neq j$. Since the x_i are drawn uniformly and independently, [Theorem 11](#) yields the bound $\Pr[\mathbf{C}] \geq 1 - \exp\left(\frac{-l(l-1)}{2^{\tau+1}}\right)$. 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 4](#)). While the computations are not hard, we give them here to provide a complete presentation.

Lemma 6. Let $n, c, l \in \mathbb{N}$ such that $c, l \leq n$ and $\mathbf{IV} \in \{0, 1\}^l$. Let further ρ be a random permutation over $\{0, 1\}^n$ and \mathcal{A} be an adversary making queries to ρ . Consider the following events:

1. Event \mathbf{E}_t (target hitting query):

\mathcal{A} makes a query Y to ρ such that $\lfloor \rho(Y) \rfloor_l = \mathbf{IV}$ or \mathcal{A} makes a query S to ρ^{-1} such that $\lfloor \rho^{-1}(S) \rfloor_l = \mathbf{IV}$.

2. Event \mathbf{E}_c (colliding queries):

\mathcal{A} makes queries $Y \neq \bar{Y}$ to ρ such that $\lfloor \rho(Y) \rfloor_c = \lfloor \rho(\bar{Y}) \rfloor_c$ or \mathcal{A} makes queries Y to ρ and \bar{S} to ρ^{-1} such that $\lfloor \rho(Y) \rfloor_c = \lfloor \rho^{-1}(\bar{S}) \rfloor_c$.

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

$$\Pr[E_t] \leq \frac{q}{2^{l-1}} \quad \text{and} \quad \Pr[E_c] \leq \frac{q(q-1)}{2^c}.$$

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

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

where $q \leq 2^{n-1}$ is used for the last inequality. Next, we bound event E_c . Let 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

$$\Pr[E_c] \leq \sum_{j=1}^q \sum_{i=1}^{j-1} \Pr[X_{ij}] \leq \sum_{j=1}^q \frac{(j-1)2^{n-c}}{2^n - j + 1} \leq 2^{n-c} \left(\frac{q(q-1)}{2(2^n - q)} \right),$$

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

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

which finishes the proof. \square

B Additional Committing Security Analysis

Here, we analyze the CMT security of the six remaining NIST LWC finalists. For ELEPHANT (cf. Appendix B.1), GIFT-COFB (cf. Appendix B.2), and PHOTON-BEETLE (cf. Appendix B.3), we give—similar to ROMULUS (cf. Section 3.1)—attacks that break committing security with essentially no cost. For XOODYAK (cf. Appendix B.4), we give a committing attack requiring about 2^{17} queries using a birthday attack as done for TINYJAMBU (cf. Section 3.2). For ISAP (cf. Appendix B.5) and SCHWAEMM (cf. Appendix B.6), we give formal proofs showing that the schemes achieve committing security of about 64-bit, similar to ASCON (cf. Section 3.3).

B.1 ELEPHANT

The AE scheme ELEPHANT [BCDM21,BCDM20] relies on a cryptographic permutation which gets masked using linear feedback shift registers similar to the masked Even-Mansour construction [GJMN16]. The security of the ELEPHANT-mode follows from the security of the masked Even-Mansour construction, which is shown to be indistinguishable from an ideal tweakable block-cipher [BCDM21]. As our results rely on the mode, we choose to present the scheme in terms of such an tweakable block-cipher, which we denote by TEM for tweakable Even-Mansour.

B.1.1 Description of ELEPHANT

The pseudocode of ELEPHANT is given in Fig. 19 and further illustration is provided in Fig. 18. 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 $\text{ENC}_{\mathcal{M}}(K, N, C)$.




Figure 18: 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.

B.1.2 Committing Attack against ELEPHANT

Since ELEPHANT follows the EtM-paradigm, we only need to focus on the underlying function $\text{ENC}_{\mathcal{T}}$. 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 [BCDM21] does not achieve committing security.

Theorem 12. *Consider ELEPHANT which is illustrated and described in Fig. 18 and Fig. 19, respectively. Let TEM 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, μ is the number of message blocks while computing $\text{ENC}_{\mathcal{M}}$ and γ is the number of ciphertext blocks while computing $\text{ENC}_{\mathcal{T}}$.²² Furthermore, α and $\bar{\alpha}$ are the number of associated data blocks for the two tuples that \mathcal{A} outputs.

Proof. We construct a CMT adversary \mathcal{A} against ELEPHANT as shown in Fig. 20. 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 $\bar{K} \leftarrow_{\$} \mathcal{K} \setminus \{K\}$ and associated data blocks $\bar{A}_2, \dots, \bar{A}_{\bar{\alpha}} \leftarrow_{\$} \{0, 1\}^n$.²³ The adversary then computes the state

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

²²Note that μ and γ might not be the same.

²³We assume that \mathcal{A} chooses the last block to exhibit a valid padding.

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

Figure 19: Pseudocode of ELEPHANT [BCDM21] in terms of ENC \mathcal{M} and ENC \mathcal{T} .

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 : for $i = 1, \dots, \gamma$
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 : for $i = 2, \dots, \bar{\alpha}$
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 : $\bar{A}_2, \dots, \bar{A}_{\bar{\alpha}} \leftarrow_{\$} \{0, 1\}^n$	18 : return \bar{S}
7 : $\bar{S} \leftarrow \mathcal{B}(C_1, \dots, C_\gamma, \bar{A}_2, \dots, \bar{A}_{\bar{\alpha}})$	
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 : $(\bar{N}, \bar{A}) \leftarrow \text{pad}_{10^*}^{-1}(\bar{A}_1, \dots, \bar{A}_{\bar{\alpha}})$	
11 : $\bar{M} \leftarrow \text{ENC}_{\mathcal{M}}(\bar{K}, \bar{N}, C)$	
12 : return (K, N, A, M), ($\bar{K}, \bar{N}, \bar{A}, \bar{M}$)	

Figure 20: ELEPHANT adversary \mathcal{A} from Theorem 12.

shown in Fig. 18. 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 $(\bar{N}, \bar{A}) \leftarrow \text{pad}_{10^*}^{-1}(\bar{A}_1 \parallel \dots \parallel \bar{A}_{\bar{\alpha}})$, 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 μ queries to compute C and $\alpha + \gamma$ to compute T . Additionally, \mathcal{A} makes μ 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 attack easily extends to CMT \mathcal{K} , CMT \mathcal{N} , CMT \mathcal{A} , and CMT \mathcal{A}^* attacks. The reasoning follows the one given for ROMULUS. For a CMT \mathcal{A}^* , the attack needs to target a different associated data block than the first as this one contains the nonce.




Figure 21: Illustration of GIFT-COFB in terms of ENC_E (top) and ENC_M (bottom). The different indices for Δ only indicate that the values are different—in the pseudocode the value Δ is constantly updated. The states S_* and Y_* , marked in red, are used in our CMT attack.

B.2 GIFT-COFB

The AE scheme GIFT-COFB uses the COFB mode [CIMN17] for authenticated encryption and instantiates the block-cipher using GIFT [BPP⁺17].

B.2.1 Description of GIFT-COFB

The pseudocode of GIFT-COFB is given in Fig. 22 while Fig. 21 provides an illustration of it. GIFT-COFB follows the CpP-approach, i.e., it first computes $S \leftarrow \text{ENC}_E(K, N, A)$ followed by $(C, T) \leftarrow \text{ENC}_M(K, S, M)$. 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 Δ to the state.

B.2.2 Committing Attack Against GIFT-COFB

The AE scheme GIFT-COFB does not achieve CMT 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 ξ . It turns out that we cannot always invert ξ —inversion only works with probability $\frac{1}{2}$. This is the reason, why the advantage depends on the number of ciphertext blocks.

GIFT-COFB.ENC(K, N, A, M)	ENC _e (K, N, A)
1: $(S, \Delta) \leftarrow \text{ENC}_e(K, N, A)$	20: $Y \leftarrow N$
2: $(C, T) \leftarrow \text{ENC}_M(K, (S, \Delta), M)$	21: $S \leftarrow \text{BC}(K, Y)$
3: return (C, T)	22: $\Delta \leftarrow \lceil S \rceil_{n/2}$
	23: $A_1, \dots, A_\alpha \xleftarrow{n} \text{pad}_{10^*}(A, n)$
ENC _M ($K, (S, \Delta), M$)	24: for $i = 1, \dots, \alpha - 1$
4: $M_1, \dots, M_\mu \xleftarrow{n} \text{pad}_{10^*}(M, n)$	25: $\Delta \leftarrow 2\Delta$
5: for $i = 1, \dots, \mu - 1$	26: $(Y, \cdot) \leftarrow \xi(S, A)$
6: $\Delta \leftarrow 2\Delta$	27: $Y \leftarrow Y \oplus (\Delta \parallel 0^{n/2})$
7: $(Y, C_i) \leftarrow \xi(S, M_i)$	28: $S \leftarrow \text{BC}(K, Y)$
8: $Y \leftarrow Y \oplus (\Delta \parallel 0^{n/2})$	29: if $ A \bmod n = 0$
9: $S \leftarrow \text{BC}(K, Y)$	30: $\Delta \leftarrow 3\Delta$
10: if $ M \bmod n = 0$	31: $j \leftarrow 1$
11: $\Delta \leftarrow 3\Delta$	32: if $ A \bmod n \neq 0$
12: if $ M \bmod n \neq 0$	33: $\Delta \leftarrow 3^2\Delta$
13: $\Delta \leftarrow 3^2\Delta$	34: $j \leftarrow 2$
14: $(Y, C_\mu) \leftarrow \xi(S, M_\mu)$	35: $(Y, \cdot) \leftarrow \xi(S, A_\alpha)$
15: $Y \leftarrow Y \oplus (\Delta \parallel 0^{n/2})$	36: $Y \leftarrow Y \oplus (\Delta \parallel 0^{n/2})$
16: $S \leftarrow \text{BC}(K, Y)$	37: $S \leftarrow \text{BC}(K, Y)$
17: $T \leftarrow \lceil S \rceil_\tau$	38: return (S, Δ)
18: $C \leftarrow \lceil C_1 \parallel \dots \parallel C_\mu \rceil_{ M }$	
19: return (C, T)	$\xi(S, I)$
	39: $Y \leftarrow \tilde{G}(S) \oplus I$
	40: $O \leftarrow S \oplus I$
	41: return (Y, O)

Figure 22: Pseudocode of GIFT-COFB [BCI⁺21] in term of ENC_e and ENC_M.

Theorem 13. Consider GIFT-COFB which is illustrated and described in Fig. 21 and Fig. 22, respectively. 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, μ is the number of message blocks, α is the number of associated data blocks for the first tuple, and $\bar{\alpha}$ for the second tuple that \mathcal{A} outputs.

While the advantage depends on μ , we stress that this value is controlled by \mathcal{A} . In particular, choosing a short ciphertext of only one block leads to an attack with success probability of $\frac{1}{2}$.

For the proof of Theorem 13 we drop the XOR of the masking values. This avoids some very cumbersome and tedious bookkeeping. Subsequent to the proof, we discuss why the result also holds if the masking values are used.

Before giving the proof of Theorem 13, we give three lemmas that we will use to prove it. The first lemma, Lemma 7, 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 8, shows that ENC_e 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 9, shows that there is an algorithm that inverts ENC_M . The latter relies heavily on Lemma 7 which is why the success probability drops exponentially in the number of ciphertext blocks.

Lemma 7. *Let ξ be the state-update-function of GIFT-COFB. Let further \mathcal{A}_ξ 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}_ξ 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}_ξ . 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}_ξ , 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}_ξ . 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. \square

Lemma 8. Consider GIFT-COFB as described in Fig. 22. 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 α 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_α , i.e., all except the first one.²⁴ Next, \mathcal{A}_E computes both S_* and Y_* (cf. Fig. 21): 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 ξ (by XORing an associated data block and inverting \tilde{G} —note that we merely need to invert \tilde{G} as the output of ξ is discarded). Finally, \mathcal{A}_E computes $A_1 \leftarrow S_* \oplus Y_*$ and outputs $A \leftarrow \text{pad}_{10^*}^{-1}(A_1 \parallel \dots \parallel A_\alpha)$. By construction, it holds that

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

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

Lemma 9. Consider GIFT-COFB described in Fig. 22. There exists an adversary \mathcal{A}_M , making q queries to the ideal cipher \mathbf{E} such that

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

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}_M is $q = \mu$ for μ being the number of ciphertext blocks that \mathcal{A}_M receives as input.

Proof. We construct an adversary \mathcal{A}_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 μ full blocks of length n . First, Adversary \mathcal{A}_M sets $S \leftarrow T$. Next, \mathcal{A}_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}_ξ correctly inverts ξ , we obtain

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

By Lemma 7—using Y is uniformly random due to \mathbf{E} being an ideal cipher—every inversion of ξ succeeds with probability $\frac{1}{2}$. Since \mathcal{A}_ξ is run μ times by \mathcal{A}_M , we get

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

The number of queries to the ideal cipher \mathbf{E} by \mathcal{A}_M is μ . □

We are now ready to give the proof of Theorem 13.

Proof (of Theorem 13). 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 $(\bar{S}, \bar{M}) \leftarrow \mathcal{A}_M(\bar{K}, \bar{N}, (C, T))$, and $\bar{A} \leftarrow \mathcal{A}_E(\bar{K}, \bar{N}, \bar{S})$. Finally, \mathcal{A} outputs $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$.

²⁴We assume that \mathcal{A}_E picks the blocks such that they exhibit a valid padding according to GIFT-COFB.

By Lemma 8 and Lemma 9, we have the following equalities with probability $\frac{1}{2^\mu}$:

$$\begin{aligned}
\text{GIFT-COFB.ENC}(\overline{K}, \overline{N}, \overline{A}, \overline{M}) &= \text{ENC}_{\mathcal{M}}(\overline{K}, \overline{N}, \text{ENC}_{\mathcal{E}}(\overline{K}, \overline{N}, \overline{A}), \overline{M}) \\
&= \text{ENC}_{\mathcal{M}}(\overline{K}, \overline{N}, \overline{S}, \overline{M}) && \text{(Lemma 8)} \\
&= (C, T) && \text{(Lemma 9)} \\
&= \text{GIFT-COFB.ENC}(K, N, A, M).
\end{aligned}$$

This yields

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

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

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 ROMULUS.

B.3 PHOTON-BEETLE

The authenticated encryption scheme PHOTON-BEETLE [BCD⁺21] is a (duplex) sponge-based AE scheme. It uses the PHOTON permutation [GPP11] as the underlying permutation and the BEETLE mode of operation [CDNY18]. In contrast to the plain duplex, the BEETLE mode uses a state-update-function to determine the next input to the underlying permutation of the sponge.

B.3.1 Description of PHOTON-BEETLE

PHOTON-BEETLE is described in Fig. 25 and illustrated in Fig. 24. It is a CpP scheme, i.e., it processes the context (K, N, A) via the function $\text{ENC}_{\mathcal{E}}$ and subsequently processes the message together with the output of $\text{ENC}_{\mathcal{E}}$ —an important property is that no part of the context is input to $\text{ENC}_{\mathcal{M}}$. In $\text{ENC}_{\mathcal{M}}$, 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.

B.3.2 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 14 below.

Theorem 14. *Consider PHOTON-BEETLE which is illustrated and described in Fig. 24 and Fig. 25, respectively. Let ρ be modeled as an ideal permutation. Then there exists an adversary \mathcal{A} , making q queries to ρ , such that*

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

where $q = \alpha + \bar{\alpha}$. Here, α 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.

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

Figure 23: GIFT-COFB adversary \mathcal{A} from Theorem 13 and the state-update-function adversary \mathcal{A}_{ξ} from Lemma 7.

Proof. We construct the following CMT adversary \mathcal{A} against PHOTON-BEETLE as shown in Fig. 26. 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. 24). Adversary \mathcal{A} then chooses different associated data \overline{A} and inverts ENC_e starting from S_* up to the initial state. This initial state is then used as the concatenation of nonce \overline{N} and key \overline{K} .²⁵ As the final step, \mathcal{A} picks a message M at random and outputs $((K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, M))$. By construction, we have $(K, N, A) \neq (\overline{K}, \overline{N}, \overline{A})$ and it holds that

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

thus \mathcal{A} wins the game CMT.

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

The attack is by construction also a valid $\text{CMT}_{\mathcal{A}}$ attack as the associated data are chosen to be different. While the adversary does not choose the key \overline{K} and nonce \overline{N}

²⁵Note that these are likely to be different than K and N but not guaranteed to be.




Figure 24: Illustration of PHOTON-BEETLE in terms of ENC_e (top) and ENC_M (bottom). The state S_* , marked in red, is used in our CMT attack.

for the second tuple, it is easy to see that the attack extends to CMT_K and CMT_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.

B.4 XOODYAK

The authenticated encryption scheme XOODYAK [DHP⁺21] is an AE scheme based on a full-state keyed duplex. XOODYAK uses the XODOO permutation [DHAK18] 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 [BDP⁺16] to the lightweight setting.

B.4.1 Description of XOODYAK

The pseudocode of XOODYAK is given in Fig. 28 and further illustration is provided in Fig. 27. XOODYAK uses a state of size 384 and is a CpP scheme, i.e., first $S \leftarrow ENC_e(K, N, A)$ is computed, followed by the computation of ciphertext and tag as $(C, T) \leftarrow ENC_M(S, M)$. In ENC_e , the inputs are XORed onto the full-state (note that the last 32 bits are reserved for padding). In contrast to this, 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.

B.4.2 Committing Attack Against XOODYAK

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

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

Figure 25: Pseudocode of PHOTON-BEETLE [BCD⁺21] in terms of ENC_e and ENC_M. Here, $\text{Shuffle}(S) = S_2 \parallel (S_1 \ggg 1)$ for $S_1, S_2 \xleftarrow{\frac{r}{2}} S$.

Theorem 15. *Consider XOODYAK which is illustrated and described in Fig. 27 and Fig. 28, respectively. Let ρ be modeled as an ideal permutation. Then there exists an adversary \mathcal{A} that makes $q = 2^{17} + 1$ queries to ρ and fulfills*

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

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 ρ . For this, $q = 2^{17} + 1$ different key-nonce pairs 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 [BS23, Section 8.3]²⁶, 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 $\bar{A} \leftarrow \lfloor \rho(\bar{\mathbf{K}} \parallel \bar{\mathbf{N}}) \rfloor_{352} \oplus \lfloor \rho(\mathbf{K} \parallel \mathbf{N}) \rfloor_{352} \oplus A$. 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, from which point on 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

²⁶Note that the prerequisite, regarding the size of domain and codomain of f , is fulfilled as $2^{384} \geq 100 \cdot 2^{32}$.

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

Figure 26: PHOTON-BEETLE adversary \mathcal{A} from Theorem 14.

also a valid CMT_K and CMT_N attack as keys and nonces, respectively, are chosen to be different. Moreover, it can be shown to be a CMT_A attack: The same associated data blocks are only chosen if the states after the first application of ρ already coincide in their rate part. Since they agree in the last 32 bits by construction, this would constitute a full-state collision of ρ , which is impossible for a permutation.

B.5 ISAP

The authenticated encryption scheme ISAP [DEM⁺21, DEM⁺17, DEM⁺20] 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.

B.5.1 Description of ISAP

The pseudocode of ISAP is given in Fig. 31 and further illustration is provided in Fig. 30. ISAP follows the EtM-approach, i.e., first the message is encrypted via ENC_M resulting in a ciphertext C and afterwards, the tag T is computed using ENC_T , which processes the context and the ciphertext. Both ENC_M and ENC_T internally uses the re-keying function ISAP.RK to derive the session key. For the security analysis [DEM⁺21], ENC_M is viewed as a keyed duplex construction while ENC_T is viewed as a suffix keyed sponge (SuKS) [DM19].

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 ρ_K , ρ_H , ρ_B , and ρ_E , each based on the same permutation but applied a different number of times.

B.5.2 Committing Security of ISAP

We show that ISAP achieves CMT security. In the security proof for ISAP [DEM⁺20], the two permutations used in ISAP.RK (namely ρ_K and ρ_B) are modeled as one permutation. We adopt this for our proof and denote the permutation used in ISAP.RK by ρ_1 and the




Figure 27: Illustration of XOODYAK in terms of ENC_e (top) and ENC_M (bottom). Here, $\mathbf{K} \parallel \mathbf{N} = \text{pad}_e((\mathbf{K} \parallel \mathbf{N} \parallel \text{enc}_8(|\mathbf{N}|)), 00000010)$, $\mathbf{M}_i = \text{pad}_M(M_i)$, $\mathbf{A}_1 = \text{pad}_e(A_1, 00000011)$, and for $i = 2, \dots, \alpha$, $\mathbf{A}_i = \text{pad}_e(A_i, 0^8)$. In ENC_M , the increased rate ($r + 8$) is required for the padding pad_M which appends 8 bits to the message blocks.

one used in ENC_T by ρ_2 .²⁷ In conformity with this, the rate in ISAP.RK is denoted by r_1 and the rate in ENC_T by r_2 . Furthermore, we consider a slightly different domain separation in ENC_T , by XORing $1 \parallel 0^{c_2-1}$ instead of $0^{c_2-1} \parallel 1$. By the same reasoning as given for ASCON, this is a purely cosmetic change which allows us to view ENC_T as a sponge with increased rate of $r + 1$. In particular, it neither affects the security of ISAP nor defeats the purpose of the domain separation.

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 each time 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 κ in the following proof.

The proof for ISAP follows the same idea as the one for ASCON, again using the collision resistance of plain sponges. There are two main differences, though. First, ISAP features a larger capacity in the initial state which allows to directly apply Theorem 9.²⁸ Second, some extra care is necessary to handle the re-keying mechanism of ISAP, which essentially results in two applications of Theorem 9.

Theorem 16. *Consider ISAP which is illustrated and described in Fig. 30 and Fig. 31, respectively. Let ρ_1 and ρ_2 be modeled by ideal permutations ρ_1 and ρ_2 , respectively. Then for any adversary \mathcal{A} making q_1 and q_2 queries to ρ_1 and ρ_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. Let $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ be the output of a CMT adversary \mathcal{A} against ISAP. Further note that IV denotes the initialization vector used in ISAP. We assume that \mathcal{A} makes queries to ρ_1 and ρ_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 \mathcal{A} into one that runs \mathcal{A}

²⁷The proof is independent of ENC_M which is why we do not need a third permutation.

²⁸In a similar vein, one can use the old indistinguishability bound [BDPV08] for ISAP, whereas ASCON requires the newer version [NO14].

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

Figure 28: Pseudocode of XOODYAK [DHP⁺21] in terms of ENC $_e$ and ENC \mathcal{M} .




Figure 29: Illustration of ISAP.RK.

to obtain $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ and—before outputting the same—makes all queries to ρ corresponding to the evaluation of $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$.

The authentication component ENC \mathcal{T} uses a session key denoted by K_A (resp. \bar{K}_A), which results from an application of ISAP.RK to the key K (resp. \bar{K}) and the intermediate state X (resp. \bar{X}) computed during ENC \mathcal{T} . We consider the event \mathbf{E} that $(N, A) = (\bar{N}, \bar{A})$ and $K_A = \bar{K}_A$. 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 4) adversary \mathcal{B} against a sponge-based hash function \mathcal{H}_2 obtained from the permutation ρ_2 with rate $\bar{r}_2 = \max\{\kappa, r_2 + 1\}$ ²⁹, capacity $\bar{c}_2 = n - \bar{r}_2$, and output length κ . Further $0^\kappa \parallel \text{IV}$ is chosen as the initial state of \mathcal{H}_2 . First, \mathcal{B} runs \mathcal{A} , which outputs $(K, N, A, M), (\bar{K}, \bar{N}, \bar{A}, \bar{M})$. For every query that \mathcal{A} makes to ρ_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 ρ_1 for \mathcal{A} . Using this, \mathcal{B} computes for both output tuples of \mathcal{A} ,

²⁹The 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.

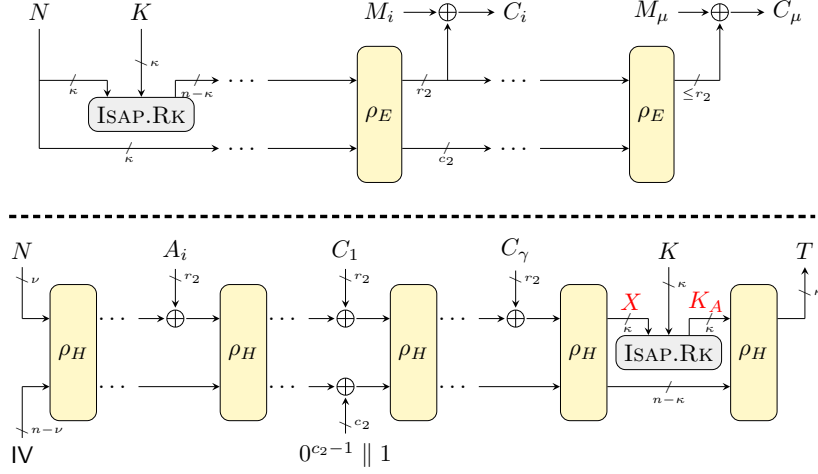


Figure 30: Illustration of ISAP in terms of $\text{ENC}_{\mathcal{M}}$ (top) and $\text{ENC}_{\mathcal{T}}$ (bottom), which both rely on the re-keying function ISAP.RK. The values X and K_A , marked in red, are used in our CMT proof.

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 \bar{X} , respectively and the session keys for $\text{ENC}_{\mathcal{T}}$ by K_A and \bar{K}_A , respectively (cf. Fig. 32). The states obtained after XORing these together are denoted by $Z = X \oplus K_A$ and $\bar{Z} = \bar{X} \oplus \bar{K}_A$.³⁰

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

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

where $\parallel 0^*$ denotes the padding with 0s up to length \bar{r}_2 . A visualization for this is provided in Fig. 32. 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 $Z = [X]_\kappa \oplus K_A \neq [X]_\kappa \oplus \bar{K}_A = \bar{Z}$. Thus, $O \neq \bar{O}$ is also given in this case. So far, 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 9, which bounds the probability of finding a collision in a general sponge-based hash function. Here, we exploit the fact that \mathcal{B} makes the same number of queries to ρ_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-based hash function \mathcal{H}_1 obtained from the permutation ρ_1

³⁰Note 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} .

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 : return (C, T)	21 : $S \leftarrow K_E \parallel N$
ENC \mathcal{T} (K, N, A, C)	22 : for $i = 1, \dots, \mu$
4 : $A_1, \dots, A_\alpha \xleftarrow{r_2} \text{pad}_{10^*}(A, r_2)$	23 : $S \leftarrow \rho_E(S)$
5 : $C_1, \dots, C_\gamma \xleftarrow{r_2} \text{pad}_{10^*}(C, r_2)$	24 : $C_i \leftarrow \lceil S \rceil_{r_2} \oplus M_i$
6 : $Y \leftarrow N \parallel \text{IV}$	25 : $C \leftarrow \lceil C_1 \parallel \dots \parallel C_\mu \rceil_{ M }$
7 : $S \leftarrow \rho_H(Y)$	26 : return C
8 : for $i = 1, \dots, \alpha$	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 : $Y \leftarrow K \parallel \text{IV}$
11 : $S \leftarrow S \oplus 0^{n-1} \parallel 1$	29 : $S \leftarrow \rho_K(Y)$
12 : for $i = 1, \dots, \gamma$	30 : for $i = 1, \dots, z - 1$
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 : $K_A \leftarrow \text{ISAP.RK}(K, \lceil S \rceil_\kappa)$	33 : $Y \leftarrow S \oplus (X_z \parallel 0^{n-r_1})$
16 : $S \leftarrow \rho_H(K_A, \lceil S \rceil_\kappa)$	34 : $S \leftarrow \rho_K(Y)$
17 : $T \leftarrow \lceil S \rceil_\tau$	35 : return $\lceil S \rceil_\kappa$
18 : return T	

Figure 31: Pseudocode of ISAP [DEM⁺21] in terms of ENC \mathcal{M} and ENC \mathcal{T} .

with rate $\bar{r}_1 = \kappa$, capacity $\bar{c}_1 = n - \bar{r}_2$, and output length κ . Further its initial state is given by $0^\kappa \parallel \text{IV}$. 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 ρ_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 ρ_2 for \mathcal{A} . Then, it computes X and \bar{X} (analogously to adversary \mathcal{B}), and we write $X_1, \dots, X_\kappa \stackrel{1}{\leftarrow} X$ and $\bar{X}_1, \dots, \bar{X}_\kappa \stackrel{1}{\leftarrow} \bar{X}$. 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 it outputs are different, but their image under \mathcal{H}_1 agrees (as $K_A = \bar{K}_A$). This implies that

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

where the last inequality holds by Theorem 9. 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. \square




Figure 32: Illustration of a proof step for ISAP ([Theorem 16](#)).. 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 16](#) is $\frac{q_1(q_1-1)}{2^\kappa} + \frac{q_2(q_2-1)}{2^\kappa}$, thus by increasing κ (i.e., the tag and key length), we can increase the committing security. Note however, that—for ISAP-A—we can only increase κ 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 κ up to 200, allowing for about 100-bit committing security.

B.6 SCHWAEMM

SCHWAEMM [[BBC+20](#), [BBdS+21](#)] is a sponge-based AE scheme. The permutation used to instantiate SCHWAEMM is SPARKLE which is inspired by the block-cipher SPARX [[DPU+16](#)]. The authentication mode is a variant of the BEETLE mode [[CDNY18](#)].

B.6.1 Description of SCHWAEMM

The pseudocode of SCHWAEMM is given in [Fig. 34](#) and further illustration is provided in [Fig. 33](#). SCHWAEMM follows the CpP-approach, i.e., first the context is processed resulting in $S \leftarrow \text{ENC}_{\mathcal{C}}(K, N, A)$ and afterwards the ciphertext is computed as $(C, T) \leftarrow \text{ENC}_{\mathcal{M}}(K, S, M)$. In SCHWAEMM the underlying permutation ρ is applied a varying number of times, depending on the position in the sponge (ρ^a and ρ^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

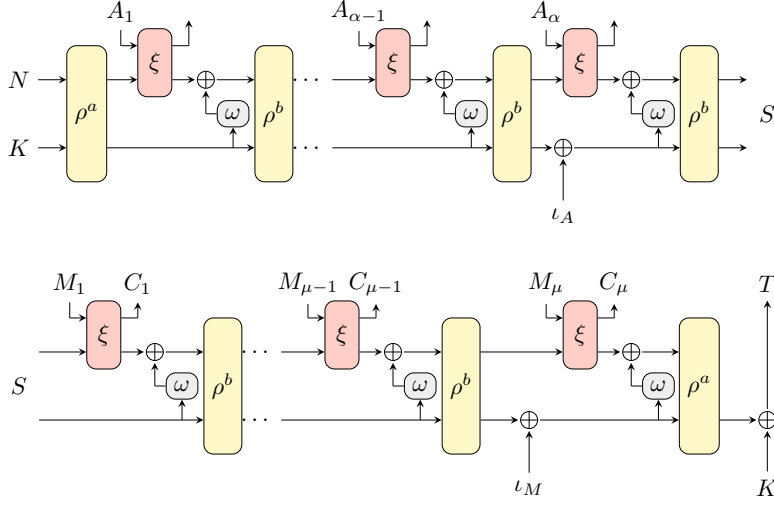


Figure 33: Illustration of SCHWAEMM in terms of ENC_c (top) and ENC_M (bottom).

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 κ bits are XORed with the key to yield the tag. As for ASCON, we refer to this as output-blinding.

B.6.2 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_c 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. 33), that makes the attack unlikely to succeed. Output-blinding is a feature we have also encountered in ASCON, as an important feature for achieving committing security. 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 bits³¹ and instead incorporate a fixed IV (of length 128 bits 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 ρ^a and ρ^b by one ideal permutation ρ , 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 17. *Consider SCHWAEMM which is illustrated and described in Fig. 33 and Fig. 34, respectively. and SCHWAEMM_{IV}, its modified version described above. Let ρ^a and ρ^b be modeled as a random permutation ρ . Then for any adversary A making $q \leq 2^{127}$*

³¹Note that the modified scheme is still in accordance to the NIST requirements [NIST15] that nonces are at least 96 bits long.

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

Figure 34: Pseudocode of SCHWAEMM [BBdS⁺21] in terms of ENC $_e$ and ENC \mathcal{M} .

queries to ρ , it holds that

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

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

Proof. Firstly, we observe that finding different inputs to SCHWAEMM $_{\text{IV}}$.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 SCHWAEMM $_{\text{IV}}$ 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 TagColl (see Fig. 15). Hence we can deduce

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

As a next step, we pass over to a plain sponge construction. For this, we construct a ShiftedColl $_{128}$ adversary \mathcal{B} against a sponge-based hash function \mathcal{H} obtained from the permutation ρ with rate 256, capacity 128, and output length 128. Further, its initial state is given by $0^{256} \parallel \text{IV}$. 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 ρ , 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 SCHWAEMM $_{\text{IV}}$ 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

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

which guarantees that $\mathcal{H}(X)$ (and $\mathcal{H}(\bar{X})$, respectively) models $\text{SCHWAEMM}_{\text{IV}}$ evaluated on (K, N, A, M) (and $(\bar{K}, \bar{N}, \bar{A}, \bar{M})$, respectively). More precisely, instead of the state-update-function, XORing of the input, and rate-whitening applied in $\text{SCHWAEMM}_{\text{IV}}$, for \mathcal{H} we XOR a suitable value which imitates these operations. A visualization for this is provided in Fig. 35.

Next, we show that if \mathcal{A} wins the game TagColl against $\text{SCHWAEMM}_{\text{IV}}$, then the constructed adversary \mathcal{B} wins the game ShiftedColl_{128} against \mathcal{H} . First observe that \mathcal{A} winning implies that $(K, N, A, M) \neq (\bar{K}, \bar{N}, \bar{A}, \bar{M})$ and the corresponding tags T and \bar{T} —computed with $\text{SCHWAEMM}_{\text{IV}}$ —agree. Note that the latter implies that $S_{\alpha+\mu+1,c} \oplus K = \bar{S}_{\alpha+\mu+1,c} \oplus \bar{K}$, hence by choice of X and \bar{X} it holds that $\mathcal{H}(X) \oplus [X]_{128} = \mathcal{H}(\bar{X}) \oplus [\bar{X}]_{128}$. Further, the fact that $(K, N, A) \neq (\bar{K}, \bar{N}, \bar{A})$, implies that $X \neq \bar{X}$: for $(K, N) \neq (\bar{K}, \bar{N})$ this is obvious while for $(K, N) = (\bar{K}, \bar{N})$ and $A \neq \bar{A}$, a simple analysis shows that X and \bar{X} differ at the point where the associated data blocks differ for the first time. This implies that \mathcal{B} wins the game ShiftedColl_{128} .

Using the indistinguishability of sponges (cf. Theorem 10), 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 10.

As a last step, we observe that for a random function $F : \{0, 1\}^{\geq 128} \rightarrow \{0, 1\}^{128}$, it is unlikely that \mathcal{B} wins the game ShiftedColl_{128} . For this, note that an adversary that wins the game 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 11, 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. \square

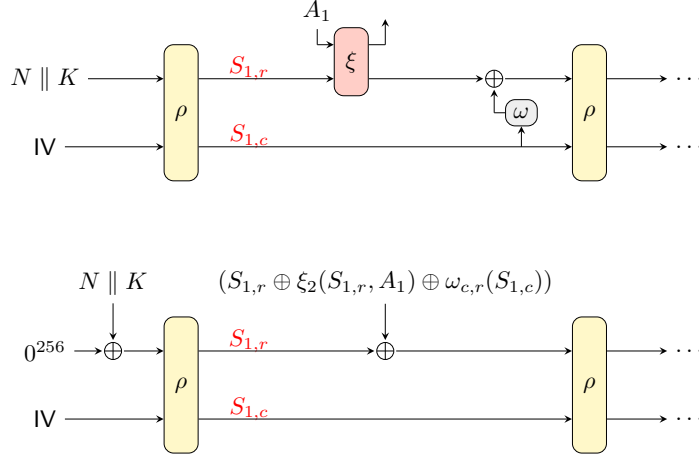


Figure 35: Illustration of a proof step for $\text{SCHWAEMM}_{\text{IV}}$ (Theorem 17). $\text{SCHWAEMM}_{\text{IV}}$ is represented as a plain sponge as is shown for the first XOR in the above figure.

Game $\text{ShiftedColl}_\kappa$

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

Figure 36: Security game $\text{ShiftedColl}_\kappa$ defined for a function $F: \{0, 1\}^{\geq \kappa} \rightarrow \{0, 1\}^\kappa$ and used in the proof for $\text{SCHWAEMM}_{\text{IV}}$ (Theorem 17).

C Deferred Proofs

C.1 Proof of Lemma 2

Proof. We construct \mathcal{A}_c as shown in Fig. 5. As input it receives (K, N, S) . It chooses an arbitrary even number of associated data block α and chooses all except the first block at random, i.e., $A_2, \dots, A_\alpha \leftarrow_s \{0, 1\}^n$.³² In addition, \mathcal{A} sets $A_{\alpha+1} \leftarrow 0^n$.³³ 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 $Y \leftarrow \tilde{E}^{-1}(K, A_{2i}, S)$, inverting the state-update-function ξ and the ideal tweakable cipher. Denote the resulting state by Y_* (see also the state marked in red in Fig. 3). Adversary \mathcal{A} sets the first associated data as $A_1 \leftarrow Y_*$ which ensures that $\text{ENC}_c(K, N, A) = S$.

The number of queries to the ideal tweakable cipher \tilde{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

C.2 Proof of Lemma 3

Proof. We construct adversary \mathcal{A}_M as shown in Fig. 5 which 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 .³⁴ Let $C_1, \dots, C_\gamma \stackrel{n}{\leftarrow} C$. The adversary \mathcal{A}_M first sets $Y \leftarrow T$ and then computes

³²We assume that these blocks are chosen to exhibit a valid padding.

³³This corresponds to the input of the last application of ξ in the upper part of Fig. 3.

³⁴This is justified by letting \mathcal{A} choose a message satisfying this.

$S \leftarrow G^{-1}(Y)$. For $i \in \{1, \dots, \gamma\}$, \mathcal{A}_M computes $Y \leftarrow \tilde{\mathbf{E}}^{-1}(K, N, S)$ ³⁵ followed by the computation of $(S, M_i) \leftarrow \mathcal{A}_\xi(Y, C_i)$ from Lemma 1. Finally, \mathcal{A}_M outputs (S, M) , where $M = M_1 \parallel \dots \parallel M_\gamma$. By construction, it holds that $\text{ENC}_M(K, N, S, M) = (C, T)$ as \mathcal{A}_M inverted all invocations of $\tilde{\mathbf{E}}$ and ξ during ENC_M —using Lemma 1 for the latter.

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

C.3 Proof of Lemma 4

Proof. For the proof of the claim, we define the following auxiliary events:

1. Event \mathbf{E}_t (target hitting query):

\mathcal{A} makes a query Y to ρ such that $\lfloor \rho(Y) \rfloor_{c^*} = \text{IV}$ or \mathcal{A} makes a query S to ρ^{-1} such that $\lfloor \rho^{-1}(S) \rfloor_{c^*} = \text{IV}$.

2. Event \mathbf{E}_c (colliding queries):

\mathcal{A} makes queries $Y \neq \bar{Y}$ to ρ such that $\lfloor \rho(Y) \rfloor_c = \lfloor \rho(\bar{Y}) \rfloor_c$ or \mathcal{A} makes queries Y to ρ and \bar{S} to ρ^{-1} such that $\lfloor \rho(Y) \rfloor_c = \lfloor \rho^{-1}(\bar{S}) \rfloor_c$.

Next, we show that if \mathcal{A} triggers PP , then it triggers one of the events defined above. For the proof assume that PP holds and denote the problematic paths \mathcal{A} finds by (P, \bar{P}) . We first consider the case that there is at least one backward edge in (P, \bar{P}) . We assume without loss of generality that P 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 also for the minimal path corresponding to P in the following and write it as

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

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 ρ^{-1} . By construction, we have $\lfloor Y_0 \rfloor_{c^*} = \text{IV}$ and $\rho^{-1}(S_1) = Y_0$, hence in particular $\lfloor \rho^{-1}(S_1) \rfloor_{c^*} = \lfloor Y_0 \rfloor_{c^*} = \text{IV}$. Thus \mathcal{A} 's query triggered event \mathbf{E}_t .

Case 2: $s \geq 2$

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

We now consider the case that the penultimate states Y_{s-1} and $\bar{Y}_{\bar{s}-1}$ are equal, but P and \bar{P} contain no backward edges. For such a pair of paths (P, \bar{P}) , 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, \bar{P}) and—for sake of simplicity—also denote them by (P, \bar{P}) . As before we use the following representation

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

³⁵Note that we drop the counter which is part of the tweak for simplicity.

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

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

We have $[Y_0]_{c^*} = \text{IV}$ by construction and $Y_0 = \bar{Y}_{\bar{s}-1}$ by assumption. This allows to deduce $[\bar{S}_{\bar{s}-1}]_{c^*} = \text{IV}$. Since $\bar{S}_{\bar{s}-1}$ is the result of querying $\bar{Y}_{\bar{s}-2}$ to ρ , this query triggered event $\bar{\mathbf{E}}_{\bar{s}-1}$.

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

The paths are

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

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 $[S_{s-1}]_c = [\bar{S}_{\bar{s}-1}]_c$. 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 $\bar{\mathbf{E}}_c$.

Collecting the above, yields

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

We apply Lemma 6 for c^* and obtain $\Pr[\mathbf{E}_t] \leq \frac{q}{2^{c^* - 1}}$ and once again for c , which gives $\Pr[\mathbf{E}_c] \leq \frac{q(q-1)}{2^c}$. \square

C.4 Proof of Theorem 3

Proof. In the proof of [BS23, Theorem 8.6], the advantage is decomposed as follows

$$\text{Adv}_{\mathcal{H}}^{\text{CR}}(\mathcal{A}) = \Pr[\text{CP}] \leq \Pr[\text{CP} \wedge \neg \text{PP}] + \Pr[\text{PP}],$$

for CP the event that \mathcal{A} finds a pair of colliding paths. Bounding the first summand works as in the original proof—hence we do not repeat it here—and results in $\Pr[\text{CP} \wedge \neg \text{PP}] \leq \frac{q(q-1)}{2^w}$. For the second summand, we apply Lemma 4, which gives $\Pr[\text{PP}] \leq \frac{q}{2^{c^* - 1}} + \frac{q(q-1)}{2^c}$. \square

C.5 Proof of Theorem 5

Proof. Adversary \mathcal{B} simply runs $(K, N, A), (\overline{K}, \overline{N}, \overline{A}) \leftarrow \mathcal{A}$, picks an arbitrary message M , and outputs $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, M)$. Since \mathcal{A} wins the game KeyColl , it holds that $K \neq \overline{K}$ and

$$\text{ENC}_c(K, N, A) = \text{ENC}_c(\overline{K}, \overline{N}, \overline{A}) =: S_*.$$

Using the fact that AE is a full-CpP scheme, we get

$$\begin{aligned} \text{ZP-AE.ENC}(K, N, A, M) &= \text{ENC}_{\mathcal{M}}(\text{ENC}_c(K, N, A), 0^z \parallel M) \\ &= \text{ENC}_{\mathcal{M}}(S_*, 0^z \parallel M) \\ &= \text{ENC}_{\mathcal{M}}(\text{ENC}_c(\overline{K}, \overline{N}, \overline{A}), 0^z \parallel M) \\ &= \text{ZP-AE.ENC}(\overline{K}, \overline{N}, \overline{A}, M). \end{aligned}$$

Thus, \mathcal{B} wins game $\text{CMT}_{\mathcal{K}}$. \square

C.6 Proof of Theorem 7

Proof. For sake of simplicity, we give the proof for the case $z = n$. The adversary \mathcal{A} generates an arbitrary context-message-pair (K, N, A, M) and computes the ciphertext $(C, T) \leftarrow \text{ZP-ELEPHANT.ENC}(K, N, A, M)$. Let $M_1, \dots, M_\mu \stackrel{n}{\leftarrow} \text{pad}_{0^*}(0^z \parallel M, n)$ be the μ message blocks, where it holds that $M_1 = 0^n$. Therefore, the ciphertext consists of μ n -bit blocks $C = C_1 \parallel \dots \parallel C_\mu$, where $C_1 = \tilde{\text{E}}(K, (0, 1), N) \oplus M_1 = \tilde{\text{E}}(K, (0, 1), N)$. In order to do these computations, \mathcal{A} needs to make μ queries to $\tilde{\text{E}}$.

Next, \mathcal{A} chooses $\overline{K} \neq K$ at random to make sure that the attack is a valid $\text{CMT}_{\mathcal{K}}$ attack. To find the second nonce \overline{N} , \mathcal{A} queries $(\overline{K}, (0, 1), C_1)$ to $\tilde{\text{E}}^{-1}$ and sets the response to be \overline{N} . This guarantees that the first ciphertext block C_1 is an encryption of the all-zero message when using \overline{N} . To find \overline{M} , \mathcal{A} computes $\overline{M}_i \leftarrow \tilde{\text{E}}(\overline{K}, (i-1, 1), \overline{N}) \oplus C_i$, for $i \in \{2, \dots, \mu\}$. This then yields that (K, N, M) and $(\overline{K}, \overline{N}, \overline{M})$ yield the same ciphertext, i.e.,

$$\text{ENC}_{\mathcal{M}}(\overline{K}, \overline{N}, 0^z \parallel \overline{M}) = C = \text{ENC}_{\mathcal{M}}(K, N, 0^z \parallel M).$$

Next, \mathcal{A} chooses arbitrary A and computes $T \leftarrow \text{ENC}_{\mathcal{T}}(K, N, A, C)$. It remains to find \overline{A} that result in the same tag, i.e., $\text{ENC}_{\mathcal{T}}(\overline{K}, \overline{N}, \overline{A}, C) = T$. The attack is effectively the existing committing attack against ELEPHANT (cf. Theorem 12); the mere difference is that \mathcal{A} targets the second associated data block since it has already fixed \overline{N} which determines part of the first associated data block. Adversary \mathcal{A} then outputs $(K, N, A, M), (\overline{K}, \overline{N}, \overline{A}, M)$. From the above we can conclude that the outputs by \mathcal{A} result in the same ciphertext (C, T) , which, together with the fact that $K \neq \overline{K}$, yields that

$$\text{Adv}_{\text{ZP-ELEPHANT}}^{\text{CMT}_{\mathcal{K}}}(\mathcal{A}) = 1.$$

The number of queries to $\tilde{\text{E}}$ by \mathcal{A} accumulate to $2\mu + 2\gamma + \alpha + \bar{\alpha}$: 2μ for computing (K, N, M) and $(\overline{K}, \overline{N}, \overline{M})$; and $2\gamma + \alpha + \bar{\alpha}$ for computing A and \overline{A} .

The proof generalizes to the case $z \leq n$ by letting \mathcal{A} choose the message such that the first bits are 0^{n-z} , which ensures that the first message block is 0^n . \square

C.7 Proof of Theorem 8

Proof. The adversary \mathcal{A} picks a key-nonce pair (K, N) at random. It then picks an arbitrary message M and computes $C \leftarrow \text{ENC}_{\mathcal{M}}(K, N, 0^z \parallel M)$ —making μ queries to ρ_3 and $\nu + 1$ queries to ρ_1 . Subsequently, the adversary applies a birthday attack on the tag, by computing $\text{ENC}_{\mathcal{T}}(K, N, \cdot, C)$ for different associated data. By [BS23, Section 8.3], a

collision is found with probability $\frac{1}{2}$ after trying $2^{\frac{\tau}{2}+2} + 1$ different associated data. Let A and \bar{A} be the associated data that led to a collision as part of the birthday attack. Finally, \mathcal{A} outputs $(K, N, A, M), (K, N, \bar{A}, M)$.

For each associated data, \mathcal{A} makes $\alpha + \gamma + 2$ queries to ρ_2 (evaluating $\text{ENC}_{\mathcal{T}}$) and $\kappa + 1$ queries to ρ_1 (evaluating ISAP.RK). In total, the queries by \mathcal{A} are $q_1 = (2^{\frac{\tau}{2}+1} + 1)(\kappa + 1) + \nu + 1$ queries to ρ_1 , $q_2 = (2^{\frac{\tau}{2}+1} + 1)(\alpha + \gamma + 2)$ queries to ρ_2 , and $q_3 = \mu$ queries to ρ_3 . \square