

# BUILDING THE NEXT GENERATION OF AUTHENTICATED ENCRYPTION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Sanketh Gora Menda

August 2025

© 2025 Sanketh Gora Menda

ALL RIGHTS RESERVED

# BUILDING THE NEXT GENERATION OF AUTHENTICATED ENCRYPTION

Sanketh Gora Menda, Ph.D.

Cornell University 2025

(sanketh: todo)

## TABLE OF CONTENTS

Table of Contents . . . . .	4
<b>1 Introduction</b>	<b>1</b>
<b>2 Context Discovery and Commitment Attacks</b>	<b>2</b>
2.1 Background . . . . .	10
2.2 Granular Committing Encryption Definitions . . . . .	15
2.3 Context Discovery Attacks against AEAD . . . . .	27
2.4 Restrictive Commitment Attacks via k-Sum Problems . . . . .	32
2.5 Related Work . . . . .	38
<b>3 The OCH Authenticated Encryption Scheme</b>	<b>42</b>
3.1 Background and Related Work . . . . .	46
3.2 Preliminaries . . . . .	49
3.3 The Xor-then-Hash Transform . . . . .	56
3.4 The OCH AEAD Scheme . . . . .	64
3.5 Security Analysis of OCH . . . . .	71
3.6 Performance . . . . .	78
3.7 (In Preparation) TPRP Security of OCT . . . . .	87
3.7.1 Updated TBC Intro . . . . .	87
3.7.2 TPRP Security of OCT . . . . .	87
<b>4 Flexible Authenticated Encryption</b>	<b>97</b>
4.1 Background . . . . .	103
4.2 The Flex Framework . . . . .	106
4.3 ARF: An Instantiation of Flexible AEAD . . . . .	114
4.3.1 ARF Overview . . . . .	115
4.4 Security Analysis of ARF . . . . .	117
4.5 Performance . . . . .	117

## CHAPTER 1

### INTRODUCTION

Our desire to send secret messages through untrusted messengers has remained constant from Julius Caesar in the battlefields of Rome to me sending silly stickers to my friends in Signal group chats. But what has changed is the explosion in scale and complexity. We now have distributed workloads that can encrypt  $2^{32}$  messages in a couple seconds [76].

scale and complexity of transmission. The internet transmits more information in the blink of an eye than all of Caesar's messengers combined in a month. To keep up, we should be able to scramble and unscramble messages just as quickly.

(**sanketh**: Rogaway, then Commitment is the next)

## CHAPTER 2

### CONTEXT DISCOVERY AND COMMITMENT ATTACKS\*

Designers of authenticated encryption with associated data (AEAD) have traditionally targeted security in the sense of confidentiality and ciphertext integrity, first in the context of randomized authenticated encryption [18], and then nonce-based [104] and misuse-resistant AEAD [105].

But in recent years researchers and practitioners have begun realizing that confidentiality and integrity as previously formalized prove insufficient in a variety of contexts. In particular, the community is beginning to appreciate the danger of schemes that are not key committing, meaning that an attacker can compute a ciphertext such that it can successfully decrypt under two (or more) keys. Non-key-committing AEAD was first shown to be a problem in the context of moderation in encrypted messaging [51, 63], and later in password-based encryption [83], password-based key exchange [83], key rotation schemes [2], and symmetric hybrid (or envelope) encryption [2].

Even more recently, new definitions have been proposed [13] that target committing to the key, associated data, and nonce. And while there have been proposals for new schemes [2, 13] that meet these varying definitions, questions still remain about which current AEAD schemes are committing and in which ways. Moreover, there have been no commitment results shown for a number of important practical AEAD schemes, such as CCM [54], EAX [24], and SIV [105]. Implementing (and standardizing) new AEAD schemes takes time and so understanding which standard AEAD schemes can be securely used in which settings is a pressing issue.

---

\*This chapter is joint work with Julia Len, Paul Grubbs, and Thomas Ristenpart. The proceedings version of this paper appears at Eurocrypt 2023 [89], and an earlier version was presented at Real World Crypto 2023 [27]. This chapter is lightly edited from full version [90].

This work makes four main contributions. First, we provide a new, more granular framework for commitment security, which expands on prior ones to better capture practical attack settings. Second, we show the first key commitment attack against the original SIV mode, which was previously an open question. Third, we introduce a new kind of commitment security notion for AEAD—what we call *context discoverability*—which is analogous to preimage resistance for cryptographic hash functions. Fourth, we give context discovery attacks against a range of schemes which, by a general implication, also yield new commitment attacks against those schemes. A summary of our new attacks, including comparison with prior ones, when relevant, is given in Figure 2.1.

**Granular commitment notions.** Recall that a nonce-based AEAD encryption algorithm  $\text{Enc}$  takes as input a key  $K$ , nonce  $N$ , associated data  $A$ , and a message  $M$ . It outputs a ciphertext  $C$ . Decryption  $\text{Dec}$  likewise takes in a  $(K, N, A)$  triple, which we call the decryption *context*, along with a ciphertext  $C$ , and outputs either a message  $M$  or special error symbol  $\perp$ .

While most prior work has focused on key commitment security, which requires commitment to only one part (the key) of the decryption context, Bellare and Hoang (BH) [13] suggest a more expansive sequence of commitment notions for nonce-based AEAD. For the first, CMT-1, an adversary wins if it efficiently computes a ciphertext  $C$  and two decryption contexts  $(K_1, N_1, A_1)$  and  $(K_2, N_2, A_2)$  such that decryption of  $C$  under either context works (does not output  $\perp$ ) and  $K_1 \neq K_2$ . CMT-1 is often called key commitment.<sup>1</sup> CMT-3 relaxes the latter winning condition to allow a win should the decryption contexts differ in any way. We therefore refer to CMT-3 as *context commitment* and schemes that meet CMT-3 as *context committing*. These notions form

---

<sup>1</sup>BH refer to this as CMTD-1, but for tidy AEAD schemes, CMT-1 and CMTD-1 are equivalent, so we prefer the compact term.

a strict hierarchy, with CMT-3 being the strongest. Despite this, most prior attacks [51, 63, 83, 2] have focused solely on key commitment (CMT-1).

Our first contribution is to refine further the definitional landscape for nonce-based AEAD schemes in a way that is particularly useful for exploring context commitment attacks. In practice, attackers will often face application-specific restrictions preventing full control over the decryption context. For example, in the Dodis, Grubbs, Ristenpart, and Woodage (DGRW) [51] attacks against Facebook’s message franking scheme, the adversary had to build a ciphertext that decrypts under two contexts with equivalent nonces. Their (in BH’s terminology) CMT-1 attack takes on a special form, and we would like to be able to formally distinguish between attacks that achieve additional adversarial goals (e.g., different keys but equivalent nonces) and those that may not.

We therefore introduce a new, parameterized security notion that generalizes the BH notions. Our  $\text{CMT}[\Sigma]$  notion specifies what we call a setting  $\Sigma = (\text{ts}, S, P)$  that includes a target specifier  $\text{ts}$ , a context selector  $S$ , and a predicate  $P$ . The parameters  $\text{ts}$  and  $S$  specify which parts of the context are attacker-controlled versus chosen by the game, and which of the latter are revealed to the attacker. Furthermore, the predicate  $P$  takes as input the two decryption contexts and decrypted messages, and outputs whether the pair of tuples satisfy a winning condition. An adversary wins if it outputs a ciphertext and two contexts satisfying the condition that each decrypt the ciphertext without error. The resulting family of commitment notions includes both CMT-1 and CMT-3 but also covers a landscape of further notions.

We highlight two sets of notions. The first set is composed of  $\text{CMT}_k$ ,  $\text{CMT}_n$ , and  $\text{CMT}_a$ , which use predicates  $(K_1 \neq K_2)$ ,  $(N_1 \neq N_2)$ , and  $(A_1 \neq A_2)$ , respectively. The first notion is equivalent to CMT-1; the latter two are new. All of them are orthogonal to each other and a scheme that meets all three simultaneously achieves CMT-3. We say



Scheme	$\text{CDY}_a^*$	$\text{CDY}_n^*$	$\text{CMT}_a^*$	$\text{CMT}_k^*$	$\text{CMT}_k$	CMT-3
GCM [53]	★✗ §2.3	★✗ §2.3	★✗ §E	☆✗ [63, 51]	☆✗ [63, 51]	☆✗ [63, 51]
SIV [107]	★✗ §2.3			★✗ §2.4	★✗ ▶▶	★✗ ▶▶
CCM [54]	★✗ §2.3				★✗ ▶▶	★✗ ▶▶
EAX [24]	★✗ §2.3	★✗ §2.3			★✗ ▶▶	★✗ ▶▶
OCB3 [80]	★✗ §2.3			☆✗ [2]	☆✗ [2]	☆✗ [2]
PaddingZeros	★✓ ▶▶	★✓ ▶▶	★✗ §E	☆✓ [2]	☆✓ [13]	★✗ ▶▶
KeyHashing	★✓ ▶▶	★✓ ▶▶	★✗ §E	☆✓ [2]	☆✓ [2]	★✗ ▶▶
CAU-C1 [13]	★✓ ▶▶	★✓ ▶▶		☆✓ [13]	☆✓ [13]	★✗ ▶▶

Figure 2.1: Summary of context discovery and commitment attacks against a variety of popular AEAD schemes. Symbol ✓ indicates a proof that any attack will take at least  $2^{64}$  time, while symbol ✗ indicates the existence of an attack that takes less than  $2^{64}$  time; symbol ★ indicates results new to this paper and ☆ indicates prior work (citation given).  $\text{CMT}_k$  and CMT-3 are from Bellare and Hoang [13], where  $\text{CMT}_k$  was called CMT-1. The notions  $\text{CDY}_a^*$ ,  $\text{CDY}_n^*$ ,  $\text{CMT}_a^*$ , and  $\text{CMT}_k^*$  are introduced in this paper, and  $\text{CDY}_a^*$ ,  $\text{CDY}_n^*$ , and  $\text{CMT}_k^*$  are implied by  $\text{CMT}_k$ . Symbol ▶▶ indicates that the result is implied from one of the other columns by a reduction shown in this paper. §E indicates Appendix E in the full version.

these notions are *permissive* because the predicates used do not make any demands on other components of the context. In contrast, *restrictive* variants, which we denote via  $\text{CMT}_k^*$ ,  $\text{CMT}_n^*$ , and  $\text{CMT}_a^*$ , require equality for other context components. For example, the first uses the predicate  $(K_1 \neq K_2) \wedge ((N_1, A_1) = (N_2, A_2))$ . These capture the types of restrictions faced in real attacks mentioned above.

**Breaking the original SIV.** While prior work has shown (in our terminology)  $\text{CMT}_k^*$  attacks for GCM [63, 51], GCM-SIV [111, 83], ChaCha20/Poly1305 [63, 83], XChaCha20/Poly1305 [83], and OCB3 [2], an open question of practical interest [112] is whether there also exists a  $\text{CMT}_k^*$  attack against Synthetic IV (SIV) mode [105]. We resolve this open question, showing an attack that works in time about  $2^{53}$ . It requires new techniques compared to prior attacks.

SIV combines a PRF  $F$  with CTR mode encryption, encrypting by first computing a tag  $T = F_K(N, A, M)$  and then applying CTR mode encryption to  $M$ , using  $T$  as the (synthetic) IV and a second key  $K'$ . The tag and CTR mode output are, together, the ciphertext. Decryption recovers the message and then recomputes the tag, rejecting the ciphertext if it does not match. Schmiege [111] and Len, Grubbs, and Ristenpart (LGR) [83] showed that when  $F$  is a universal hash-based PRF, in particular GHASH for AES-GCM-SIV, one can achieve a fast  $\text{CMT}_k^*$  attack.

Their attack does not extend to other versions of SIV, perhaps most notably the original version that uses for  $F$  the S2V[CMAC] PRF [105]. This version has been standardized [68] and is available in popular libraries like Tink [4]. For brevity here we describe the simpler case where  $F$  is just CMAC; the body will expand on the details. At first it might seem that CMAC's well-known lack of collision resistance (for adversarially-chosen keys), should extend to allow a simple  $\text{CMT}_k^*$  attack: find  $K_1, K_2$  such that  $T = \text{CMAC}_{K_1}(N, A, M) = \text{CMAC}_{K_2}(N, A, M')$  for  $M \neq M'$ . But the problem is that we need  $M, M'$  to also satisfy

$$M \oplus \text{CTR}_{K'_1}(T) = M' \oplus \text{CTR}_{K'_2}(T) \quad (2.1)$$

where  $\text{CTR}_K(T)$  denotes running counter mode with initialization vector  $T$  and block cipher key  $K$ . When using a GHASH-based PRF, the second condition “plays well” with the algebraic structure of the first condition, making it computationally easy to satisfy both simultaneously. But, here that does not work.

The core enabler for our attack is that we can recast the primary collision finding goal as a generalized birthday bound attack. For block-aligned messages, we show how the two constraints above can be rewritten as a single equation that is the xor-sum of four terms, each taking values over  $\{0, 1\}^n$ . Were the terms independently and

uniformly random, one would immediately have an instance of a 4-sum problem, which can be solved using Wagner’s k-tree algorithm [116] in time  $\mathcal{O}(2^{n/3})$ . But our terms are neither independent nor uniformly random. Nevertheless, our main technical lemma shows that, in the ideal cipher model, the underlying block cipher and the structure of the terms (which are dictated by the details of CMAC-SIV) allows us to analyze the distribution of these terms and show that we can still apply the k-tree algorithm and achieve the same running time. This technique may be of independent interest.

Using this, we construct a  $\text{CMT}_k^*$  attack against  $\text{S2V}[\text{CMAC}]\text{-SIV}$  that works in time about  $2^{53}$ , making it practical and sufficiently damaging to rule out SIV as suitable for contexts where key commitment matters.

**Context discoverability.** Next we introduce a new type of security notion for AEAD. The cryptographic hashing community has long realized the significance of definitions for both collision resistance and preimage resistance [33], the latter of which, roughly speaking, refers to the ability of an attacker to find some input that maps to a target output. In analyzing  $\text{CMT}_k$  security for schemes, we realized that in many cases we can give very strong attacks that, given any ciphertext, can find a context that decrypts it—a sort of preimage attack against AEAD. To avoid confusion, we refer to this new security goal for AEAD as *context discoverability (CDY)*, as the adversary is tasked with efficiently computing (“discovering”) a suitable context for some target ciphertext.

While we have not seen real attacks that exploit context discoverability, since CDY is to CMT what preimage resistance is to collision resistance, we believe that they are inevitable. We therefore view it beneficial to get ahead of the curve and analyze the CDY security before concrete attacks surface.

We formalize a family of CDY definitions similarly to our treatment for CMT. Our




Figure 2.2: **(Top)** Selected relationships between permissive CMT notions and restrictive CDY notions. Solid arrows represent implications. **(Bottom)** Selected relationships between CMT-3 and the notions we introduce in this paper. Solid arrows represent implications. The dotted arrow from CMT-3 to  $CDY^*$  holds assuming “context compression” as defined in Theorem 1.

$CDY[\Sigma]$  notion is parameterized by a setting  $\Sigma = (ts, S)$  that specifies a target specifier  $ts$  and a context selector  $S$ . Like for  $CMT[\Sigma]$ ,  $ts$  and  $S$  specify the parts of the context that the attacker can choose and which parts are chosen by the game and either hidden or revealed to the attacker. Unlike CMT, however, the attacker is always given a target ciphertext and needs to only produce one valid decrypting context.

Similar to  $CMT_k^*$ ,  $CMT_n^*$ ,  $CMT_a^*$ , we define the notions  $CDY_k^*$ ,  $CDY_n^*$ ,  $CDY_a^*$ . The notion  $CDY_k^*$  captures the setting where an adversary is given arbitrary ciphertext  $C$ , nonce  $N$ , and associated data  $A$ , and must produce a key  $K$  such that  $C$  decrypts under  $(K, N, A)$ . Similarly,  $CDY_n^*$  and  $CDY_a^*$  require the adversary to provide a nonce and associated data, respectively, given the other components chosen arbitrarily. These model restricted attack settings where parts of the context are not in the adversary’s control.

We also define  $CDY^*[ts]$  which generalizes this intuition to any target specifier  $ts$ . For example, in  $CDY^*[ts = \{n\}]$  the adversary is given arbitrary ciphertext and nonce  $N$ , and must produce a key  $K$  and associated data  $A$  such that the ciphertext decrypts

under  $(K, N, A)$ .

We next analyze the relations between these sets of notions. In particular, we show that if an AEAD scheme is “context compressing”—ciphertexts are decryptable under more than one context—then CMT-3 security implies  $\text{CDY}^*$ . This is analogous to collision resistance implying preimage resistance, though the details are different. Further, we observe that almost all deployed AEAD schemes are context compressing since they “compress” the nonce and associated data into a shorter tag. This allows us to focus on finding  $\text{CDY}^*[\Sigma]$  attacks for AEAD schemes to show that these schemes also do not meet  $\text{CMT}[\Sigma]$  security. Selected relationships are shown in Figure 2.2.

This opens up a new landscape of analysis, which we explore. We characterize a large class of AEAD schemes that use non-preimage resistant MACs and, based on this weakness, develop fast  $\text{CDY}_a^*$  attacks. The set includes CCM, EAX, SIV, GCM, and OCB3. For EAX and CCM, this represents the first attacks of any kind for committing security. For EAX and GCM, we are also able to give  $\text{CDY}_n^*$  attacks, which is perhaps even more surprising a priori, given that an adversary in this case only controls the nonce.

All this sheds light on the deficiencies of several popular design paradigms for AEAD, when viewed from the perspective of context commitment security. These definitions also allow us to precisely communicate attacks and threat models. For example,  $\text{CDY}$  might suffice for some applications while others might want the more computationally expensive CMT security.

**Revisiting commitment-enhancing mechanisms** Finally, in Appendix E in the full version we use this new framework to analyze proposed mechanisms for commitment security. First, we look at the folklore padding zeros transform which pre-

fixes zeroes to a message before encrypting and verifies the existence of these zeroes at decryption. This transform was recommended in an early OPAQUE draft specification [78, §3.1.1] and was shown by Albertini et al. [2, §5.3] to achieve FROB security and by Bellare and Hoang [13] to achieve CMT-1 security. We show that this transform does not achieve our  $\text{CMT}_a^*$  notion (and thus CMT-3) for all AEAD schemes, ruling it out as a candidate commitment security transform. We then make similar observations about the CommitKey transform which appends to the ciphertext a hash commitment to the key and the nonce. Finally, we conclude by considering the practical key commitment security of the recent CAU-C1 scheme from BH [13]. While a naive adaptation of DGRW’s [51] “invisible salamanders” attack to this scheme takes about  $2^{81}$  time, we show a more optimized attack which takes a little more than  $2^{64}$  time, showing that 64-bit key-committing security does not preclude practical attacks.

**Next steps and open problems.** Our results resolve a number of open problems about AEAD commitment security, and overall highlight the value of new definitional frameworks that surface different avenues for attack. That said, we leave several open problems, such as whether different flavors of context discovery or commitment attacks can be found against popular schemes—the blank entries in Figure 2.1. Our attack techniques do not seem to work against these schemes, but whether positive security results can be shown is unclear.

## 2.1 Background

**Notation.** We refer to elements of  $\{0, 1\}^*$  as *bitstrings*, denote the length of a bitstring  $x$  by  $|x|$  and the left-most (i.e., “most-significant”) bit by  $\text{msb}(x)$ . Given two bitstrings  $x$  and  $y$ , we denote their concatenation by  $x \parallel y$ , their bitwise xor by  $x \oplus y$ , and their

bitwise and by  $x \& y$ . Given a number  $n$ , we denote its  $m$ -bit encoding as  $\text{encode}_m(n)$ . For a finite set  $X$ , we use  $x \leftarrow^s X$  to denote sampling a uniform, random element from  $X$  and assigning it to  $x$ .

Sometimes, we operate in the finite field  $\text{GF}(2^n)$  with  $2^n$  elements. This field is defined using an irreducible polynomial  $f(\alpha)$  in  $\text{GF}(2)[\alpha]$  of degree  $n$ . The elements of the field are polynomials  $x_0 + x_1\alpha + x_2\alpha^2 + \dots + x_{n-1}\alpha^{n-1}$  of degree  $n-1$  with binary coefficients  $x_i \in \text{GF}(2)$ . These polynomials can be represented by the  $n$ -bit string  $x_0x_1 \dots x_{n-1}$  of their coefficients. Both addition and subtraction of two  $n$ -bit strings, denoted  $x + y$  and  $x - y$ , respectively, is their bitwise xor  $x \oplus y$ . Multiplication of two  $n$ -bit strings, denoted  $x \cdot y$ , corresponds to the multiplication of the corresponding polynomials  $x$  and  $y$  followed by modular reduction with the irreducible polynomial  $f(\alpha)$ .

**Probability.** An  $n$ -bit random variable  $X$  is one whose value is probabilistically assigned, defined by *probability mass function*  $p_X(x) := \Pr[X = x]$ . The  $n$ -bit uniform random variable  $U$  is the random variable with the probability mass function  $p_U(x) = \frac{1}{2^n}$  for all  $x \in \{0, 1\}^n$ . Given two  $n$ -bit random variables  $X$  and  $Y$ , we define the *total variation distance* between them

$$\Delta(X, Y) := \max_{i \in \{0, 1\}^n} \left| \Pr(X = i) - \Pr(Y = i) \right|.$$

A *random function*  $F$  from  $n$ -bit strings to  $m$ -bit strings is a collection  $\{X_i : i \in \{0, 1\}^n\}$  of  $m$ -bit random variables  $X_i$ , one for each  $n$ -bit input, such that for all  $i \in \{0, 1\}^n$ ,  $F(i) := X_i$ . A random function  $F$  from  $n$ -bit strings to  $m$ -bit strings is *uniformly random* if, for all  $i \in \{0, 1\}^n$ ,  $F(i)$  is the  $m$ -bit uniform random variable. We say that two random functions  $F_1$  and  $F_2$  from  $n$ -bit strings to  $m$ -bit strings are *independent* if, for all  $i \in \{0, 1\}^n$  and for all  $j \in \{0, 1\}^n$ ,  $F_1(i)$  and  $F_2(j)$  are independent  $m$ -bit random variables.

**Code-based games** To formalize security experiments, we use the *code-based games* framework of Bellare and Rogaway [22]; with refinements from Ristenpart, Shacham, and Shrimpton [102]. A *procedure*  $P$  is a sequence of code-like statements that accepts some input and produces some output. We use superscripts like  $P^Q$  to denote that procedure  $P$  calls procedure  $Q$ . We use  $(G \Rightarrow x)$  to denote the event that the procedure  $G$  outputs  $x$ , over the random coins of the procedure. Finally, given a game  $G$  and an adversary  $\mathcal{A}$ , we denote the *advantage* of  $\mathcal{A}$  at  $G$  by  $\text{Adv}_G(\mathcal{A}) := \Pr[G(\mathcal{A}) \Rightarrow \text{true}]$ .

**Cost of attacks.** We represent cryptanalytic attacks by procedures and compute their cost using a *unit-cost RAM model*. Specifically, following [102], we use the convention that each pseudocode statement of a procedure runs in unit time. This lets us write the running time of a procedure as the maximum number of statements executed, with the maximum taken over all inputs of a given size. Similarly, we define the number of queries as the maximum number of queries executed over inputs of a given size. We recognize that this is a simplification of the real-world (e.g., see Wiener [117]), but for the attacks discussed in this paper, we nevertheless believe that it provides a good estimate.

**Pseudorandom functions.** A *pseudorandom function (PRF)* is a function  $F : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{Y}$  defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , and output space  $\mathcal{Y} \subseteq \{0, 1\}^*$ , that is indistinguishable from a uniform random function. More formally, we define the PRF advantage of an adversary  $\mathcal{A}$  as

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) := \left| \Pr[K \leftarrow \mathcal{K} : \mathcal{A}(F(K, \cdot))] - \Pr[R \leftarrow \text{Func} : \mathcal{A}(R)] \right|,$$

and say that  $F$  is a PRF if this advantage is small for all adversaries  $\mathcal{A}$  that run in a feasible amount of time.



**Hash functions.** A *hash function* is a function  $H : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{Y}$ , defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , and hash space  $\mathcal{Y} \subseteq \{0, 1\}^*$ . We define the collision-resistance advantage of adversary  $\mathcal{A}$  for  $H$  as

$$\text{Adv}_{\text{H}}^{\text{coll}}(\mathcal{A}) := \Pr \left[ K \leftarrow \mathcal{K}, (M_1, M_2) \leftarrow \mathcal{A}(K) : (M_1 \neq M_2) \text{ and } (H(K, M_1) = H(K, M_2)) \right].$$

**Block ciphers and the ideal cipher model.** An  $n$ -bit *block cipher*, or a block cipher with *block length*  $n$  bits, is a function  $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ , where for each key  $k \in \{0, 1\}^n$ ,  $E(k, \cdot)$  is a permutation on  $\{0, 1\}^n$ . Since it is a permutation, it has an inverse which we denote by  $E^{-1}(k, \cdot)$ . To simplify notation, we sometimes use the shorthands  $E_k(\cdot) := E(k, \cdot)$  and  $E_k^{-1}(\cdot) := E^{-1}(k, \cdot)$ .

An  $n$ -bit *ideal block cipher* [77] is a random map  $E : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ , such that for each key  $k \in \{0, 1\}^n$ ,  $E_k(\cdot)$  is a permutation on  $\{0, 1\}^n$ . Alternatively, we can think of an ideal block cipher as one where for each key  $k \in \{0, 1\}^n$ ,  $E_k(\cdot)$  is uniformly, randomly sampled from the set of permutations on  $n$ -bits.

**Authenticated encryption schemes.** An *AEAD scheme* is a triple of algorithms  $\text{AEAD} = (\text{Kg}, \text{Enc}, \text{Dec})$ , defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , nonce space  $\mathcal{N} \subseteq \{0, 1\}^*$ , associated data space  $\mathcal{A} \subseteq \{0, 1\}^*$ , message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , and ciphertext space  $\mathcal{C} \subseteq \{0, 1\}^*$ .

1.  $\text{Kg} : \emptyset \rightarrow \mathcal{K}$  is a randomized algorithm that takes no input and returns a fresh secret key  $K$ .
2.  $\text{Enc} : (\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}) \rightarrow (\mathcal{C} \cup \{\perp\})$  is a deterministic algorithm that takes a 4-tuple of a key  $K$ , nonce  $N$ , associated data  $A$ , and message  $M$  and returns a ciphertext  $C$  or an error (denoted by  $\perp$ ).

3.  $\text{Dec} : (\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C}) \rightarrow (\mathcal{M} \cup \{\perp\})$  is a deterministic algorithm that takes a 4-tuple of a key  $K$ , nonce  $N$ , associated data  $A$ , and ciphertext  $C$  and returns a plaintext  $M$  or an error (denoted by  $\perp$ ).

We call the non-message inputs to  $\text{Enc}$ —the key, nonce, and associated data—the *encryption context* and the non-ciphertext inputs to  $\text{Dec}$ —the key, nonce, and associated data—the *decryption context*. And, for a given message, say that an encryption context is *valid* if  $\text{Enc}$  succeeds (i.e., does not output  $\perp$ ). Similarly, for a given ciphertext, say that a decryption context is *valid* if  $\text{Dec}$  succeeds (i.e., does not output  $\perp$ ).

For traditional AEAD correctness, we need  $\text{Enc}$  to be the inverse of  $\text{Dec}$ . In other words, for any 4-tuple  $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ , it holds that

$$\text{Dec}(K, N, A, \text{Enc}(K, N, A, M)) = M.$$

In addition, we impose *tidyness* [94], *ciphertext validity*, and *length uniformity* assumptions. Tidyness requires that for any 4-tuple  $(K, N, A, C) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C}$ , it holds that

$$\text{Dec}(K, N, A, C) = M \neq \perp \implies \text{Enc}(K, N, A, M) = C.$$

Ciphertext validity requires that for every ciphertext  $C \in \mathcal{C}$  there exists at least one valid decryption context  $(K, N, A) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A}$ ; that is  $\text{Dec}(K, N, A, C) \neq \perp$ . Length uniformity requires that the length of a ciphertext depends only on the length of the message and the length of the associated data.

Finally, for AEAD security, we use the traditional *privacy* and *authenticity* definitions [104, §3].

**Committing authenticated encryption.** A number of prior notions for committing AEAD have been proposed. In Figure 2.3 we provide the CMT-1 and CMT-3 games from

<p>CMT-1(<math>\mathcal{A}</math>):</p> <p><math>((K_1, N_1, A_1), (K_2, N_2, A_2), C) \leftarrow \mathcal{A}</math></p> <p><math>M_1 \leftarrow \text{AEAD.Dec}(K_1, N_1, A_1, C)</math></p> <p><math>M_2 \leftarrow \text{AEAD.Dec}(K_2, N_2, A_2, C)</math></p> <p>// decryption success</p> <p>If <math>M_1 = \perp</math> or <math>M_2 = \perp</math></p> <p>    Return false</p> <p>// commitment condition</p> <p>If <math>K_1 = K_2</math></p> <p>    Return false</p> <p>Return true</p>	<p>CMT-3(<math>\mathcal{A}</math>):</p> <p><math>((K_1, N_1, A_1), (K_2, N_2, A_2), C) \leftarrow \mathcal{A}</math></p> <p><math>M_1 \leftarrow \text{AEAD.Dec}(K_1, N_1, A_1, C)</math></p> <p><math>M_2 \leftarrow \text{AEAD.Dec}(K_2, N_2, A_2, C)</math></p> <p>// decryption success</p> <p>If <math>M_1 = \perp</math> or <math>M_2 = \perp</math></p> <p>    Return false</p> <p>// commitment condition</p> <p>If <math>(K_1, N_1, A_1) = (K_2, N_2, A_2)</math></p> <p>    Return false</p> <p>Return true</p>
---	---

Figure 2.3: The CMT-1 and CMT-3 games from [13], with differences between them are highlighted.

Bellare and Hoang [13]. The FROB game from Farshim, Orlandi, and Rosie [58] adapted to the AEAD setting by Grubbs, Lu, and Ristenpart [63], is the same except that the final highlighted predicate is changed to “ $K_1 = K_2$  or  $N_1 \neq N_2$ ”. The FROB game asks the adversary to produce a ciphertext that decrypts under two different keys with the same nonce. The CMT-1 game is more permissive and removes the condition that the nonce be the same. The CMT-3 game is even more permissive and relaxes the different key condition to different keys, nonces, or associated data. Bellare and Hoang [13] show that CMT-3 implies CMT-1, which implies FROB. We will expand on these definitions with a more general framework next.

## 2.2 Granular Committing Encryption Definitions

We provide a more general framework for defining commitment security for encryption. As motivation, we observe that while the CMT-1 and the stronger CMT-3 notions provide good security goals for constructions, they do not precisely capture the way in which attacks violate security—which parts of the decryption context does the attacker need to control, which parts have been pre-selected by some other party, and which parts are known to the attacker.

These considerations are crucial for determining the exploitability of commitment vulnerabilities in practice. For instance, the vulnerability in Facebook attachment franking [57] exploited by Dodis et al. [51, §3] only works if the nonces are the same; and the key rotation attack described by Albertini et al. [2] only works with keys previously imported to the key management service. And, looking ahead, we propose a variant of the Subscribe with Google attack described by Albertini et al. [2] in which a malicious publisher provides a full decryption context only knowing the honestly published ciphertext.

We provide a more general framework for commitment security notions that more precisely captures attack settings. As we will see in subsequent sections, our definitions provide a clearer explanatory framework for vulnerabilities.

**Committing security framework.** We find it useful to expand the set of security notions to more granularly capture the ways in which the two decryption contexts are selected that generalizes context commitment security. In Figure 2.4 we detail the  $\text{CMT}[\Sigma]$  game, parameterized by a *setting*  $\Sigma = (\text{ts}, S, P)$  that specifies a *target specifier*  $\text{ts}$ , a *context selector*  $S$ , and a *predicate*  $P$  (to be defined next.) The adversary helps compute a ciphertext and two decryption contexts  $(C, (K_1, N_1, A_1), (K_2, N_2, A_2))$ , what we call a *commitment attack instance* (*cat*). The adversary wins if  $C$  decrypts under both decryption contexts, and the two decryption contexts satisfy the predicate  $P$ . The parameterization allows attack settings in terms of which portions of the commitment attack instance are attacker controlled versus chosen in some other way, and which of the latter are revealed to the attacker.

We now provide more details. A commitment attack instance is a tuple  $(C, (K_1, N_1, A_1), (K_2, N_2, A_2))$  consisting of a ciphertext  $C \in \mathcal{C}$ ; two keys  $K_1, K_2 \in \mathcal{K}$ ; two nonces  $N_1, N_2 \in \mathcal{N}$ ; and two associated data  $A_1, A_2 \in \mathcal{A}$ . A target specifier  $\text{ts}$

<div style="border: 1px solid black; padding: 10px;"> <p><math>\text{CMT}[\text{ts}, \text{S}, \text{P}](\mathcal{A})</math>:</p> <p><math>\text{cat}_c \leftarrow \text{S}</math></p> <p><math>\text{cat}_a \leftarrow \mathcal{A}(\text{Reveal}_{\text{ts}}(\text{cat}_c))</math></p> <p><math>\text{cat} \leftarrow \text{Merge}_{\text{ts}}(\text{cat}_c, \text{cat}_a)</math></p> <p>If <math>\text{cat} = \perp</math>:</p> <p style="padding-left: 20px;">Return false</p> <p><math>(C, (K_1, N_1, A_1), (K_2, N_2, A_2)) \leftarrow \text{cat}</math></p> <p><math>M_1 \leftarrow \text{AEAD.Dec}(K_1, N_1, A_1, C)</math></p> <p><math>M_2 \leftarrow \text{AEAD.Dec}(K_2, N_2, A_2, C)</math></p> <p>If <math>M_1 = \perp</math> or <math>M_2 = \perp</math>:</p> <p style="padding-left: 20px;">Return false</p> <p>Return <math>\text{P}((K_1, N_1, A_1), (K_2, N_2, A_2))</math></p> </div>	<table border="1"> <thead> <tr> <th>Notion</th><th>Predicate P</th></tr> </thead> <tbody> <tr> <td><math>\text{CMT}_k</math></td><td><math>(K_1 \neq K_2)</math></td></tr> <tr> <td><math>\text{CMT}_n</math></td><td><math>(N_1 \neq N_2)</math></td></tr> <tr> <td><math>\text{CMT}_a</math></td><td><math>(A_1 \neq A_2)</math></td></tr> <tr> <td><math>\text{CMT}_k^*</math></td><td><math>(K_1 \neq K_2) \wedge (N_1, A_1) = (N_2, A_2)</math></td></tr> <tr> <td><math>\text{CMT}_n^*</math></td><td><math>(N_1 \neq N_2) \wedge (K_1, A_1) = (K_2, A_2)</math></td></tr> <tr> <td><math>\text{CMT}_a^*</math></td><td><math>(A_1 \neq A_2) \wedge (K_1, N_1) = (K_2, N_2)</math></td></tr> </tbody> </table>	Notion	Predicate P	$\text{CMT}_k$	$(K_1 \neq K_2)$	$\text{CMT}_n$	$(N_1 \neq N_2)$	$\text{CMT}_a$	$(A_1 \neq A_2)$	$\text{CMT}_k^*$	$(K_1 \neq K_2) \wedge (N_1, A_1) = (N_2, A_2)$	$\text{CMT}_n^*$	$(N_1 \neq N_2) \wedge (K_1, A_1) = (K_2, A_2)$	$\text{CMT}_a^*$	$(A_1 \neq A_2) \wedge (K_1, N_1) = (K_2, N_2)$
Notion	Predicate P														
$\text{CMT}_k$	$(K_1 \neq K_2)$														
$\text{CMT}_n$	$(N_1 \neq N_2)$														
$\text{CMT}_a$	$(A_1 \neq A_2)$														
$\text{CMT}_k^*$	$(K_1 \neq K_2) \wedge (N_1, A_1) = (N_2, A_2)$														
$\text{CMT}_n^*$	$(N_1 \neq N_2) \wedge (K_1, A_1) = (K_2, A_2)$														
$\text{CMT}_a^*$	$(A_1 \neq A_2) \wedge (K_1, N_1) = (K_2, N_2)$														

Figure 2.4: **(Left)** The  $\text{CMT}[\Sigma]$  commitment security game, parameterized by  $\Sigma = (\text{ts}, \text{S}, \text{P})$ , a target selector  $\text{ts}$ , context selector  $\text{S}$ , and predicate  $\text{P}$ . **(Right)** Predicates for the permissive notions  $\text{CMT}_k$ ,  $\text{CMT}_n$ ,  $\text{CMT}_a$  and restrictive notions  $\text{CMT}_k^*$ ,  $\text{CMT}_n^*$ ,  $\text{CMT}_a^*$ , where  $\text{ts} = \emptyset$ .

is a subset of labels  $\{C, k_1, n_1, a_1, k_2, n_2, a_2\} \times \{\cdot, \hat{\cdot}\}$ . The left set labels the components of a commitment attack instance, called component labels, and the right set denotes whether the specified component is revealed to the adversary (no hat means revealed and hat means not revealed.) For example,  $\text{ts} = \{k_1, \hat{k}_2\}$  indicates the  $K_1$  and  $K_2$  in the context, and that  $K_1$  is revealed to the attacker.

A context selector  $\text{S}$  is a randomized algorithm that takes no input and produces the challenger-defined elements of a commitment attack instance, denoted  $\text{cat}_c$ , as specified by the target specifier  $\text{ts}$ . The reveal function  $\text{Reveal}_{\text{ts}}$  parameterized by  $\text{ts}$ , takes a subset of a commitment attack instance and reveals the components that  $\text{ts}$  tells it to reveal; i.e., the specified components with no hat. The merge function  $\text{Merge}_{\text{ts}}(\text{cat}_c, \text{cat}_a)$  parameterized by the target specifier  $\text{ts}$ , takes two subsets of commitment attack instances  $\text{cat}_c$  (challenger-defined elements) and  $\text{cat}_a$  (adversary-defined elements) and works as follows. First, it checks for every component specified by  $\text{ts}$  that  $\text{cat}_c$  has a corresponding value. Second, it checks that for every component specified by  $\text{ts}$ , if  $\text{cat}_a$  has a value, that it matches the value in  $\text{cat}_c$ . If either of these checks fail, it outputs

⊥. Otherwise, it returns their union  $\text{cat}_c \cup \text{cat}_a$ . Finally, the predicate  $P$  takes two decryption contexts output by  $\text{Merge}_{\text{ts}}(\text{cat}_c, \text{cat}_a)$ , and outputs **true** if they satisfy some criteria (e.g., that  $K_1 \neq K_2$ ), and **false** otherwise.

We associate to a setting  $\Sigma = (\text{ts}, S, P)$ , AEAD  $\Pi$ , and adversary  $\mathcal{A}$  the CMT advantage defined as

$$\text{Adv}_{\Pi}^{\text{CMT}[\Sigma]}(\mathcal{A}) := \Pr [ \text{CMT}[\Sigma](\mathcal{A}) \Rightarrow \text{true} ] .$$

Taking a concrete security approach, we will track the running time used by  $\mathcal{A}$  and provide explicit advantage functions. Adapting our notions to support asymptotic definitions of security is straightforward: in our discussions we will often say a scheme is  $\text{CMT}[\Sigma]$  secure as informal shorthand that no adversary can win the  $\text{CMT}[\Sigma]$  game with “good” probability using “reasonable” running time.

**Capturing CMT-1, CMT-3, and more via predicates** To understand our definitional framework further, we can start by seeing how to instantiate it to coincide with prior notions. Let  $\text{ts} = \emptyset$  indicate the empty target selector, meaning that  $\mathcal{A}$  chooses the ciphertext and two decryption contexts fully. Then the set of  $\Sigma$  settings that use the empty target selector defines a family of security goals, indexed solely by predicates, which we denote by  $\text{CMT}[P]$ . This family includes CMT-1 by setting  $P := (K_1 \neq K_2)$  and CMT-3 by setting  $P := (K_1, N_1, A_1) \neq (K_2, N_2, A_2)$ . Not all instances in this family are interesting: consider, for example, when  $P$  always outputs **true** or **false**. Nevertheless, the flexibility here allows for more granular specification of adversarial ability. For instance, the predicate that requires  $(K_1 \neq K_2) \wedge (N_1 = N_2)$  captures a setting like that of the Dodis et al. [51] attack against Facebook’s message franking, which requires that both decryption contexts have the same nonce.

Three games of particular interest are those with predicates that focus on inequality

of the three individual context components:  $(K_1 \neq K_2)$ ,  $(N_1 \neq N_2)$ , and  $(A_1 \neq A_2)$ . For notational brevity, we let game  $\text{CMT}_k := \text{CMT}[P = (K_1 \neq K_2)]$  and similarly  $\text{CMT}_n := \text{CMT}[P = (N_1 \neq N_2)]$  and  $\text{CMT}_a := \text{CMT}[P = (A_1 \neq A_2)]$ . Then  $\text{CMT}_k$  corresponds to CMT-1, but  $\text{CMT}_n$  and  $\text{CMT}_a$  are new. They are also orthogonal to CMT-1, in the sense that we can give schemes that achieve CMT-1 but not  $\text{CMT}_a$  nor  $\text{CMT}_n$  security (see Theorem 9 in the full version.) All three are, however, implied by being CMT-3 secure, and a scheme that simultaneously meets  $\text{CMT}_k$ ,  $\text{CMT}_n$ , and  $\text{CMT}_a$  also enjoys CMT-3 security (see Lemmas 7 and 8 in the full version.)

Note that  $\text{CMT}_k$ ,  $\text{CMT}_n$ , and  $\text{CMT}_a$  are *permissive*: as long as the relevant component is distinct across the two contexts, it does not matter whether the other components are distinct. Also, of interest are *restrictive* versions; for example, we can consider  $\text{CMT}_k^* := \text{CMT}[(K_1 \neq K_2) \wedge (N_1, A_1) = (N_2, A_2)]$  which requires that the nonces and associated data are the same. Similarly, we can define restrictive notions  $\text{CMT}_n^*$  and  $\text{CMT}_a^*$ . Restrictive versions are useful as they correspond to attacks that have limited control over the decryption context. Interestingly, these restrictive notions are not equivalent to the corresponding permissive notions, nor does a scheme that simultaneously meets  $\text{CMT}_k^*$ ,  $\text{CMT}_n^*$ , and  $\text{CMT}_a^*$  achieve CMT-3 security (see Theorem 10 in the full version.)

**Targeted attacks.** Returning to settings with target specifier  $\text{ts} \neq \emptyset$ , we can further increase the family of notions considered to capture situations where a portion of the context is pre-selected. For instance, in the key rotation example of Albertini et al. [2] mentioned earlier, we would have  $\text{ts} = \{k_1, k_2\}$  and  $S = \{K_1 \leftarrow \mathcal{K}; K_2 \leftarrow \mathcal{K}; \text{Return } (K_1, K_2)\}$  to indicate that the malicious sender has to use the two randomly generated keys.

However, not all targeted attack settings are interesting. For some target specifiers

ts, we can specify a context selector  $S$  such that no adversary can achieve non-zero advantage. In particular, if we have  $ts = \{C, k_1, n_1, a_1\}$  and have  $S$  pick ciphertext  $C$  and context  $(K_1, N_1, A_1)$  such that  $\text{AEAD.Dec}(K_1, N_1, A_1, C)$  returns  $\perp$ , then no adversary can win the game, making the security notion trivial (all schemes achieve it.)

**Hiding target components.** Finally, our game considers target specifiers  $ts$  that indicate that some values chosen by  $S$  should remain hidden from  $\mathcal{A}$ . For example, the Subscribe with Google attack described by Albertini et al. [2] can be reframed as a meddler-in-the-middle attack as follows. A publisher creates premium content  $M_1$  and encrypts it using a context  $(K_1, N_1, A_1)$  to get a ciphertext  $C$ . The ciphertext  $C$  is published, but the context  $(K_1, N_1, A_1)$  is hidden. A malicious third-party, only looking at the ciphertext  $C$ , tries to construct a valid decryption context  $(K_2, N_2, A_2)$  and uses that to sell fake paywall bypasses. We can formalize this setting by having the target specifier  $ts = \{C, \hat{k}_1, \hat{n}_1, \hat{a}_1\}$ , with the context selector  $S$  as

$$\begin{aligned} K_1 &\leftarrow \$\mathcal{K}; N_1 \leftarrow \$\mathcal{N}; A_1 \leftarrow \$\mathcal{A}; M_1 \leftarrow \$\mathcal{M} \\ \text{Return } &(\text{AEAD.Enc}(K_1, N_1, A_1, M_1), K_1, N_1, A_1) \end{aligned}$$

and with  $\text{Reveal}_{ts}(C, K_1, N_1, A_1)$  outputting  $C$ .

**Context discoverability security.** Dodis et al [51, §5] and Albertini et al. [2, §3.3] have pointed out that traditional CMT games are analogous to collision-resistance for hash functions, in the sense that the goal is to find two different *encryption contexts*  $(K_1, N_1, A_1, M_1)$  and  $(K_2, N_2, A_2, M_2)$  such that they produce the same ciphertext  $C$ . Under this lens, CMT with targeting (and no hiding) is like second preimage resistance, and CMT with targeting and hiding is like preimage resistance. But, the analogy to preimage resistance is not perfect, since we are not asking for *any* preimage but rather one that is not the same as the original. Further, this restriction is unnecessary. Going back to the meddler-in-the-middle example above, it suffices for an on-path attacker to



$\text{CDY}[\text{ts}, \text{S}](\mathcal{A})$ : $\text{dat}_c \leftarrow \text{S}$ $\text{dat}_a \leftarrow \mathcal{A}(\text{Reveal}_{\text{ts}}(t))$ $\text{dat} \leftarrow \text{Merge}_{\text{ts}}(\text{dat}_c, \text{dat}_a)$ If $\text{dat} = \perp$ : Return false $(C, (K, N, A)) \leftarrow \text{dat}$ $M \leftarrow \text{AEAD.Dec}(K, N, A, C)$ If $M = \perp$ : Return false Return true	$\text{CDY}[\{\text{k}, \text{n}\}, \text{S}](\mathcal{A})$ : $(C, K, N) \leftarrow \text{S}$ $A \leftarrow \mathcal{A}(C, K, N)$ $M \leftarrow \text{AEAD.Dec}(K, N, A, C)$ If $M = \perp$ : Return false Return true	$\text{CDY}[\{\text{k}, \text{a}\}, \text{S}](\mathcal{A})$ : $(C, K, A) \leftarrow \text{S}$ $N \leftarrow \mathcal{A}(C, K, A)$ $M \leftarrow \text{AEAD.Dec}(K, N, A, C)$ If $M = \perp$ : Return false Return true
---	--	--

Figure 2.5: **(Left)** The  $\text{CDY}[\text{ts}, \text{S}]$  commitment security game, parameterized by a target specifier  $\text{ts}$  and a context selector  $\text{S}$ . **(Middle)** The variant of  $\text{CDY}[\Sigma]$  used in the definition of  $\text{CDY}_a^*$ . **(Right)** The variant of  $\text{CDY}[\Sigma]$  used in the definition of  $\text{CDY}_n^*$ .

produce any valid context. Thus, we find it useful to define a new preimage resistance-inspired notion of commitment security.

In Figure 2.5 we define the game  $\text{CDY}[\text{ts}, \text{S}]$ , parameterized by a setting  $\Sigma = (\text{ts}, \text{S})$  that specifies a target specifier  $\text{ts}$  and a context selector  $\text{S}$ . In more detail, a *discoverability attack instance* ( $\text{dat}$ ) is a ciphertext and a decryption context  $(C, (K, N, A))$ . Here, a target specifier  $\text{ts}$  is a subset of  $\{\text{k}, \text{n}, \text{a}\} \times \{\cdot, \hat{\cdot}\}$  and a context selector  $\text{S}$  is a randomized algorithm that takes no input and produces a ciphertext and the elements of a decryption context specified by the target specifier  $\text{ts}$ . The reveal function  $\text{Reveal}_{\text{ts}}$  and the merge function  $\text{Merge}_{\text{ts}}(\text{dat}_c, \text{dat}_a)$  work similarly to their CMT counterparts. Finally, the goal of the adversary is to produce *one* valid decryption context for the target ciphertext.

We associate to a setting  $\Sigma = (\text{ts}, \text{S})$ , AEAD scheme  $\Pi$ , and adversary  $\mathcal{A}$  the CDY advantage defined as

$$\text{Adv}_{\Pi}^{\text{CDY}[\Sigma]}(\mathcal{A}) = \Pr [\text{CDY}[\Sigma](\mathcal{A}) \Rightarrow \text{true}].$$

**Restricted CDY and its variants.** To more accurately capture attack settings and to

prove relations, we find it useful to define restricted variants of the  $\text{CDY}[\Sigma]$  game. A class of games of particular interest are ones that allow targeting under *any* context selector; we call this class *restricted CDY*. For a target specifier  $\text{ts}$ , let  $\text{CDY}^*[\text{ts}]$  be the game where the adversary is given a ciphertext and elements of a decryption context specified by  $\text{ts}$ , all selected arbitrarily, and needs to produce the remaining elements of a decryption context such that  $\text{AEAD.Dec}(K, N, A, C) \neq \perp$ . Formally, for an AEAD scheme  $\Pi$  and adversary  $\mathcal{A}$ , we define the  $\text{CDY}^*$  advantage as

$$\text{Adv}_{\Pi}^{\text{CDY}^*[\text{ts}]}(\mathcal{A}) = \Pr [\text{ for all } S, \text{CDY}[\text{ts}, S](\mathcal{A}) \Rightarrow \text{true} ] .$$

In addition, we find it useful to define three specific variants of  $\text{CDY}^*$  that allow targeting two-of-three components of a decryption context. Let  $\text{CDY}_a^*$  be the game where the adversary is given an arbitrary ciphertext  $C$ , key  $K$ , and nonce  $N$ , and has to produce associated data  $A$  such that  $\text{AEAD.Dec}(K, N, A, C) \neq \perp$ . Formally, for an AEAD scheme  $\Pi$  and adversary  $\mathcal{A}$ , we define the  $\text{CDY}_a^*$  advantage as

$$\text{Adv}_{\Pi}^{\text{CDY}_a^*}(\mathcal{A}) = \Pr [\text{ for all } S, \text{CDY}[\{k, n\}, S](\mathcal{A}) \Rightarrow \text{true} ] .$$

The  $\text{CDY}_k^*$  and  $\text{CDY}_n^*$  games are defined similarly where the adversary has to produce a valid key and nonce respectively such that decryption succeeds when the remaining inputs to decryption are pre-selected. Formally, for an AEAD scheme  $\Pi$  and adversary  $\mathcal{A}$ , we define the  $\text{CDY}_k^*$  and  $\text{CDY}_n^*$  advantage as

$$\text{Adv}_{\Pi}^{\text{CDY}_k^*}(\mathcal{A}) = \Pr [\text{ for all } S, \text{CDY}[\{n, a\}, S](\mathcal{A}) \Rightarrow \text{true} ] ,$$

$$\text{Adv}_{\Pi}^{\text{CDY}_n^*}(\mathcal{A}) = \Pr [\text{ for all } S, \text{CDY}[\{k, a\}, S](\mathcal{A}) \Rightarrow \text{true} ] .$$

Note that the context selector can only select *valid* ciphertexts, which sidesteps issues with formatting. Without this constraint, a context selector could select a ci-

phertext that has invalid padding for a scheme that requires valid padding, thereby making the notion trivial (all schemes achieve it.)

Furthermore, specific variants like  $\text{CDY}_a^*$  may be trivial even with this constraint. For instance, if the ciphertext embeds the nonce, then one can pick some key  $K$ , some ciphertext  $C$  embedding some nonce  $N_1$ , some other nonce  $N_2$ , then no  $\text{CDY}_a^*$  adversary can pick associated data  $A$  such that  $C$  decrypts correctly under  $(K, N_2, A)$ . However, in the context of this restricted CDY notion, we think this is desired behavior and delegate capturing nuances like this to the unrestricted CDY notion (which can capture this by restricting to context selectors which ensure that the nonce embedded is the same as the nonce provided.)

**With context compression, CMT-3 implies restricted CDY** A  $\text{CDY}[\Sigma]$  attack does not always imply a  $\text{CMT}[\Sigma]$  attack. Consider, for example, the “identity” AEAD that has  $\text{Enc}(K, N, A, M) \Rightarrow K \parallel N \parallel A \parallel M$  which has an immediate  $\text{CDY}[\Sigma]$  attack but is  $\text{CMT}[\Sigma]$  secure since a ciphertext can only be decrypted under one context.<sup>2</sup> However, continuing with the hash function analogy, we wonder if a “compression” assumption could make this implication hold. In Theorem 1 we show this statement for  $\text{CDY}^*[\text{ts} = \emptyset]$  and CMT-3. And note that this generalizes to  $\text{CDY}^*[\text{ts}]$  for any  $\text{ts}$  with an appropriate compression assumption. Notably, it holds for  $\text{CDY}_a^*$  if we assume compression over associated data rather than the full context.

**Theorem 1.** *Fix some AEAD  $\Pi$ . Then for any adversary  $\mathcal{A}$  that wins the  $\text{CDY}^*[\text{ts} = \emptyset]$  game, we can give an adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\Pi}^{\text{CDY}^*[\text{ts}=\emptyset]}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\Pi}^{\text{CMT-3}}(\mathcal{B}) + \text{ProbBadCtx}_{\Pi}, \quad (2.2)$$

where  $\text{ProbBadCtx}_{\Pi}$  is the probability that a random decryption context, when used for

---

<sup>2</sup>While the “identity” AEAD is not secure in the sense of privacy [104, §3], one can construct a secure counterexample by using a wide pseudorandom permutation [29].

$\underline{\mathcal{B}}:$ $K_1 \leftarrow \mathcal{K}; N_1 \leftarrow \mathcal{N}; A_1 \leftarrow \mathcal{A}$ $M_1 \leftarrow \mathcal{M}$ $\text{ctx}_1 \leftarrow (K_1, N_1, A_1)$ $C \leftarrow \Pi.\text{Enc}(K_1, N_1, A_1, M_1)$ $\text{ctx}_2 \leftarrow \mathcal{A}(C)$ If $\text{ctx}_2 = \perp$ : Return $\perp$ $(K_1, N_2, A_2) \leftarrow \text{ctx}_2$ If $(K_1, N_1, A_1) = (K_2, N_2, A_2)$ Return $\perp$ Return $(C, (K_1, N_1, A_1), (K_2, N_2, A_2))$	$\underline{\mathcal{S}}:$ $K_1 \leftarrow \mathcal{K}; N_1 \leftarrow \mathcal{N}; A_1 \leftarrow \mathcal{A}$ $M_1 \leftarrow \mathcal{M}$ $C \leftarrow \Pi.\text{Enc}(K_1, N_1, A_1, M_1)$ Return $C$
---	---

Figure 2.6: Pseudocode for the CMT-3 adversary  $\mathcal{B}$  and CDY\* context selector  $\mathcal{S}$ , used in proof of Theorem 1.

*encrypting a random message, is the only valid decryption context for the resulting ciphertext.*

*Proof.* This proof is adapted from Bellare and Rogaway [21, p.147], where they prove a similar theorem for hash functions. We construct an adversary  $\mathcal{B}$  that randomly samples a context  $(K_1, N_1, A_1)$ , encrypts a random message to get a ciphertext  $C$ , then asks the CDY adversary  $\mathcal{A}$  to produce a decryption context for  $C$  to get  $(K_2, N_2, A_2)$ . This ciphertext generation can be viewed as a valid CDY context selector  $\mathcal{S}$  so  $\mathcal{B}$  wins if the returned context is different from the one it sampled; i.e.,  $(K_1, N_1, A_1) \neq (K_2, N_2, A_2)$ . The pseudocode for  $\mathcal{B}$  and  $\mathcal{S}$  is given in Figure 2.6 and the success probability is analyzed below.

Per the above discussion the advantage of  $\mathcal{B}$  is

$$\text{Adv}_{\Pi}^{\text{CMT-3}}(\mathcal{B}) = \Pr[(\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \neq \text{ctx}_2)], \quad (2.3)$$

where without loss of generality, we are assuming that  $\mathcal{A}$  always produces a valid context or fails and produces  $\perp$ . But, before simplifying this equation, we need to define some terminology. First, let us define the set of valid decryption contexts for a

ciphertext as

$$\Gamma(C) := \{(K, N, A) : (\Pi.\text{Dec}(K, N, A, C) \neq \perp)\}.$$

Now, for a given message  $M$ , let us also define the set of “bad” decryption contexts which when used for encrypting  $M$ , remain the only valid decryption context for the resulting ciphertext

$$\text{BadCtxs}(M) := \{(K, N, A) : |\Gamma(\Pi.\text{Enc}(K, N, A, M))| = 1\}.$$

Finally, let us define the probability that a random decryption context is bad

$$\text{ProbBadCtx}_\Pi := \Pr[(K, N, A) \in \text{BadCtxs}(M)],$$

over the choice  $(K, N, A, M) \leftarrow_s (\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M})$ . Using this notation we can rewrite [Equation 2.3](#), where the probabilities are over the choice  $(K, N, A, M) \leftarrow_s (\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M})$ , as

$$\begin{aligned} \text{Adv}_\Pi^{\text{CMT-3}}(\mathcal{B}) &= \Pr[(\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \neq \text{ctx}_2)] \\ &\geq \Pr[(\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \neq \text{ctx}_2) \wedge (\text{ctx}_1 \notin \text{BadCtxs}(M))]. \end{aligned}$$

Using conditional probability, we can rewrite this term as

$$\Pr[\text{ctx}_1 \neq \text{ctx}_2 \mid (\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \notin \text{BadCtxs}(m))] \cdot \Pr[(\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \notin \text{BadCtxs}(m))].$$

Recall that if  $\text{ctx}_1 \notin \text{BadCtxs}(m)$ , then the adversary must choose one of at least two valid contexts, each of which are equally likely to be  $\text{ctx}_1$  (even conditioned on  $C$ ). Thus the probability that it picks  $\text{ctx}_1$  is at most  $1/2$ , and so

$$\begin{aligned} \text{Adv}_\Pi^{\text{CMT-3}}(\mathcal{B}) &\geq \frac{1}{2} \cdot \Pr[(\mathcal{A}(C) \neq \perp) \wedge (\text{ctx}_1 \notin \text{BadCtxs}(m))] \\ &\geq \frac{1}{2} \cdot (\Pr[\mathcal{A}(C) \neq \perp] - \Pr[\text{ctx}_1 \in \text{BadCtxs}(m)]). \end{aligned}$$

$\underline{\mathcal{B}}:$ $K_1 \leftarrow \mathcal{K}; N_1 \leftarrow \mathcal{N}; A_1 \leftarrow \mathcal{A}$ $M_1 \leftarrow \mathcal{M}$ $C \leftarrow \Pi.\text{Enc}(K_1, N_1, A_1, M_1)$ $K_2 \leftarrow K_1 + 1; N_2 \leftarrow N_1 + 1$ $A_2 \leftarrow \mathcal{A}(C, K_2, N_2)$ If $A_2 = \perp$ Return $\perp$ Return $(C, (K_1, N_1, A_1), (K_2, N_2, A_2))$	$\underline{\mathcal{S}}:$ $K_1 \leftarrow \mathcal{K}; N_1 \leftarrow \mathcal{N}; A_1 \leftarrow \mathcal{A}$ $M_1 \leftarrow \mathcal{M}$ $C \leftarrow \Pi.\text{Enc}(K_1, N_1, A_1, M_1)$ $K_2 \leftarrow K_1 + 1; N_2 \leftarrow N_1 + 1$ Return $(C, K_2, N_2)$
--	---

Figure 2.7: Pseudocode for the CMT-3 adversary  $\mathcal{B}$  and  $\text{CDY}_a^*$  context selector  $\mathcal{S}$ , used in proof of Theorem 2.

Putting it all together, we get that

$$\text{Adv}_{\Pi}^{\text{CMT-3}}(\mathcal{B}) \geq \frac{1}{2} \cdot (\text{Adv}_{\Pi}^{\text{CDY}^*[\text{ts}=\emptyset]}(\mathcal{A}) - \text{ProbBadCtx}_{\Pi}),$$

and finally rearranging gives the desired result.  $\square$

**CMT-3 implies restricted variants of CDY** We now show that if an attack against any of  $\text{CDY}_k^*$ ,  $\text{CDY}_n^*$ , or  $\text{CDY}_a^*$  implies an attack against CMT-3. Theorem 2 shows this for  $\text{CDY}_a^*$ , but it readily generalizes to  $\text{CDY}_k^*$  and  $\text{CDY}_n^*$ .

**Theorem 2.** *Fix some AEAD  $\Pi$  with key space  $|\mathcal{K}| \geq 2$  and nonce space  $|\mathcal{N}| \geq 2$ . Then for any adversary  $\mathcal{A}$  that wins the  $\text{CDY}_a^*$  game, we can give an adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\Pi}^{\text{CDY}_a^*}(\mathcal{A}) = \text{Adv}_{\Pi}^{\text{CMT-3}}(\mathcal{B}),$$

*and the runtime of  $\mathcal{B}$  is that of  $\mathcal{A}$ .*

*Proof.* We prove this by constructing  $\mathcal{B}$  such that it succeeds whenever  $\mathcal{A}$  succeeds. The adversary  $\mathcal{B}$  randomly samples a context  $(K_1, N_1, A_1)$ , encrypts a random message to get a ciphertext  $C$ , selects some other key  $K_2$  and nonce  $N_2$  and asks the  $\text{CDY}_a^*$  adversary  $\mathcal{A}$  to produce an associated data  $A_2$  such that  $(K_2, N_2, A_2)$  can decrypt  $C$ . This ciphertext and partial context construction can be viewed as a valid context selector  $\mathcal{S}$ .

The pseudocode for the adversary  $\mathcal{B}$  and the context selector  $S$  are given in Figure 2.7. And, notice that by construction,  $\mathcal{B}$  wins whenever  $\mathcal{A}$  succeeds.  $\square$

This approach of constructing  $\mathcal{B}$  readily generalizes to  $\text{CDY}_n^*$  and  $\text{CDY}_k^*$ . Further, notice that the  $\mathcal{B}$  constructed in Figure 2.7 wins  $\text{CMT}_k$  and  $\text{CMT}_n$ ; and similar relations hold for adversaries  $\mathcal{B}$  constructed from  $\text{CDY}_n^*$  and  $\text{CDY}_k^*$  adversaries. Corollary 3 captures these implications.

**Corollary 3.** *Fix some AEAD  $\Pi$  with key space  $|\mathcal{K}| \geq 2$ , nonce space  $|\mathcal{N}| \geq 2$ , and associated data space  $|\mathcal{A}| \geq 2$ . Then the following three statements hold. First, for any adversary  $\mathcal{A}_1$  that wins the  $\text{CDY}_a^*$  game, we can give an adversary  $\mathcal{B}_1$  such that*

$$\text{Adv}_{\Pi}^{\text{CDY}_a^*}(\mathcal{A}_1) = \text{Adv}_{\Pi}^{\text{CMT}_k}(\mathcal{B}_1) = \text{Adv}_{\Pi}^{\text{CMT}_n}(\mathcal{B}_1).$$

*Second, for any adversary  $\mathcal{A}_2$  that wins the  $\text{CDY}_n^*$  game, we can give an adversary  $\mathcal{B}_2$  such that*

$$\text{Adv}_{\Pi}^{\text{CDY}_n^*}(\mathcal{A}_2) = \text{Adv}_{\Pi}^{\text{CMT}_k}(\mathcal{B}_2) = \text{Adv}_{\Pi}^{\text{CMT}_a}(\mathcal{B}_2).$$

*Third, for any adversary  $\mathcal{A}_3$  that wins the  $\text{CDY}_k^*$  game, we can give an adversary  $\mathcal{B}_3$  such that*

$$\text{Adv}_{\Pi}^{\text{CDY}_k^*}(\mathcal{A}_3) = \text{Adv}_{\Pi}^{\text{CMT}_n}(\mathcal{B}_3) = \text{Adv}_{\Pi}^{\text{CMT}_a}(\mathcal{B}_3).$$

*And the runtimes of  $\mathcal{B}_1$ ,  $\mathcal{B}_2$ , and  $\mathcal{B}_3$  are that of  $\mathcal{A}_1$ ,  $\mathcal{A}_2$ , and  $\mathcal{A}_3$ , respectively.*

## 2.3 Context Discovery Attacks against AEAD

We show context discovery attacks on many AEAD schemes which delegate their authenticity to a non-preimage resistant MAC. Specifically, we show  $\text{CDY}_a^*$  attacks

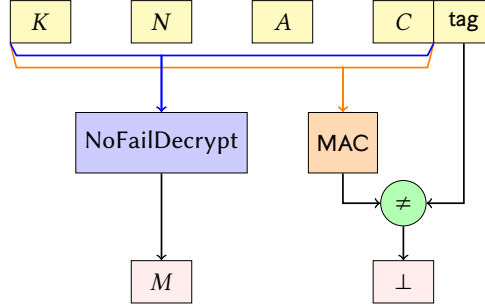


Figure 2.8: Decryption structure of AEAD schemes which delegate their authenticity to a MAC. Should the MAC tag comparison fail, the routine outputs an error ( $\perp$ ), otherwise a message is always output by NoFailDecrypt.

on EAX [24], SIV [107], CCM [54], GCM [53], and OCB3 [80], and  $\text{CDY}_n^*$  attacks on EAX [24] and GCM [53].

We say that an AEAD delegates its authenticity to a MAC if during decryption, a message is output whenever the MAC comparison succeeds. To formalize this, we define NoFailDecrypt as a class of decryption algorithms that never fail. In other words, given a key, nonce, associated data, and ciphertext, they always produce a message. For example, ECB and CTR decryption are NoFailDecrypt algorithms since a valid ciphertext decrypts under any choice of key, nonce, and associated data. On the other hand, CBC with PKCS7 padding is not a NoFailDecrypt algorithm since there most ciphertexts do not decrypt under all choices of key, nonce, and associated data because the decrypted plaintext has incorrect padding.

With this terminology, we say that an AEAD delegates its authenticity to a MAC if it can be written as a combination of a MAC and a NoFailDecrypt algorithm such that if the MAC check fails, decryption fails; if instead the MAC check passes, then decryption outputs the result of NoFailDecrypt (which never fails). This structure is illustrated in Figure 2.8. As a concrete example, for EAX [24] (described in Figure 2.9), the MAC corresponds to checking the OMAC tag, and the NoFailDecrypt corresponds to the



CTR decryption. In this section, we are particularly interested in schemes that compose this structure with a non-preimage resistant MAC like CMAC [55], GMAC [53, §6.4], or OMAC [24, Fig 1].

The  $\text{CDY}_a^*$  attacks we show on these schemes have the following outline. Following the definition of the game, the challenger provides the adversary with a ciphertext  $C \parallel \text{tag}$ , a target key  $K$ , and a target nonce  $N$ , and asks it to find an associated data  $A$  such that  $\text{Decrypt}(K, N, A, C \parallel \text{tag}) \neq \perp$ . Then, the adversary exploits the lack of preimage resistance to find an associated data  $A$  such that  $\text{MAC}(K, N, A, C) = \text{tag}$  and returns  $A$ . Since, in these schemes, the tag check passing guarantees decryption success, we get that decryption succeeds.

For EAX [24] and GCM [53], we also show  $\text{CDY}_n^*$  attacks. They proceed in a similar fashion to the  $\text{CDY}_a^*$  attacks but now the adversary finds a nonce  $N$  such that  $\text{MAC}(K, N, A, C) = \text{tag}$ . But, when the nonce length is shorter than a block (which is always true with GCM, and may be true with EAX), the  $\text{CDY}_n^*$  attacks are slower than the  $\text{CDY}_a^*$  attacks.

The remainder of the section describes the attacks on EAX. The attacks on SIV, CCM, GCM, and OCB3 are in Appendix B of the full version.

**$\text{CDY}_a^*$  and  $\text{CDY}_n^*$  attacks on EAX** We consider EAX over a 128-bit block cipher as defined in Bellare, Rogaway, and Wagner [24]. For simplicity, we restrict to 128-bit tag, 128-bit nonce,<sup>3</sup> and block-aligned messages and associated data. We note however that this is only to make the exposition simpler and is not necessary for the attack. Pseudocode for the scheme with these parameter choices is given in Figure 2.9.

---

<sup>3</sup>EAX [24, Figure 4] supports an arbitrary length nonce; 128 bits is the default in the popular Tink library [4], see [5].

<p><b>OMAC(<math>K, M</math>):</b></p> <p>// Compute Constants  <math>L \leftarrow E_K(0^{128})</math>  <math>B \leftarrow 2 \cdot L</math>  // split into <math>n</math>-bit blocks  // &amp; xor <math>B</math> to the last block  Let <math>M_1, \dots, M_m \leftarrow M</math>  <math>M_m \leftarrow M_m \oplus B</math>  // CBC-MAC Evaluation  <math>C_0 \leftarrow 0^{128}</math>  For <math>i = 1..m</math>:  <math>C_i \leftarrow E_K(C_{i-1} \oplus M_i)</math>  Return <math>C_m</math></p>	<p><b>EAX-Decrypt(<math>K, N, A, C</math>):</b></p> <p>// Separate the Tag  <math>C \parallel \text{tag} \leftarrow C</math>  // Compute and Check Tag  <math>\mathcal{N} \leftarrow \text{OMAC}(K, 0^{128} \parallel N)</math>  <math>\mathcal{H} \leftarrow \text{OMAC}(K, 0^{127} 1 \parallel A)</math>  <math>\mathcal{C} \leftarrow \text{OMAC}(K, 0^{126} 10 \parallel C)</math>  If <math>\text{tag} \neq (\mathcal{N} \oplus \mathcal{H} \oplus \mathcal{C})</math>:  Return <math>\perp</math>  // CTR Decryption  <math>r \leftarrow  C /16</math> // num blocks  For <math>i = 0..(r-1)</math>:  <math>M_i \leftarrow C_i \oplus E_K(\mathcal{N} + i)</math>  Return <math>M</math></p>	<p><b><math>\mathcal{A}(C, K, N)</math>:</b></p> <p><math>C \parallel \text{tag} \leftarrow C</math>  // Compute <math>\xi</math>  <math>\xi \leftarrow \text{tag}</math>  <math>\xi \leftarrow \xi \oplus \text{OMAC}_K(0^{128} \parallel N)</math>  <math>\xi \leftarrow \xi \oplus \text{OMAC}_K(0^{126} 10 \parallel C)</math>  // Reconstruct <math>A</math> and Return  <math>A \leftarrow E_K^{-1}(\xi)</math>  <math>A \leftarrow A \oplus E_K(0^{127} 1) \oplus (2 \cdot E_K(0^{128}))</math>  Return <math>(K, N, A)</math></p>
--	--	---

Figure 2.9: **(Left)** Pseudocode for OMAC [24, Fig 1], used in EAX, with block-aligned inputs. **(Middle)** Pseudocode for EAX Mode [24] decryption with 128-bit tag, 128-bit nonce, and block-aligned messages and associated data. **(Right)** Pseudocode for an  $\text{CDY}_a^*$  attack on EAX.

Let's start by contextualizing the  $\text{CDY}_a^*$  game. The challenger provides us with an  $m$ -block ciphertext  $C = C_1 \dots C_m \parallel \text{tag}$ , a 128-bit target key  $K$ , and a 96-bit target nonce  $N$ . And the goal is to find a 1-block associated data  $A$  such that  $\text{EAX-Decrypt}(K, N, A, C) \neq \perp$ . Notice from Figure 2.9 that decryption passing reduces to the tag check passing. In other words, we can rewrite the goal as finding an associated data  $A$  such that

$$\text{tag} = \text{OMAC}_K(0^{128} \parallel N) \oplus \text{OMAC}_K(0^{126} 10 \parallel C) \oplus \text{OMAC}_K(0^{127} 1 \parallel A). \quad (2.4)$$

We can rearrange terms to get

$$\text{OMAC}_K(0^{127} 1 \parallel A) = \text{tag} \oplus \text{OMAC}_K(0^{128} \parallel N) \oplus \text{OMAC}_K(0^{126} 10 \parallel C).$$

Notice that the right-hand side is composed entirely of known terms, thus we can evaluate it to some constant  $\xi$ . Using the assumption that  $A$  is 1-block, we can expand  $\text{OMAC}_K$  to get

$$E_K(E_K(0^{127} 1) \oplus A \oplus (2 \cdot E_K(0^{128}))) = \xi.$$

Decrypting both sides under  $K$ , and solving for  $A$  gives

$$A = E_K^{-1}(\xi) \oplus E_K(0^{127}1) \oplus (2 \cdot E_K(0^{128})).$$

The full pseudocode for this attack is given in Figure 2.9.

This attack generalizes to other parameter choices. It works as is against an arbitrary-length message, an arbitrary-length tag, and an arbitrary-length nonce. In addition, this attack can also be adapted as a  $\text{CDY}_n^*$  attack. We start by rewriting Equation 2.4 as

$$\text{OMAC}_K(0^{128} \parallel N) = \text{tag} \oplus \text{OMAC}_K(0^{126}10 \parallel C) \oplus \text{OMAC}_K(0^{127}1 \parallel C),$$

and solving for  $N$  as we did for  $A$  above. Since  $N$  is 1 block (128 bits), the reduction is similar, and the success probability remains one. If the nonce length was shorter, then assuming an idealized model like the ideal cipher model, the success probability reduces by a multiplicative factor of  $2^{-f \cdot 128}$  where  $f$  is the fraction of bytes we do not have control over. For example, if we only had control over 14 of the 16 bytes in an encoded block, then the success probability would reduce by  $2^{-16}$ .

This attack can also be adapted to provide partial control over the output plaintext. Notice that the output plaintext is a CTR decryption under the chosen key with the OMAC of the nonce as IV. Assuming an idealized model where the block cipher is an ideal cipher and OMAC is a random function, for every new choice of key and nonce, we get a random output plaintext. So, by trying  $2^m$  key and nonce pairs, we can expect to control  $m$  bits of the output plaintext.

<p><u>SIV-1b-Decrypt(<math>K, C</math>):</u></p> <p><math>c \leftarrow 1^{n-64} 01^{31} 01^{31}</math></p> <p><math>C_1 \parallel \text{tag} \leftarrow C</math></p> <p><math>I \leftarrow \text{tag}</math></p> <p><math>K_1 \parallel K_2 \leftarrow K</math></p> <p>// CTR Decryption</p> <p><math>\text{ctr} \leftarrow I \ \&amp; \ c</math></p> <p><math>M \leftarrow C_1 \oplus E_{K_2}(\text{ctr})</math></p> <p>// IV Check</p> <p><math>I' \leftarrow \text{CMAC}^*(K_1, M)</math></p> <p>If <math>I \neq I'</math>:</p> <p style="padding-left: 20px;">Return <math>\perp</math></p> <p>Return <math>M</math></p>	<p><u>CMAC<math>^*(K, M)</math>:</u></p> <p><math>S \leftarrow \text{CMAC}(K, 0^n)</math></p> <p>Return <math>\text{CMAC}(K, S \oplus M)</math></p> <p><u>CMAC(<math>K, X</math>):</u></p> <p><math>K_s \leftarrow 2 \cdot E_K(0^n)</math></p> <p>Return <math>E_K(K_s \oplus X)</math></p>
--	--

Figure 2.10: **(Left)** Pseudocode for SIV Mode [107] decryption with an  $n$ -bit message and no associated data. **(Right)** Pseudocode for CMAC $^*$  [107] and CMAC [55] with an  $n$ -bit input.

## 2.4 Restrictive Commitment Attacks via k-Sum Problems

The previous section’s  $\text{CDY}_a^*$  and  $\text{CDY}_n^*$  attacks against GCM, EAX, OCB3, SIV, and CCM immediately give rise to *permissive*  $\text{CMT}_k$  attacks against each scheme. This follows from our general result showing that  $\text{CMT}_k$  security implies  $\text{CDY}_a^*$  and  $\text{CDY}_n^*$  (Corollary 3). But this does not imply the ability to build restrictive  $\text{CMT}_k^*$ ,  $\text{CMT}_n^*$ , or  $\text{CMT}_a^*$  attacks that require the non-adversarially controlled parts of the two decryption contexts to be identical (see Theorem 10 in the full version.)

Prior work has provided (in our terminology)  $\text{CMT}_k^*$  attacks for GCM [63, 51], AES-GCM-SIV [111, 83], ChaCha20/Poly1305 [63, 83], XChaCha20/Poly1305 [83], and OCB3 [2]. An open question of practical interest [112] is whether there is a  $\text{CMT}_k^*$  attack against SIV. We resolve this open question, showing an attack that works in time about  $2^{n/3}$ . It requires new techniques related to the fast solution of  $k$ -sum problems, as we explain below.

**Attack on 1-block SIV** We consider SIV over an  $n$ -bit block cipher (for  $n \geq 64$ ) as

defined in the draft NIST specification [107]. For ease of exposition, we restrict to the case of an  $n$ -bit message and no associated data, and describe how to generalize this to the multi-block case in Appendix D in the full version. Pseudocode for the scheme with these parameter choices is given in Figure 2.10.

Here, the  $\text{CMT}_k^*$  adversary seeks to produce a ciphertext  $C = C_1 \parallel \text{tag}$  and two  $2n$ -bit keys  $K = K_1 \parallel K_2$  and  $K' = K'_1 \parallel K'_2$  such that  $\text{SIV-Decrypt}(K, C) \neq \perp$  and  $\text{SIV-Decrypt}(K', C) \neq \perp$ . Notice from Figure 2.10 that this reduces to two simultaneous IV checks passing which can be written as

$$\text{tag} = \text{CMAC}^*(K_1, C_1 \oplus E_{K_2}(\text{tag} \& \text{c})) = \text{CMAC}^*(K'_1, C_1 \oplus E_{K'_2}(\text{tag} \& \text{c}))$$

where  $\text{c} = 1^{n-64}01^{31}01^{31}$  is a constant specified by the SIV standard. Our attack strategy will be to choose  $\text{tag}$  arbitrarily, so we can treat this as a constant value. Towards solving for the remaining variable  $C_1$ , we can substitute in the definition of  $\text{CMAC}^*$  to get

$$\begin{aligned} \text{tag} &= E_{K_1}((2 \cdot E_{K_1}(0^n)) \oplus E_{K_1}(2 \cdot E_{K_1}(0^n)) \oplus C_1 \oplus E_{K_2}(\text{tag} \& \text{c})) \\ &= E_{K'_1}((2 \cdot E_{K'_1}(0^n)) \oplus E_{K'_1}(2 \cdot E_{K'_1}(0^n)) \oplus C_1 \oplus E_{K'_2}(\text{tag} \& \text{c})), \end{aligned}$$

which we can rearrange the two equalities by solving for the variable  $C_1$ , giving us the following:

$$\begin{aligned} C_1 &= E_{K_1}^{-1}(\text{tag}) \oplus (2 \cdot E_{K_1}(0^n)) \oplus E_{K_1}(2 \cdot E_{K_1}(0^n)) \oplus E_{K_2}(\text{tag} \& \text{c}) \\ &= E_{K'_1}^{-1}(\text{tag}) \oplus (2 \cdot E_{K'_1}(0^n)) \oplus E_{K'_1}(2 \cdot E_{K'_1}(0^n)) \oplus E_{K'_2}(\text{tag} \& \text{c}). \end{aligned} \quad (2.5)$$

The above implies that it suffices now to find  $K_1, K_2, K'_1, K'_2$  that satisfy Equation 2.5.

To ease notation, we define four helper functions, one for each term:

$$\begin{aligned}
F_1(K_1) &:= E_{K_1}^{-1}(\text{tag}) \oplus 2 \cdot E_{K_1}(0^n) \oplus E_{K_1}(2 \cdot E_{K_1}(0^n)), \\
F_2(K_2) &:= E_{K_2}(\text{tag} \& c), \\
F_3(K_1) &:= E_{K_1'}^{-1}(\text{tag}) \oplus 2 \cdot E_{K_1'}(0^n) \oplus E_{K_1'}(2 \cdot E_{K_1'}(0^n)), \\
F_4(K_2') &:= E_{K_2'}(\text{tag} \& c),
\end{aligned}$$

and recast [Equation 2.5](#) as a 4-sum problem

$$F_1(K_1) \oplus F_2(K_2) \oplus F_3(K_1') \oplus F_4(K_2') = 0.$$

If these were independent random functions, then we could directly apply Wagner's k-tree algorithm [116] for finding a 4-way collision (also referred to as the generalized birthday problem). But even modeling  $E$  as an ideal cipher, the functions are neither random nor independent. For example,  $F_1(x) = F_3(x)$  always.

Towards resolving this, we first ensure that the keys  $K_1$ ,  $K_2$ ,  $K_1'$ , and  $K_2'$  are domain separated. This can be easily arranged: see [Figure 2.11](#) for the pseudocode of our  $\text{CMT}_k^*$  adversary  $\mathcal{A}$  against SIV. We now turn to lower bounding  $\mathcal{A}$ 's advantage, which consists of two primary steps.

The first is that we argue that, in  $\text{CMT}_k^*$  when running our adversary against SIV, the helper-function outputs are statistically close to uniform. Then, we show that Wagner's approach works for such values.

We observe that  $F_2$  and  $F_4$  trivially behave as independent random functions in the ideal cipher model for  $E$ . The analysis for  $F_1$  and  $F_3$  is more involved. We use the following lemma, which bounds the distinguishing advantage between a uniform  $n$ -bit string and the output of a query to either  $F_1$  or  $F_3$ .

```

 $\mathcal{A}()$ :
 $c \leftarrow 1^{n-64}01^{31}01^{31}$ 
// Arbitrarily pick a tag
 $\text{tag} \leftarrow \{0, 1\}^n \setminus \{0^n\}$ 
// Define helper functions
Def  $F_1(K_1) \leftarrow E_{K_1}^{-1}(\text{tag}) \oplus 2 \cdot E_{K_1}(0^n) \oplus E_{K_1}(2 \cdot E_{K_1}(0^n))$ 
Def  $F_2(K_2) \leftarrow E_{K_2}(\text{tag} \& c)$ 
Def  $F_3(K'_1) \leftarrow E_{K'_1}^{-1}(\text{tag}) \oplus 2 \cdot E_{K'_1}(0^n) \oplus E_{K'_1}(2 \cdot E_{K_1}(0^n))$ 
Def  $F_4(K'_2) \leftarrow E_{K'_2}(\text{tag} \& c)$ 
// Generate lists
For  $i = 1, \dots, q$ :
   $x \leftarrow \text{encode}_{128-2}(i)$ 
  // Domain separate the keys
   $K_1 \leftarrow 00 \parallel x$ ;  $K_2 \leftarrow 01 \parallel x$ ;  $K'_1 \leftarrow 10 \parallel x$ ;  $K'_2 \leftarrow 11 \parallel x$ 
  // Query a row
   $L_1[i] \leftarrow F_1(K_1)$ ;  $L_2[i] \leftarrow F_2(K_2)$ ;  $L_3[i] \leftarrow F_3(K'_1)$ ;  $L_4[i] \leftarrow F_4(K'_2)$ 
// Find an 4-way collision using Wagner's k-tree algorithm [116]
 $\text{res} \leftarrow \mathcal{A}.\text{fourWayCollision}(L_1, L_2, L_3, L_4)$ 
If  $\text{res} = \emptyset$ :
  Return  $\perp$ 
// Repackage the collision into ciphertext and keys
 $(x_1, x_2, x_3, x_4) \leftarrow \text{res}$ 
 $C_1 \leftarrow F_1(x_1) \oplus F_2(x_2)$ 
 $K_1 \leftarrow 00 \parallel x_1$ ;  $K_2 \leftarrow 01 \parallel x_2$ ;  $K'_1 \leftarrow 10 \parallel x_3$ ;  $K'_2 \leftarrow 11 \parallel x_4$ 
Return  $C_1 \parallel \text{tag}, K_1 \parallel K_2, K'_1 \parallel K'_2$ 

```

Figure 2.11: Pseudocode for  $\text{CMT}_k^*$  attack on SIV-1b, where `fourWayCollision` is defined in Appendix C in the full version.

**Lemma 4.** Let  $\text{tag} \in \{0, 1\}^n \setminus \{0^n\}$  and  $\sigma$  be an  $n$ -bit random permutation with inverse  $\sigma^{-1}$  and  $U$  be the uniform random variable over  $n$  bit strings. Define  $n$ -bit random variables (over the choice of  $\sigma$ )

$$A := \sigma^{-1}(\text{tag}), \quad B := 2 \cdot \sigma(0^n), \quad C := \sigma(2 \cdot \sigma(0^n)),$$

where  $\cdot$  denotes multiplication in  $\text{GF}(2^n)$ . Then no adversary that makes one query to a procedure  $P$  can distinguish between  $P \mapsto (U, U, U)$  and  $P \mapsto (A, B, C)$  with probability greater than  $6 \cdot 2^{-n}$ .

The proof proceeds by constructing identical-until-bad games and applying the

*fundamental lemma of game playing* [22] to discern the distinguishing advantage. The proof appears in Appendix C in the full version.

We combine this with the following technical statement about applying Wagner’s  $k$ -tree algorithm [116] to almost-random lists.

**Theorem 5.** *Let  $L$  be a list of  $\ell$  4-tuples  $x = (x_1, x_2, x_3, x_4)$ , where each entry  $x$  is distinguishable from an 4-tuple of independent uniformly random values with probability at most  $\xi$ . Let  $L_1, L_2, L_3$ , and  $L_4$  be lists of 1-index ( $x_1$ ), 2-index ( $x_2$ ), 3-index ( $x_3$ ), and 4-index ( $x_4$ ) elements of  $L$  respectively. Then Wagner’s  $k$ -tree algorithm [116] finds a solution  $(y_1, y_2, y_3, y_4) \in L_1 \times L_2 \times L_3 \times L_4$  such that*

$$y_1 \oplus y_2 \oplus y_3 \oplus y_4 = 0,$$

*with probability at least*

$$(1 - \ell \cdot \xi) \left( 1 - \exp\left(-\frac{\ell^2 \cdot 2^{-n/3}}{8}\right) \right) \left( 1 - \exp\left(1 - \frac{\ell^4 \cdot 2^{-4n/3}}{8} - \frac{2}{\ell^4 \cdot 2^{-4n/3}}\right) \right),$$

*and time at most*

$$20\ell + 4\ell^2 \cdot 2^{-n/3} + 4\text{Sort}(\ell) + 2\text{Sort}((1/2)\ell^2 \cdot 2^{-n/3}),$$

*where  $\text{Sort}(k)$  denotes the time to sort a list of  $k$  items.*

The proof proceeds by analyzing the algorithm step-by-step and at each step applying Chernoff bounds [61] to compute a lower bound on the success probability. The proof appears in Appendix C in the full version.

With Lemma 4 and Theorem 5, we can now prove a lower bound on the advantage of the  $\text{CMT}_k^*$  adversary in Figure 2.11.

**Theorem 6.** *Let  $\mathcal{A}$  be the  $\text{CMT}_k^*$  adversary against SIV over an  $n$ -bit ideal cipher  $E$ ,*



detailed in Figure 2.11. It makes  $10q$  queries to  $E$  and takes at most

$$35q + 4q^2 \cdot 2^{-n/3} + 4\text{Sort}(q) + 2\text{Sort}((1/2)q^2 \cdot 2^{-n/3}) + 11,$$

time, where  $\text{Sort}(k)$  is the cost of sorting a list of  $k$  items. Then the advantage

$$\begin{aligned} \text{Adv}_{\text{SIV}}^{\text{CMT}_k^*}(\mathcal{A}) \geq (1 - 8q \cdot 2^{-n}) & \left( 1 - \exp\left(-\frac{q^2 \cdot 2^{-n/3}}{8}\right) \right) \\ & \left( 1 - \exp\left(1 - \frac{q^4 \cdot 2^{-4n/3}}{8} - \frac{2}{q^4 \cdot 2^{-4n/3}}\right) \right). \end{aligned} \quad (2.6)$$

*Proof.* By construction, the adversary  $\mathcal{A}$  (Figure 2.11) wins whenever it finds a collision, so it suffices to lower bound this probability. First, the domain separation over the keys ensures that the two helper functions never query the ideal cipher with the same key. This, by the properties of the ideal cipher, ensures independence of the outputs. Second,  $F_2$  and  $F_4$  call the ideal cipher only once on a fixed output under a new key each invocation, so their outputs are indistinguishable from an  $n$ -bit uniform random value. Third,  $F_1$  and  $F_3$  call the ideal cipher three times under the same key each invocation. However, applying Lemma 4 gives us that their outputs are distinguishable from an  $n$ -bit uniform random value with probability at most  $6 \cdot 2^{-n}$ . So, by the union bound, a row of outputs  $(F_1(K_1), F_2(K_2), F_3(K'_1), F_4(K'_2))$  is distinguishable from four independent, uniformly random outputs with probability at most  $8 \cdot 2^{-n}$ . Then, Theorem 5 tells us that the function `fourWayCollision` called by  $\mathcal{A}$  finds a collision with probability at least that of Equation 2.6.

It remains to analyze the cost of the adversary  $\mathcal{A}$ . First, it costs 2 operations to initialize `c` and `tag`. Second, since each loop iteration costs 15 operations, the loop costs  $15q$  operations. Third, from Theorem 5, finding a 4-way collision on four lists of size  $q$  using Wagner's  $k$ -tree algorithm [116] costs at most

$$20q + 4q^2 \cdot 2^{-n/3} + 4\text{Sort}(q) + 2\text{Sort}((1/2)q^2 \cdot 2^{-n/3})$$

operations. Fourth, repackaging the collision and returning costs 9 operations. So, the runtime is at most

$$35q + 4q^2 \cdot 2^{-n/3} + 4\text{Sort}(q) + 2\text{Sort}((1/2)q^2 \cdot 2^{-n/3}) + 11.$$

Finally, since each loop iteration makes 10 ideal cipher queries, the algorithm makes  $10q$  queries.  $\square$

In the following corollary, we show that when the adversary makes approximately  $2^{n/3}$  queries, it can win  $\text{CMT}_k^*$  against SIV with high probability, taking time approximately  $2^{n/3}$ .

**Corollary 7.** *Let  $\mathcal{A}$  be the  $\text{CMT}_k^*$  adversary against SIV over an  $n$ -bit ideal cipher  $E$ , detailed in Figure 2.11 with  $q = 10 \cdot 2^{n/3}$ . It makes  $100 \cdot 2^{n/3}$  queries to  $E$  and takes at most*

$$750 \cdot 2^{n/3} + 4\text{Sort}(10 \cdot 2^{n/3}) + 2\text{Sort}(50 \cdot 2^{n/3}) + 11,$$

*time, where  $\text{Sort}(n)$  is the cost of sorting a list of  $n$  items. Then*

$$\text{Adv}_{\text{SIV}}^{\text{CMT}_k^*}(\mathcal{A}) \geq (1 - 80 \cdot 2^{-2n/3}) (1 - \exp(-12.5 \cdot 2^{n/3})) (1 - \exp(-1249)).$$

## 2.5 Related Work

Key commitment for authenticated encryption was introduced in Farshim, Orlandi, and Rosie [58] through full robustness (FROB), which in turn was inspired by key robustness notions in the public key setting by Abdalla, Bellare, and Neven [1] and refined by Farshim et al. [60]. The FROB game asks that a ciphertext only be able to decrypt under a single key. However, the FROB game was defined for randomized authenticated encryption. Grubbs, Lu, and Ristenpart [63] adapted the FROB game to work with as-

sociated data, where they ask that a ciphertext only be able to decrypt under a single key (with no constraints on the associated data.) This notion was further generalized by Bellare and Hoang [13] to the nonce-based setting, with their committing security 1 (CMT-1) definition. The CMT-1 game asks that a ciphertext only be able to decrypt under a single key (with no constraints on the nonce nor the associated data.)

The real-world security implications of key commitment were first highlighted by Dodis et al. [51] where they exploited the lack of key commitment when encrypting attachments in Facebook Messenger’s message franking protocol [57] to send abusive images that cannot be reported. Albertini et al. [2] generalized this attack from images to other file formats and called attention to more settings where lack of key commitment can be exploited to defeat integrity. While both these attacks targeted integrity, Len, Grubbs, and Ristenpart [83] introduced partitioning oracle attacks and showed how to use them for password guessing attacks by exploiting lack of key commitment to obtain large speedups over standard dictionary attacks, endangering confidentiality.

Proposals for constructing key committing ciphers also started in the Farshim, Orlandi, and Rosie paper [58] where they showed that single-key Encrypt-then-MAC, Encrypt-and-MAC, and MAC-then-Encrypt constructions produce key committing ciphers, when the MAC is collision-resistant. Grubbs, Lu, and Ristenpart [63] showed that the Encode-then-Encipher construction [20] was key committing. Dodis et al. [51] proposed a faster compression function-based key committing AEAD construction termed *encryptment*, and also discussed the closely related Duplex construction [30], which is also key committing. Albertini et al. [2] formally analyzed the folklore padding zeroes and key hashing transforms and showed that they produce key committing AEAD at a lower performance cost than prior constructions. Bellare and Hoang [13] constructed key committing variants of GCM and GCM-SIV termed CAU-C1 and CAU-

SIV-C1, and generic transforms UtC and RtC that can be used to turn unique-nonce secure and nonce-reuse secure AEAD schemes respectively into key committing AEAD schemes.

The potential risk of delegating authenticity of an AEAD entirely to a non-collision-resistant MAC is folklore. Farshim, Orlandi, and Rosie [58] who introduced the notion of committing AEAD also cautioned against using non-collision-resistant MACs and CBC-MAC in particular.

On February 7, 2023, NIST announced the selection of the Ascon family for lightweight cryptography standardization [113]. The finalist version of Ascon [49] specifies two AEAD parameter sets ASCON-128 and ASCON-128A. Both parameter sets specify a 128 bit tag, which by the birthday bound, upper bounds the committing security at 64 bits. But, since the underlying algorithm is a variant of the Duplex construction with a 320-bit permutation, and the same specification specifies parameters for a hash function with 128-bit collision resistance, one can specify an AEAD with 128-bit committing security by tweaking parameters.

**Concurrent work.** In independent and concurrent work made public very recently, Chan and Rogaway [38] introduced a new definitional framework for committing AE. Their goal is to capture multiple different types of commitment attacks—what they call *misattributions*, or an adversary being able to construct distinct pairs  $(K, N, A, M)$  and  $(K', N', A', M')$  that both “explain” a single ciphertext  $C$ —in a unified way. Their main definition only captures commitment to an entire  $(K, N, A, M)$  tuple; but in [38, Appendix A], they briefly describe an extension to only require commitments to a subset of the values.

The extended version of their framework is similar to our  $\text{CMT}[\Sigma]$  definition.

While both frameworks aim to capture granular win conditions beyond CMT-3, they are orthogonal. Their framework models the multi-key setting with many randomly chosen unknown-to-the-adversary, known-to-the-adversary, and chosen-by-the-adversary keys. While our  $\text{CMT}[\Sigma]$  captures the distinction between *permissive* and *restrictive* notions, and settings that impose restrictions on the nonce and associated data. We also introduce the notion of context discoverability and describe its relation to  $\text{CMT}[\Sigma]$ .

Chan and Rogaway [38] also independently observed that AEAD with non-preimage resistant MACs are vulnerable to commitment attacks and show attacks on GCM and OCB3 similar to the ones we give in Section 2.3.

## CHAPTER 3

### THE OCH AUTHENTICATED ENCRYPTION SCHEME\*

Recall that in an authenticated encryption with associated data (AEAD) scheme [103], encryption takes a key, nonce, associated data, and message to deterministically return a ciphertext. The classical security requirement is unique-nonce AE (NAE) security. This means privacy of the message, and authenticity of the message and associated data, assuming encryption never reuses a nonce. These were the security and functionality goals for which the most widely used AEAD schemes were designed over two decades ago, including AES-GCM [88] and ChaCha20-Poly1305 [97].

Practitioners and researchers have in the last few years reached consensus that NAE, and the widely used schemes that target it, suffer from a variety of critical limitations. First, NAE does not guarantee what has come to be called key commitment: an adversary can construct a ciphertext so that it is decryptable by multiple maliciously chosen keys. A spate of work [84, 3, 89, 51, 64, 111, 71] highlights how lack of key commitment leads to attacks in important applications. A stronger goal that implies key commitment is context commitment (CMT) [16]: a ciphertext can only be decrypted under a particular key, nonce, and associated data triple. Second, the concrete security of in-use schemes limits the number of bytes that can be encrypted under any given key, forcing complicated key rotation logic when used at the scale of modern workloads [76]. Finally, widely used schemes do not support nonce hiding [19]. Nonce hiding is important when nonces contain sensitive information such as device identifiers, and also fills a gap in classical NAE. (The latter actually fails to hide the message for certain choices of public nonces [19].)

---

\*This chapter is joint work with Mihir Bellare, Viet Tung Hoang, Julia Len, and Thomas Ristenpart. The proceedings version of this paper will appear at CCS 2025 [92], and an earlier version was presented at Real World Crypto 2024 [28]. This chapter is lightly edited from the full version we are preparing for public release.

The academic literature includes techniques for rectifying each of these problems in isolation, but no AEAD schemes have been proposed that achieve *all* of them in a high-performance way.

**Our contributions.** We set out to build AEAD schemes that could replace today’s NAE schemes and protect data for decades to come. To do so we provide a new, clean-slate approach to secure, efficient AEAD schemes and apply it to build OCH (offset codebook with hashing): a permutation-based AEAD. OCH is the first scheme to provide the following combination of features:

- *128-bit NAE security:* OCH’s formal confidentiality and authenticity guarantees scale to even settings where adversaries have unrealistically large resources, having to get close to  $2^{128}$  queries to have a meaningful attack.
- *128-bit CMT security:* OCH is secure against the strongest form of commitment-style attacks, scaling up to adversaries that can perform close to  $2^{128}$  computations.
- *Nonce versatility and hiding:* OCH can be used with random nonces of  $\geq 192$  bits to avoid collisions, or structured nonces such as counters. It can hide nonces.
- *Speed:* OCH is fast, maximally parallelizable, and utilizes CPU pipelines to achieve high throughput. For example, OCH easily achieves  $> 1$  GB/s throughput on a variety of platforms. It is almost as fast as AES-GCM and OCB [79, 108]—the fastest NAE schemes on many platforms—despite those schemes not meeting any of the above security goals.

See Figure 3.1 for a comparison of OCH and prior schemes.

UCH relies on a large number of design ideas, many of which are new to this work and of broader applicability. We summarize UCH’s design at a high level, highlighting novel components and starting from the bottom. We build a tweakable Even-Mansour

Scheme	128-bit NAE	128-bit CMT	192-bit Nonce	Nonce privacy	1GB/s encrypt
AES128-GCM [88, 53]	✗	✗	✗	✗	✓
AES256-GCM [88, 53]	✗	✗	✗	✗	✓
ChaCha20/Poly1305 [98]	✓	✗	✗	✗	✓
CommitKey[GCM256] [66]		✗	✓	✗	✓
XChaCha20/Poly1305 [8]	✓	✗	✓	✗	✓
AES-256-GEM [7]		✗	✓	✗	✓
XAES-256-GCM [115]		✗	✓	✗	✓
DNDK-GCM256 [65]		✗	✓	✗	✓
CAU-C1 [14]		✗	✗	✗	✓
CTX[GCM256] [39]		✓		✗	✓
Ascon-AEAD128 [114]	✗	✗	✗	✗	✗
Aegis-256 [118, 46]		✗	✓	✗	✓
Deoxys-I-256 [73]		✗	✗	✗	✓
TurboSHAKE128 [43]	✓	✓	✓	✗	✗
<i>AreionOCH</i> (§3.4)	✓	✓	✓	✓	✓

Figure 3.1: Comparison of widely used AEAD schemes (top group), transforms of existing schemes (middle), and clean-state designs (bottom). NAE and CMT refer to multi-user nonce-based AE security and context commitment security, respectively. The last column lists whether the scheme can encrypt a billion bytes per second, for 1024 byte messages with 13 bytes of associated data (corresponding to TLS) on an Intel Raptor Lake CPU.

blockcipher [75, 82, 41] from an underlying cryptographic permutation; the construction, however, is not a generic application of prior work but customized to enable fast amortized generation of offsets for use in a new  $\Theta$ CB-like [79] authenticated encryption (AE) that: (1) supports nonce-hiding, (2) omits support for associated data, and (3) does *not* provide authenticity. Instead, it meets a new, weaker-than-NAE security goal for authenticated encryption (without associated data) called NAX. Intuitively, tags need not be unforgeable but just be computationally almost-xor-universal.

We then provide a new generic transform called Xor-then-Hash (XtH) that uses a collision-resistant PRF to convert a (nonce-hiding) NAX-secure AE scheme into a



(nonce-hiding) context-committing AEAD scheme. We build the CR-PRF from a permutation using a Sponge construction [32, 31]. XtH can also be applied to standard AEAD schemes, in which case it can be viewed as a new generic transform for rendering AEAD context committing, and is more efficient than the prior CTY [16] and CTX [39] transforms, particularly for the short associated data strings most often used in practice.

The final component is a specialized mode for handling very short messages. Here, we provide a new, nonce-hiding variant of the synthetic IV mode [106] that can be built from the same tweakable blockcipher as  $\Theta\text{CX}$ .

We provide a formal analysis of all of the above. To make it both tractable and easier to verify, our analysis is highly modular. This also means that our intermediate design components can be more readily used elsewhere with formal analyses already provided.

**Limitations.** Non-goals for OCH are achieving nonce-misuse-resistance [106] or robust AEAD security [69]. These provide improved confidentiality and authenticity in the face of practically relevant threats, such as accidental nonce reuse or accidental release of unverified plaintext. But achieving them requires more than one cryptographic pass over plaintexts, and would bar the ability to achieve performance competitive with one-pass NAE schemes.

**Summary.** We designed the first AEAD scheme that simultaneously provides high-security parameter NAE security, context commitment security, and support for nonce hiding, while performing almost as well as insufficiently secure schemes such as AES-GCM and ChaCha20/Poly1305. Our implementation includes optimizations for various platforms. We will be open-sourcing and making it public before publication of this

paper.

### 3.1 Background and Related Work

We provide a high-level overview of modern design goals, interleaving discussion of related work.

**High security scaling.** We desire schemes that have good (multi-user) security bounds, so that security holds even when used across many instances and for today’s large-scale workloads. For example, AES-based schemes with birthday bound security fail at around  $2^{64}$  AES invocations, which limits scalability. TLS 1.3 does not allow encryption of more than  $2^{24.5}$  records under a single key due to this birthday bound issue [101, §5.5]. See [76] for a more detailed discussion of modern scaling challenges. By providing significantly higher security, we can instead avoid hard usage limits in practice. Thus, we have the following goal:

**128-bit NAE:** *AEAD should achieve 128-bit NAE security, meaning breaking confidentiality and authenticity security would require close to an intractably large  $2^{128}$  adversarial queries.*

In theory, it is not hard to build AEAD that achieves 128-bit NAE. and in fact the academic literature has a surfeit of options for building AEAD schemes. For example, one might build an Encrypt-then-MAC AEAD using Rijndael-256 in counter mode for encryption, together with HMAC-SHA256 for message authentication. But this will be a lot slower than what we achieve. We further review the landscape of approaches for building AEAD in ??.

**Context commitment.** NAE security does *not* guarantee what has come to be called key commitment, as first formalized by [59]: an adversary can construct a ciphertext so that it is decryptable by multiple maliciously chosen keys. An increasingly large body of work shows how lack of key commitment leads to attacks in practical settings including: abuse reporting in encrypted messaging [51, 64], password-based encryption [84], password-based key exchange [84], anonymous public key encryption [84], key rotation schemes [3], and symmetric hybrid (also called envelope) encryption [3]. Subsequent works introduced security goals stronger than key commitment, such as context commitment [14, 89, 39]. It requires that ciphertexts must only be decryptable under a single key, nonce, and associated data triple—what we refer to as a context, following [89]. Attacks have shown that even schemes that achieve key commitment can fail to meet this stronger goal [89]. We refer to context commitment as CMT.

Given all the security issues stemming from lack of commitment, the academic and practitioner consensus (c.f., [76, 40, 28, 27, 84, 100, 17]) is that general-purpose AEAD should achieve at least key commitment and, ideally, context commitment as well.

Prior academic work, as well as deployed implementations, suggest transforms that convert an existing non-CMT AEAD to be CMT [3, 76, 52, 45, 14, 39, 66]. But these transforms don’t address other deficiencies in the underlying AEAD schemes, such as not achieving 128-bit NAE. One could apply known frameworks for building new schemes that are 128-bit NAE (such as Encrypt-then-MAC, as mentioned above), and then apply the transforms. But, again, these won’t be as fast as our approach.

Ascon [48, 114], the winner of a recent lightweight AEAD competition, appears to achieve CMT security, but only against adversaries limited to much less than  $2^{64}$  computations. Other schemes that achieve 128-bit CMT and 128-bit NAE are the recent SHAKE and TurboSHAKE schemes [43] built from a cryptographic hash function.

While elegant and fast on short messages, these schemes are inherently serial and not as fast as ours for larger messages.

**128-bit CMT:** *AEAD should achieve 128-bit CMT security, meaning an adversary must use  $2^{128}$  computations to find a ciphertext that decrypts under two different contexts.*

**Large nonces and nonce hiding.** Finally, we want to support both stateful and randomized nonces, or nonces that are some combination of the two. Here there are two issues. First is that nonces for widely used schemes are too small to allow encrypting, with the same key, a large number of messages with random nonces. For example, AES-GCM’s maximum effective nonce length is 96 bits (longer nonces get hashed to 96 bits), and so one cannot use random nonces to process more than  $2^{32}$  messages before security breaks down. But large scale deployments may encrypt that many messages in a few seconds, forcing more complicated software engineering to work around the limitations of AES-GCM (frequent rekeying, state management for counter-based nonces, etc.).

A second issue is that, as Bellare, Ng, and Tackmann [19] observed, publicly transmitted non-random nonces can be privacy damaging, such as leaking the number of encryptions performed, the identity of devices performing encryptions, or even partial information about messages for some ways of choosing nonces. While some prior work has given schemes that achieve nonce-hiding [19], no schemes widely used in practice provide it. Ideally, a scheme would be versatile, allowing sufficiently large random nonces ( $\geq 192$  bits to avoid collisions) that need not be hidden, stateful or structured nonces that are kept secret, or nonces that combine both. We refer to this as supporting versatile nonces.

**Versatile nonces:** *AEAD should support random nonces of sufficient size ( $\geq 192$  bits),*

and allow hiding part or all of a nonce.

### 3.2 Preliminaries

**Notation.** We let  $\mathbb{N}_0$  be the set of non-negative integers. We refer to elements of  $\{0, 1\}^*$  as *bitstrings*, denote the length of a bitstring  $x$  by  $|x|$  and refer to bitstrings of length  $n$  as  $n$ -bit strings. We denote the empty string with  $\varepsilon$ , the  $n$ -bit all-zero and all-one strings by  $0^n$  and  $1^n$ , the concatenation of two bitstrings  $x, y \in \{0, 1\}^*$  by  $x\|y$ , and the bitwise xor of two  $n$ -bit strings  $x, y \in \{0, 1\}^n$  by  $x \oplus y$ . Given an  $n$ -bit string  $x = x_{n-1} \cdots x_1 x_0 \in \{0, 1\}^n$ , we define its most significant bit  $\text{msb}(x) := x_{n-1}$ , its logical left shift  $(x \ll i) := x_{n-1-i} \cdots x_0 0^i$  where  $i < n$ , and its number representation  $\text{str2num}_n(x) := \sum_{i=0}^{n-1} x_i 2^i$ . We note that  $\text{str2num}_n$  is invertible and denote its inverse  $\text{num2str}_n$ , which maps all numbers  $0 \leq A < 2^n$  to  $n$ -bit strings such that  $\text{str2num}_n(\text{num2str}_n(A)) = A$ . An  $n$ -tuple  $(x_1, \dots, x_n) \in (\{0, 1\}^*)^n$  of bitstrings is a sequence of  $n$  bitstrings, and we denote the set of all bitstring tuples with  $\{0, 1\}^{**}$ .

All the sets we consider in this work are finite. For a finite set  $X$ , we use  $x \leftarrow_s X$  to denote sampling a uniform, random element from  $X$  and assigning it to  $x$ .

We use *code-based games* [23, 95] for our definitions and proofs. A *procedure*  $P$  is a sequence of code-like statements. We assume all the variables in a procedure are initialized, and all the inputs and outputs are in the specified domain and range, respectively. We use the symbol  $\perp$  to denote errors, and assume it is always distinguishable. For a procedure  $P$ , we let  $y \leftarrow_s P^{O_1, \dots}(x_1, \dots)$  denote running  $P$  on inputs  $x_1, \dots$ , with oracle access to  $O_1, \dots$ , and assigning the output to  $y$ . We use  $(P \Rightarrow x)$  to denote the event that procedure  $P$  outputs  $x$ , over the coins of  $P$ .

**Permutations.** A  $w$ -bit permutation is function  $\Pi : \{0, 1\}^w \rightarrow \{0, 1\}^w$ , with an inverse  $\Pi^{-1} : \{0, 1\}^w \rightarrow \{0, 1\}^w$ . Let  $\text{Perm}(\mathcal{M})$  be the set of all permutations on  $\mathcal{M}$ , and let  $\widetilde{\text{Perm}}(\mathcal{T}, \mathcal{M})$  be the set of all permutations on  $\mathcal{M}$  tweaked by  $\mathcal{T}$ , that is, the set of all families of permutations on  $\mathcal{M}$  indexed by  $T \in \mathcal{T}$ .

**Ideal permutation model.** In this paper, we design schemes based on an underlying permutation  $\Pi$ . To analyze them, we use the *ideal permutation model* [6, 74, 44] where we model this permutation as a uniformly random permutation  $\Pi \leftarrow \$ \text{Perms}(\mathcal{M})$  over the message space  $\mathcal{M}$ . In this model, the adversary is additionally given access to this permutation and its inverse with the oracles  $\text{PERM}$  and  $\text{PERMINV}$ , and probabilities are additionally taken over the random choice of the underlying permutation  $\Pi$ . We assume this underlying permutation is public and provide the adversary access to the inverse oracle  $\text{PERMINV}$  even if the construction does not use the inverse.

**Multi-user security.** We consider security in the *multi-user or multi-key setting* [93, 26, 87] where the adversary can query oracles under different keys  $K_i \leftarrow \$ \mathcal{K}$  chosen independently and uniformly at random from the keyspace  $\mathcal{K}$ . and the adversary wins if it compromises the security of *any one* of those keys. This captures settings like TLS, Signal, and server key-rotation where compromising one key is sufficient to cause harm. Following convention, for the remainder of this paper, we will use term “multi-user” but note that it is misleading since users routinely have multiple keys.

**Tweakable blockciphers.** A tweakable blockcipher (TBC) [85] is a function  $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$ , defined over a key space  $\mathcal{K} \subset \{0, 1\}^*$ , tweak space  $\mathcal{T} \subset \{0, 1\}^*$ , and input space  $\mathcal{M} \subset \{0, 1\}^*$ , where for each key  $K \in \mathcal{K}$  and each tweak  $T \in \mathcal{T}$ , the function  $\tilde{E}_K^T(\cdot) := \tilde{E}(K, T, \cdot)$  is a permutation on  $\mathcal{M}$ , and we denote its inverse with  $\tilde{D}_K^T(\cdot) := \tilde{D}(K, T, \cdot)$ . We require keys  $|K| =: k$  and inputs  $|M| =: n$  to be fixed length,

<b>Game <math>\text{TPRP}_E^{\text{real}}(\mathcal{A})</math>:</b> $P \leftarrow \$ \text{Perm}(\mathcal{M})$ $b \leftarrow \$ \mathcal{A}^{\text{TPRP}[\text{Kg}, \text{Enc}, \text{Dec}], P^\pm}$ <b>return</b> $b$ <hr/> <b>TPRPKg():</b> $u \leftarrow u + 1; K_u \leftarrow \$ \mathcal{K}$ <hr/> <b>TPRPENC<math>^{P^\pm}(i, T, X)</math>:</b> <b>return</b> $\tilde{E}_{K_i}^{P^\pm}(T, X)$ <hr/> <b>TPRPDEC<math>^{P^\pm}(i, T, Y)</math>:</b> <b>return</b> $\tilde{D}_{K_i}^{P^\pm}(T, Y)$	<b>Game <math>\text{TPRP}_E^{\text{rand}}(\mathcal{A})</math>:</b> $P \leftarrow \$ \text{Perm}(\mathcal{M})$ $b \leftarrow \$ \mathcal{A}^{\text{TPRP}[\text{Kg}, \text{Enc}, \text{Dec}], P^\pm}$ <b>return</b> $b$ <hr/> <b>TPRPKg():</b> $u \leftarrow u + 1; \tilde{\pi}_u \leftarrow \$ \widetilde{\text{Perm}}(\mathcal{T}, \mathcal{M})$ <hr/> <b>TPRPENC<math>(i, T, X)</math>:</b> <b>return</b> $\tilde{\pi}_i(T, X)$ <hr/> <b>TPRPDEC<math>(i, T, Y)</math>:</b> <b>return</b> $\tilde{\pi}_i^{-1}(T, Y)$
---	--

Figure 3.2: Games  $(\text{TPRP}^{\text{real}}, \text{TPRP}^{\text{rand}})$  for a TBC  $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$ , based on a permutation  $P : \mathcal{M} \rightarrow \mathcal{M}$ . If we set the tweakspace  $\mathcal{T} = \{\varepsilon\}$ , these games reduce to  $(\text{SPRP}^{\text{real}}, \text{SPRP}^{\text{rand}})$  for a blockcipher  $E : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{M}$ .

Game $\text{PRF}_F^{\text{real}}(\mathcal{A})$	Game $\text{PRF}_F^{\text{rand}}(\mathcal{A})$	Game $\text{CR}_F(\mathcal{A})$
$b \leftarrow \$ \mathcal{A}^{\text{Kg}, \text{Func}}$	$b \leftarrow \$ \mathcal{A}^{\text{Kg}, \text{Func}}$	$(\ell, (K_1, X_1), (K_2, X_2)) \leftarrow \$ \mathcal{A}^F$
<b>return</b> $b$	<b>return</b> $b$	$Y_1 \leftarrow F_{K_1}(\ell, X_1)$
<b>Kg()</b>	<b>Kg()</b>	$Y_2 \leftarrow F_{K_2}(\ell, X_2)$
$u \leftarrow u + 1; K_u \leftarrow \$ \mathcal{K}$	$u \leftarrow u + 1$	<b>if</b> $Y_1 = \perp$ <b>or</b> $Y_2 = \perp$
<b>Func<math>(i, \ell, X)</math></b>	<b>Func<math>(i, \ell, X)</math></b>	<b>return false</b>
<b>return</b> $F(K_i, \ell, X)$	<b>if</b> $\text{Tbl}[i, \ell, X] = \perp$	<b>if</b> $(K_1, X_1) = (K_2, X_2)$
	$\text{Tbl}[i, \ell, X] \leftarrow \$ \{0, 1\}^\ell$	<b>return false</b>
	<b>return</b> $\text{Tbl}[i, \ell, X]$	<b>return</b> $(Y_1 = Y_2)$

Figure 3.3: Games  $(\text{PRF}^{\text{real}}, \text{PRF}^{\text{rand}})$  and CR for a function  $F$ .

but allow structured tweaks. We recover traditional blockciphers by setting  $\mathcal{T} = \{\varepsilon\}$ .

In this work, we analyze tweakable blockciphers  $\tilde{E}$  built from a public permutation  $P$ , and consider their security in the *ideal permutation model*, where we model the permutation  $P$  as a perfectly random permutation  $P \leftarrow \$ \text{Perm}(\mathcal{M})$ .

For a TBC  $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$  based on a permutation  $P \leftarrow \$ \text{Perm}(\mathcal{M})$  and an adversary  $\mathcal{A}$  we define multi-user (strong) TPRP security as  $\text{Adv}_E^{\text{tprp}}(\mathcal{A}) := \Pr[\text{TPRP}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{TPRP}^{\text{rand}}(\mathcal{A}) \Rightarrow 1]$ , using a pair of games  $\text{TPRP}^{\text{real}}$  and  $\text{TPRP}^{\text{rand}}$  defined in ??.

**PRFs and CR-PRFs.** A variable-output-length pseudorandom function (PRF) is a function  $F :: \mathcal{K} \times \mathcal{L} \times \mathcal{X} \rightarrow \mathcal{Y}$ , defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , output length space  $\mathcal{L} \subseteq \mathbb{N}$ , input space  $\mathcal{X} \subseteq \{0, 1\}^{**}$ , and output space  $\mathcal{Y} \subseteq \{0, 1\}^*$ . We allow the input space to include tuples, require the lengths of keys  $|K| =: k$  be a public constant, and require the length of outputs match the provided output length  $|F(K, \ell, X)| = \ell$ . To simplify notation, we sometimes use the shorthand  $F_K^\ell(\cdot) := F(K, \ell, \cdot)$ .

We define multi-user PRF security [11], extended to variable output lengths, using a pair of games  $\text{PRF}^{\text{real}}$  and  $\text{PRF}^{\text{rand}}$  defined in Figure 3.3, as

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) := \Pr[\text{PRF}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{PRF}^{\text{rand}}(\mathcal{A}) \Rightarrow 1],$$

and define collision-resistance security, similarly extended to variable output lengths, using the game CR defined in Figure 3.3, as

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) := \Pr[\text{CR}(\mathcal{A}) \Rightarrow \text{true}].$$

We let  $\widetilde{\text{Funcs}}(\mathcal{K}, \mathcal{L}, \mathcal{X}, \mathcal{Y})$  be the set of all be the set of all families of variable-output-length functions  $f : \mathcal{X} \times \mathcal{L} \rightarrow \mathcal{Y}$  satisfying  $|f(\ell, X)| = \ell$  for all  $\ell \in \mathcal{L}$ , indexed by  $K \in \mathcal{K}$ .

**Universal hashing.** Let  $G :: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  be a function, defined over a key space  $\mathcal{K}$ , input space  $\mathcal{X} \subseteq \{0, 1\}^*$ , and output space  $\mathcal{Y} \subseteq \{0, 1\}^*$ , where the length of all inputs  $|X|$  is upper bounded by a constant, and the length of all outputs  $|Z| = \eta$  is a constant  $\eta$ . We will refer to such functions as hash functions and define the shorthand  $G_K(\cdot) := G(K, \cdot)$ . Then, we say that the family of hash functions  $\{G_K(\cdot)\}_{K \in \mathcal{K}}$  is  $\delta$ -uniform if for all input  $X \in \mathcal{X}$  and all outputs  $Y \in \mathcal{Y}$ , it holds that

$$\Pr_{K \leftarrow \mathcal{K}} [G_K(X) = Y] \leq \delta.$$

And, we say that  $\{G_K(\cdot)\}_{K \in \mathcal{K}}$  is  $\epsilon$ -AXU if for all all distinct inputs  $X_1, X_2 \in \mathcal{X}$  and all



outputs  $Y \in \mathcal{Y}$ , it holds that

$$\Pr_{K \leftarrow \mathcal{K}} [\mathbf{G}_K(X_1) \oplus \mathbf{G}_K(X_2) = Y] \leq \epsilon .$$

**Domain separation.** We prefix distinct uses of the same function, including PRFs and AXU hashes, with unique one-byte labels, typeset in monospace like 1 and 2, for domain separation [12].

**AEAD syntax.** We use a special case of the general AE3 syntax [25, 16], namely that of the CAESAR call for submissions [36], dubbed AE5 in [96]. Thus we define an AEAD scheme as a tuple of algorithms  $\text{AEAD} = (\text{Enc}, \text{Dec})$ , defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , public nonce space  $\mathcal{N}_p \subseteq \{0, 1\}^*$ , secret nonce space  $\mathcal{N}_s \subseteq \{0, 1\}^*$ , associated data space  $\mathcal{A} \subseteq \{0, 1\}^*$ , message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , and ciphertext space  $\mathcal{C} \subseteq \{0, 1\}^*$ , where:

1.  $\text{Enc} :: \mathcal{K} \times \mathcal{N}_p \times \mathcal{N}_s \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$  is a deterministic algorithm that takes a 5-tuple of a key  $K$ , public nonce  $N_p$ , secret nonce  $N_s$ , associated data  $A$ , and message  $M$ , and returns a ciphertext  $C$ .
2.  $\text{Dec} :: \mathcal{K} \times \mathcal{N}_p \times \mathcal{A} \times \mathcal{C} \rightarrow (\mathcal{M} \times \mathcal{N}_s) \cup \{\perp\}$  is a deterministic algorithm that takes a 4-tuple of a key  $K$ , public nonce  $N_p$ , associated data  $A$ , and ciphertext  $C$ , and returns a message  $M$  and a secret nonce  $N_s$  or an error  $\perp$ .

We call the non-message inputs to  $\text{Enc}$ —the key, public nonce, secret nonce, and associated data—the *encryption context* and the non-ciphertext inputs to  $\text{Dec}$ —the key, public nonce, and associated data—the *decryption context*.

We impose correctness and *tidiness* [95] requirements on AEAD which requires  $\text{Enc}$  and  $\text{Dec}$  to be inverses of each other. Namely for all  $(K, N_p, N_s, A, M, C) \in \mathcal{K} \times \mathcal{N}_p \times$

$\mathcal{N}_s \times \mathcal{A} \times \mathcal{M} \times \mathcal{C}$ :

$$\text{Enc}(K, N_p, N_s, A, M) = C \implies \text{Dec}(K, N_p, A, C) = (M, N_s),$$

$$\text{Dec}(K, N_p, A, C) = (M, N_s) \neq \perp \implies \text{Enc}(K, N_p, N_s, A, M) = C.$$

We also require that the lengths of keys  $|K| =: k$ , public nonces  $|N_p| =: n_p$ , and secret nonces  $|N_s| =: n_s$  be public constants, and the length of the ciphertext  $|C| =: \text{clen}(|M|)$  to be computable by a deterministic public function  $\text{clen}$  of the length of the message  $|M|$ . And define the *ciphertext overhead* to be the function that on input  $m$  returns the difference  $\text{clen}(m) - m$ .

**Tag-based AEAD syntax.** An AEAD scheme  $\text{AEAD} = (\text{Enc}, \text{Dec})$  is *tag-based* [39, 16] if it specifies a tag length  $\tau$  and its ciphertexts  $C$  can be split into a ciphertext core  $C^*$  and an authentication tag  $T$  via  $C^* \| T \leftarrow C$ , meaning the last  $\tau$  bits of the ciphertext is the tag. Additionally, we ask that there is a partial decryption algorithm  $\text{ParDec}$  in terms of which  $\text{Dec}(K, N_p, A, C)$  is defined by

$$C^* \| T \leftarrow C ; (M, N_s, T_d) \leftarrow \text{ParDec}(K, N_p, A, C^*)$$

If  $T = T_d$  then return  $(M, N_s)$  else return  $\perp$ .

We refer to  $T_d$  as the expected tag.

**AEAD security.** We define this in the multi-user setting following Bellare and Hoang [16]. We use the pair of games  $\text{NAE}^{\text{real}}$  and  $\text{NAE}^{\text{rand}}$  which are defined in Figure 4.1, and then set

$$\text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{A}) := \Pr[\text{NAE}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{NAE}^{\text{rand}}(\mathcal{A}) \Rightarrow 1].$$

To prevent trivial victories, we require adversaries to be *valid*—it may not query

Game $\text{NAE}_{\text{AEAD}}^{\text{real}}(\mathcal{A})$	Game $\text{NAE}_{\text{AEAD}}^{\text{rand}}(\mathcal{A})$
$b \leftarrow \mathcal{A}^{\text{NAEKg, NAEEnc, NAEVer}}$	$b \leftarrow \mathcal{A}^{\text{NAEKg, NAEEnc, NAEVer}}$
<b>return</b> $b$	<b>return</b> $b$
$\text{NAEKg}()$	$\text{NAEKg}()$
$u \leftarrow u + 1; K_u \leftarrow \mathcal{K}$	$u \leftarrow u + 1$
$\text{NAEEnc}(i, N_P, N_S, A, M)$	$\text{NAEEnc}(i, N_P, N_S, A, M)$
<b>if</b> $(i, N_P, N_S) \in \mathcal{Q}_N$ <b>then return</b> $\perp$ $C \leftarrow \text{AEAD.Enc}(K_i, N_P, N_S, A, M)$ $\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_P, A, C)\}$ $\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_P, N_S)\}$ <b>return</b> $C$	<b>if</b> $(i, N_P, N_S) \in \mathcal{Q}_N$ <b>then return</b> $\perp$ $C \leftarrow \{0, 1\}^{\text{clen}( M )}$ $\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_P, A, C)\}$ $\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_P, N_S)\}$ <b>return</b> $C$
$\text{NAEVer}(i, N_P, A, C)$	$\text{NAEVer}(i, N_P, A, C)$
<b>if</b> $(i, N_P, A, C) \in \mathcal{Q}_D$ <b>then return</b> $\perp$ $(M, N_S) \leftarrow \text{AEAD.Dec}(K_i, N_P, A, C)$ <b>if</b> $(M, N_S) \neq \perp$ <b>then return true</b> <b>else return false</b>	<b>if</b> $(i, N_P, A, C) \in \mathcal{Q}_D$ <b>then return</b> $\perp$ <b>return false</b>

Figure 3.4: Games  $(\text{NAE}^{\text{real}}, \text{NAE}^{\text{rand}})$  for an AEAD scheme [16].

$\text{NAEEnc}$  with the same input twice, and it may not query  $\text{NAEVer}$  on an output of  $\text{NAEEnc}$ . The latter condition is enforced with the set  $\mathcal{Q}_D$ . In addition, we require adversaries to be *nonce-respecting*—for a given user  $i$ , it may not query  $\text{NAEEnc}$  with the same nonce  $(N_P, N_S)$ . This is enforced with the set  $\mathcal{Q}_N$ .

We say that an adversary is *orderly* [35, 19, 14, 16] if its  $\text{NAEVer}$  queries are made after all the  $\text{NAEEnc}$  queries, all the  $\text{NAEVer}$  queries are made at once (i.e.,  $\text{NAEVer}$  queries may depend on  $\text{NAEEnc}$  queries but not on other  $\text{NAEVer}$  queries), and it returns **true** if any of the  $\text{NAEVer}$  queries returns **true**. We find it useful to restrict to orderly adversaries and the following lemma of [16] upper bounds the advantage loss of such a restriction.

**Lemma 8** (Lemma 3.2 in [16]). *Let AEAD be an AEAD scheme with constant ciphertext overhead  $\tau$ . Let  $\mathcal{A}$  be an NAE adversary making  $q$  queries, with at most  $q_0$  verification queries, to  $u$  users and at most  $B$  queries per user. Then, we can construct an orderly adversary  $\mathcal{B}$  such that  $\text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{A}) \leq \text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{B}) + \frac{q_0}{2^\tau}$ . Adversary  $\mathcal{B}$  makes at most  $q$*

Game CMT( $\mathcal{A}$ )
$((K_1, N_{P_1}, A_1), (K_2, N_{P_2}, A_2), C) \leftarrow \mathcal{A}$ $(M_1, N_{S_1}) \leftarrow \text{AEAD.Dec}(K_1, N_{P_1}, A_1, C)$ $(M_2, N_{S_2}) \leftarrow \text{AEAD.Dec}(K_2, N_{P_2}, A_2, C)$ <b>if</b> $(M_1, N_{S_1}) = \perp$ <b>or</b> $(M_2, N_{S_2}) = \perp$ <b>then return false</b> <b>return</b> $(K_1, N_{P_1}, A_1) \neq (K_2, N_{P_2}, A_2)$

Figure 3.5: Game CMT for an AEAD scheme [14, 39].

queries, with at most  $q_v$  verification queries, to  $u$  users and at most  $B$  queries per user, and has running time similar to  $\mathcal{A}$ .

**Committing security.** We focus on the strongest commitment notion *context commitment* [14, 39] which requires the ciphertext to commit to the entire decryption context. In Figure 4.2 we define security. The outputs  $(K_1, N_{P_1}, A_1), (K_2, N_{P_2}, A_2)$  of  $\mathcal{A}$  are required to be in  $\mathcal{K} \times \mathcal{N}_{\mathcal{P}} \times \mathcal{A}$ . We define the following advantage measure

$$\text{Adv}_{\text{AEAD}}^{\text{cmt}}(\mathcal{A}) := \Pr[\text{CMT}(\mathcal{A}) \Rightarrow \text{true}].$$

### 3.3 The Xor-then-Hash Transform

In this section, we introduce a new transform for building a context committing AEAD called Xor-then-Hash (XtH). It maps a tag-based AEAD scheme AEAD and a collision-resistant hash function  $F$  to a context committing AEAD scheme  $\text{XtH}[F, \text{AEAD}]$ .

Our starting point is the CTY transform of [16], which is a refinement of the CTX transform of [39]. Consider a (non-context-committing) tag-based AEAD scheme with encryption  $\text{AEAD.Enc}(K, N_P, N_S, A, M)$  outputting a ciphertext-tag pair  $(C, T)$ . The CTY transform works by running  $(C, T) \leftarrow \text{AEAD.Enc}(K, N_P, N_S, \varepsilon, M)$ , where  $\varepsilon$  denotes empty associated data, and then  $T^* \leftarrow H(K, N_P, N_S, A, T)$ . The output is  $(C, T^*)$ .

This was shown to achieve NAE security in the random oracle model assuming an NAE-secure AEAD scheme and to achieve CMT security assuming  $H$  is CR.

Our primary goal is to reduce the number of bits hashed by  $H$  to generate  $T^*$ ; our baseline is CTY, which hashes  $|K| + |N_P| + |N_S| + |A| + |T|$  bits. Omitting any of  $K$ ,  $(N_P, N_S)$ ,  $A$ , or  $T$  is insecure — for example, omitting  $(N_P, N_S)$  might be tempting as  $T$  already depends on the nonce, and so NAE security would still be achieved. But CMT security would not. Another straw approach would be to somehow compress  $(N_P, N_S)$ ,  $A$ ,  $T$  before hashing, for example by applying a universal hash function  $G$  to them before applying  $H$ . Since universal hashes are faster than cryptographic hash functions, this would be a net performance improvement. However, this again would not achieve CMT, as an adversary that knows the key used for  $G$  could find  $A, A'$  such that  $G((N_P, N_S), A, T) = G((N_P, N_S), A', T)$ .

First, we observe that instead of hashing the associated data  $A$  and the tag  $T$ , it suffices to hash their bitwise XOR  $A \oplus T$  and still achieve CMT security. Assume for the moment that  $|A| = |T|$  is fixed. Then we claim that the tag computation  $T^* = H(K, N_P, N_S, A \oplus T)$  achieves CMT. At first glance, this seems incorrect, since an adversary can easily find  $A, A'$  and  $T, T'$  such that  $A \oplus T = A' \oplus T'$ . But actually the adversary does not have the freedom to choose the tags  $T$  or  $T'$  arbitrarily, as tags are a deterministic function of the key, nonce, and ciphertext  $C$ . This observation leads to a simple but perhaps counter-intuitive proof of CMT security assuming  $H$  is collision resistant (across all its inputs, including  $K$ ). Later in this section, we will extend this idea to  $|A| \neq |T|$  by replacing the bitwise XOR with an unequal-length generalization  $\text{ixor}$ . As a result, we have reduced the number of bits we need to hash for CMT security to  $|K| + |N_P| + |N_S| + \max\{|A| + 1, |T|\}$  bits. This means that for short associated data  $|A| < |T|$ , which are popular in practice, we incur no cryptographic overhead from  $A$ .

$\text{XtH}[F, \text{AEAD}].\text{Enc}(K, N_P, N_S, A, M)$	$\text{XtH}[F, \text{AEAD}].\text{Dec}(K, N_P, A, C^* \parallel T^*)$
$L \leftarrow F_K^k(1)$	$L \leftarrow F_K^k(1)$
$C^* \parallel T \leftarrow \text{AEAD}.\text{Enc}(L, N_P, N_S, \varepsilon, M)$	$(M, N_S, T) \leftarrow \text{AEAD}.\text{ParDec}(L, N_P, \varepsilon, C^*)$
$T^* \leftarrow F_K^{2\kappa}(2, N_P, \text{ixor}(A, T))$	$T' \leftarrow F_K^{2\kappa}(2, N_P, \text{ixor}(A, T))$
<b>return</b> $C^* \parallel T^*$	<b>if</b> $T' = T^*$ <b>then return</b> $(M, N_S)$
	<b>else return</b> $\perp$

Figure 3.6: **Definition of  $\text{XtH}[F, \text{AEAD}]$ .** Here,  $\text{AEAD} = (\text{Enc}, \text{ParDec})$  is a tag-based AEAD with key length  $k$ ,  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  is a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$ , and the function  $\text{ixor}$  is defined in eq. (3.1).

Second, we observe that if the inner tag  $T$  has length at least  $\min(\kappa + n_s, 2\kappa)$  for security parameter  $\kappa$ , then we can omit hashing the secret nonce  $N_S$ , and rely on the inner tag  $T$  to commit to  $N_S$ . This length assumption isn't necessary for CMT security but necessary for NAE security. It corresponds to a birthday attack on the secret nonces to find an inner tag collision and thereby distinguish the outer tag from a random one. We describe this attack in more detail in ???. Notice that we never need to extend the length of  $T$  by more than  $n_s$ , so we never hash more bits than individually hashing  $T$  and  $N_S$ ; and if  $n_s > \kappa$ , we hash fewer bits. Combining with the previous observation, we have reduced the number of bits we need to hash for CMT security to  $|K| + |N_P| + \max\{|A| + 1, \min(\kappa + n_s, 2\kappa)\}$  bits.

We also take this opportunity to avoid the random oracle model by using  $H$  to derive a subkey  $L$  for the base AEAD, in addition to the tag  $T^*$ .

**Specification of  $\text{XtH}$ .** Let  $\kappa$  be the security parameter. Let  $\text{AEAD} = (\text{Enc}, \text{ParDec})$  be a tag-based AEAD with key length  $k$  and fixed secret nonce and public nonce lengths, and let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$ . The transform is specified in Figure 3.6.

The function  $\text{ixor}(A, T)$  on the associated data  $A$  and the tag  $T$  is designed to min-

imize output length while allowing recovery of  $A$  given both  $T$  and  $\text{ixor}(A, T)$ . The latter property is crucial for security (see Theorem 9). If  $T$  and  $A$  were promised to be equal length bitstrings, then a simple xor would suffice. However, in practice, the length of  $A$  is not fixed, so we define the following encoding

$$\text{ixor}(A, T) := (A \parallel 10^{\ell-|A|-1}) \oplus (T \parallel 0^{\ell-|T|}), \quad (3.1)$$

where  $\ell = \max(|A| + 1, |T|)$ .

At first glance, hashing  $\text{ixor}(A, T)$  instead of  $A \parallel T$  and not hashing  $N_S$  seems counter-intuitive. To see why XtH provides CMT security, note that (1)  $T^*$  is a commitment of  $\text{ixor}(A, T)$ ,  $N_P$ , and  $K$ , and (2)  $(K, C^* \parallel T^*, N_P, \text{ixor}(A, T))$  is a commitment of  $(N_S, A, M)$  because we can recover the latter from the former. The following result shows that XtH achieves CMT security; we provide the proof in ??.

**Theorem 9.** *Let AEAD be a tag-based AEAD and let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$ . Then for any CMT adversary  $\mathcal{A}$  we can construct a CR adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\text{XtH}[F, \text{AEAD}]}^{\text{cmt}}(\mathcal{A}) \leq \text{Adv}_F^{\text{cr}}(\mathcal{B}),$$

*and  $\mathcal{B}$  emulates  $\mathcal{A}$  and makes two additional calls to each of  $F$  and  $\text{AEAD.ParDec}$ .*

**NAE security of XtH.** We now transition to proving the NAE security of XtH. The traditional approach would be to follow prior work and prove that XtH preserves NAE security, thereby requiring the underlying AEAD scheme also be NAE-secure. However, it turns out that NAE security is unnecessarily strong for XtH. This is because XtH does not reveal the raw authentication tag of the underlying AEAD. Still, XtH reveals authenticated tag collisions for different secret nonces  $N_{S_1} \neq N_{S_2}$  for the same key  $K$  and public nonce  $N_P$ . This then begs the question of what property is sufficient for the

Game $\text{NAX}^{\text{real}}(\mathcal{A})$	Game $\text{NAX}^{\text{rand}}(\mathcal{A})$
$b \leftarrow \mathcal{A}^{\text{Kg,Enc,Ver}}$ <i>// <math>\mathcal{A}</math> must be orderly</i>	$b \leftarrow \mathcal{A}^{\text{Kg,Enc,Ver}}$ <i>// <math>\mathcal{A}</math> must be orderly</i>
<b>return</b> $b$	<b>return</b> $b$
Procedure $\text{NAXKg}()$	Procedure $\text{NAXKg}()$
$u \leftarrow u + 1; K_u \leftarrow \mathcal{K}$	$u \leftarrow u + 1$
Procedure $\text{NAXEnc}(i, N_p, N_s, A, M, X)$	Procedure $\text{NAXEnc}(K, N_p, N_s, A, M, X)$
<b>if</b> $(i, N_p, N_s) \in \mathcal{Q}_N$ <b>then return</b> $\perp$	<b>if</b> $(i, N_p, N_s) \in \mathcal{Q}_N$ <b>then return</b> $\perp$
$C^* \parallel T \leftarrow \text{AEAD.Enc}(K_i, N_p, N_s, A, M)$	$C^* \leftarrow \{0, 1\}^{ M }$
$\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_p, A, C^*, X)\}$	$\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_p, A, C^*, X)\}$
$\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_p, N_s)\}$	$\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_p, N_s)\}$
<b>if</b> $T = \perp$ <b>then return</b> $\perp$	<b>return</b> $C^*, \text{false}$
$\text{win} \leftarrow (T \oplus X \in W_{i, N_p})$	
$W_{i, N_p} \leftarrow W_{i, N_p} \cup \{T \oplus X\}$	
$S_{i, N_p}[N_s] \leftarrow T \oplus X$	
<b>return</b> $C^*, \text{win}$	
Procedure $\text{NAXVer}(i, N_p, N_s^*, A, C^*, X)$	Procedure $\text{NAXVer}(i, N_p, N_s^*, A, C^*, X)$
<b>if</b> $(i, N_p, A, C^*, X) \in \mathcal{Q}_D$ <b>then return</b> $\perp$	<b>if</b> $(i, N_p, A, C^*, X) \in \mathcal{Q}_D$ <b>then return</b> $\perp$
$(M, N_s, T_d) \leftarrow \text{AEAD.ParDec}(K_i, N_p, A, C^*)$	<b>return false</b>
<b>if</b> $(M, N_s, T_d) = \perp$ <b>then return false</b>	
<b>if</b> $S_{i, N_p}[N_s^*] = \perp$ <b>then return false</b>	
<b>if</b> $T_d \oplus X = S_{i, N_p}[N_s^*]$ <b>then return true</b>	
<b>else return false</b>	

Figure 3.7: **NAX security**. Differences to NAE (Figure 4.1) are highlighted.

AEAD used in our transform.

To explore this, we define a new notion that we call NAX security. For schemes without secret nonces, NAX is strictly weaker than NAE. For schemes with secret nonces, it is incomparable. We then show that the XtH transform only requires NAX security of the base AE scheme to provide NAE security.

**The NAX notion.** The NAX notion formalizes the assumptions XtH makes on the underlying AEAD. First, it assumes that the authentication tag has AXU-like (almost XOR universality) security, which is strictly weaker than NAE security; the latter requires random-looking tags. Thus, as we'll see, NAX security can be achieved more efficiently. Second, it assumes that for a given user  $i$  and public nonce  $N_p$ , it is hard to






Figure 3.8: **Illustration of  $\Theta\text{CH}[\tilde{E}, G, F]$ .** (Top) encryption using OCH-core when the secret nonce length  $|N_S| = w$  and the message length  $|M|$  not a multiple of  $n$ . (Bottom) encryption using OCH-tiny when  $|M| + |N_S| < w$ .

find secret nonces  $N_{S_1}$  and  $N_{S_2}$  that result in the same inner tag  $T$ . The latter condition is because XtH does not hash the secret nonce and corresponds to the birthday attack on secret nonces described above.

NAX requires that schemes are tag-based (§3.2). We capture multi-user NAX security using a pair of games  $\text{NAX}^{\text{real}}$  and  $\text{NAX}^{\text{rand}}$ , defined in Figure 3.7. We then define

NAX advantage as

$$\text{Adv}_{\text{AEAD}}^{\text{nax}}(\mathcal{A}) := \Pr[\text{NAX}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{NAX}^{\text{rand}}(\mathcal{A}) \Rightarrow 1].$$

These games are similar to those of the NAE notion (Figure 4.1) with some key differences, highlighted in Figure 3.7. In particular, in each NAXENC query, in addition to the usual  $(i, N_p, N_s, A, M)$ , the adversary must also specify a  $\tau$ -bit offset  $X$ . In the real game, it reveals the ciphertext core  $C^*$  (instead of revealing the full ciphertext  $C^* \parallel T$ ) and a boolean  $\text{win} \leftarrow (T \oplus X \in W_{i, N_p})$ , where  $W_{i, N_p}$  is the set of previously encountered  $T \oplus X$  values. The game also saves  $S_{i, N_p}[N_s] \leftarrow T \oplus X$ . In the random game, the NAXENC oracle simply outputs a uniformly random ciphertext core  $C^*$  and  $\text{win} \leftarrow \text{false}$ . Next, for each NAXVER query, in addition to the usual  $(i, N_p, A)$ , the adversary must now specify a ciphertext core  $C^*$  (instead of the full ciphertext), a secret nonce guess  $N_s^*$ , and an offset  $X$ . In the real game, it returns **true** if and only if  $\text{AEAD.ParDec}$  succeeds in producing a tag  $T_d$  such that  $(T_d \oplus X) = S_{i, N_p}[N_s]$ . In the random game, it always returns **false**.

Like NAE, we require adversaries to be *valid*—that is, not query NAXENC with the same input twice or query NAXVER on an output of NAXENC; and to be *nonce-respecting*—that is, for a given user  $i$ , not query NAXENC with the same nonce  $(N_p, N_s)$  twice. Finally, unlike NAE, we additionally require NAX adversaries to be *orderly*—that is, make all the NAXVER queries at once after all the NAXENC queries—this is a mild restriction since all adversaries can be turned into orderly adversaries with small loss (see Theorem 20), but nevertheless simplifies proofs and weakens NAX relative to NAE.

We emphasize that NAX does not imply NAE, due to the strictly weaker assumption on the tag, and give an explicit scheme that is NAX secure but has a two query NAE

attack in ??.

Proposition 10 below shows that, on the other hand, NAE implies NAX. The term  $q \cdot \min(B, 2^{n_s}) \cdot 2^{-\tau}$  corresponds to the birthday attack on secret nonces described above. We prove this proposition in ??.

**Proposition 10.** *Let AEAD be a tag-based AEAD with tag length  $\tau$  and secret nonce length  $n_s$ . Then for any NAX adversary  $\mathcal{A}$  making  $q$  queries in total and at most  $B$  queries per user, we can construct an NAE adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\text{AEAD}}^{\text{nax}}(\mathcal{A}) \leq \text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{B}) + \frac{q \cdot \min(B, 2^{n_s})}{2^\tau},$$

where adversary  $\mathcal{B}$  has the same query complexity as  $\mathcal{A}$  and similar time complexity.

**XtH promotes NAX to NAE.** In Theorem 11 we show that XtH promotes a NAX secure scheme to a NAE secure scheme, and preserves NAE security. The proof is in ??.

**Theorem 11.** *Let  $\mathcal{K}_{\text{XtH}} := \{0, 1\}^{2\kappa}$  and  $\mathcal{A}_{\text{XtH}}$  be the key space and associated data space for XtH, respectively. Let AEAD be a tag-based AEAD with key length  $k$ , tag length  $\tau$ , public nonce length  $n_p$ , and secret nonce length  $n_s$ . Let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$ . Let  $\mathcal{A}$  be an NAE adversary making  $q$  queries, with  $q_v$  verification queries. Assume  $\mathcal{A}$  queries  $u$  users, making at most  $B$  queries (encryption and verification) per user. Then, we can construct a PRF adversary  $\mathcal{B}$  and an NAX adversary  $\mathcal{C}$  such that*

$$\text{Adv}_{\text{XtH}[F, \text{AEAD}]}^{\text{nae}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prf}}(\mathcal{B}) + \text{Adv}_{\text{AEAD}}^{\text{nax}}(\mathcal{C}) + \frac{2qB}{2^{2\kappa}}.$$

Adversary  $\mathcal{B}$  makes at most  $q$  PRF queries to  $u$  users with at most  $B$  queries per user.

Adversary  $\mathcal{C}$  makes at most  $q$  NAX queries, with at most  $q_v$  verification queries, to  $u$  users

and at most  $B$  queries (encryption and verification) per user. Adversaries  $\mathcal{B}$  and  $\mathcal{C}$  have running time similar to  $\mathcal{A}$ .

### 3.4 The OCH AEAD Scheme

We specify the OCH scheme in two steps. We first specify a TBC-based scheme  $\Theta\text{CH}$ , and then we instantiate it with a permutation-based TBC called OCT.

**Underlying parameters and components.** For security parameter  $\kappa$  choose a block size  $w \geq 2\kappa$ , public nonce length  $n_p$ , and a secret nonce length  $n_s \leq w$ . The nonce space is  $\mathcal{N} = \{0, 1\}^{n_p+n_s}$ . Let  $\text{maxBlocks}$  be an upper bound on the number of blocks that  $\Theta\text{CH}$  supports (later we set this to  $2^{63}$ ).

$\Theta\text{CH}$  uses three underlying components whose instantiations we will discuss later. First is a TBC with a structured tweak space. Let  $\tilde{E} :: \{0, 1\}^{\tilde{k}} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  for some key length  $\tilde{k}$  and with tweak space

$$\mathcal{T}_{\Theta\text{CH}} := \mathcal{N} \times \{0, 1, \dots, \text{maxBlocks}\} \times \{\varepsilon, \star, \$, \star \$\} \quad (3.2)$$

where we can equivalently view  $\{\varepsilon, \star, \$, \star \$\}$  as  $\{0, 1, 2, 3\}$ . When the tweak uses  $\varepsilon$  we omit it from the notation. The second component is an  $\epsilon$ -AXU hash function  $G :: \{0, 1\}^h \times \{0, 1\}^* \rightarrow \{0, 1\}^\tau$ , and the third is a variable-output-length collision-resistant PRF  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  with length space  $\mathcal{L} = \{k, 2\kappa\}$  where  $k := \tilde{k} + h$ . The length space reflects that we will be generating a TBC key and an AXU hash key, as well as tags, using  $F$ .

**The TBC-based scheme  $\Theta\text{CH}$ .** We take a top-down approach in our explanation of the scheme.  $\Theta\text{CH}$  is the composition of two encryption algorithms, one for very

short messages and one for longer messages. Diagrams depicting encryption appear in Figure 3.8 and pseudocode appears in Figure 3.10.

When encrypting very short messages, e.g., messages  $M$  with  $n_s + |M| < n$ , we use a construction  $\Theta\text{CH-tiny}$ . It takes an approach inspired by SIV [106]. First, it generates a context-committing tag  $T^*$  using the AXU hash  $G$  and the CR PRF  $F$ , similar to  $\text{XtH}$ . Then it uses this tag  $T^*$  as a synthetic IV to encrypt the message  $M$  and the secret nonce  $N_s$ . On decryption, it uses the provided tag  $T^*$  to recover the message  $M$  and the secret nonce  $N_s$ , then recomputes and checks the tag. Notice that the tag  $T^*$  non-invertibly depends on the secret nonce so enables full nonce-hiding security, and since  $T^*$  is provided to decryption, we can still decrypt.

For all other situations, we use another construction  $\Theta\text{CH-core}$  which composes the  $\text{XtH}$  transform with a new base AEAD scheme  $\Theta\text{CX}$  that is tailored for the  $\text{XtH}$  transform. We explain it next. Note first, however, that the schemes  $\Theta\text{CH-tiny}$  and  $\Theta\text{CH-core}$  share the same TBC key, but use disjoint tweaks to achieve domain separation.

**The  $\Theta\text{CX}$  construction.** For all but tiny messages,  $\Theta\text{CH}$  utilizes an only-NAX-secure mode called  $\Theta\text{CX}$  to which we apply  $\text{XtH}$ .  $\Theta\text{CX}$  uses the same TBC and also the AXU hash  $G$ . When using secret nonces  $n_s > 0$ , it proceeds in a careful way, applying the TBC using the public tweak  $\tilde{E}_K^{N_0,1}(P_1)$  to an input  $P_1$  that contains the secret nonce and the first  $(w - n_s)$  bits of plaintext (assuming  $n_s < w$ ). Looking ahead to security, this secret nonce handling is a bit delicate as we must ensure that the first block of ciphertext is appropriately “randomized” by both the secret and public nonces. Subsequent blocks of plaintext are then encrypted in an OCB-style mode of operation.

Tag computation deviates from prior OCB-style modes, because we will use  $\text{XtH}$ .

$\Theta\text{CX}[\tilde{E}, G].\text{Enc}((\tilde{K}, H), N_P, N_S, M)$	$\Theta\text{CX}[\tilde{E}, G].\text{ParDec}((\tilde{K}, H), N_P, C^*)$
$P_1, \dots, P_m, P_* \leftarrow N_S \parallel M \quad \#  P_i  = n$ $N_0 \leftarrow N_P \parallel 0^{n_s}$ $C_1 \leftarrow \tilde{E}_{\tilde{K}}^{N_0, 1}(P_1)$ $\text{Chk} \leftarrow P_1[1 : \kappa]$ $N \leftarrow N_P \parallel N_S$ <b>for</b> $i \leftarrow 2..m$ <b>do</b> $\quad C_i \leftarrow \tilde{E}_{\tilde{K}}^{N, i}(P_i)$ $\quad \text{Chk} \leftarrow \text{Chk} \oplus P_i[1 : \kappa]$ $C \leftarrow C_1 \parallel \dots \parallel C_m$ <b>if</b> $P_* \neq \varepsilon$ <b>then</b> $\#  P_*  < n$ $\quad \text{Pad} \leftarrow \tilde{E}_{\tilde{K}}^{N, m, *}(0^n)$ $\quad C_* \leftarrow P_* \oplus \text{Pad}[1 :  P_* ]$ $\quad \text{Chk} \leftarrow (\text{Chk} \parallel 0^{\max(0,  P_*  - \kappa)}) \oplus P_*$ $C \leftarrow C \parallel C_*$ $\text{mlen} \leftarrow \text{encode}_r(m \parallel (P_* \neq \varepsilon))$ $T \leftarrow G_H(6, N_S, \text{Chk}, \text{mlen})$ <b>return</b> $C \parallel T$	$C_1, \dots, C_m, C_* \leftarrow C^* \quad \#  C_i  = n$ $N_0 \leftarrow N_P \parallel 0^{n_s}$ $P_1 \leftarrow \tilde{D}_{\tilde{K}}^{N_0, 1}(C_1); N_S \leftarrow P_1[1 : n_s]$ $\text{Chk} \leftarrow P_1[1 : \kappa]$ $N \leftarrow N_P \parallel N_S$ <b>for</b> $i \leftarrow 2..m$ <b>do</b> $\quad P_i \leftarrow \tilde{D}_{\tilde{K}}^{N, i}(C_i)$ $\quad \text{Chk} \leftarrow \text{Chk} \oplus P_i[1 : \kappa]$ $M \leftarrow P_1[n_s : n] \parallel \dots \parallel P_m$ <b>if</b> $C_* \neq \varepsilon$ <b>then</b> $\#  C_*  < n$ $\quad \text{Pad} \leftarrow \tilde{E}_{\tilde{K}}^{N, m, *}(0^n)$ $\quad M_* \leftarrow C_* \oplus \text{Pad}[1 :  C_* ]$ $\quad \text{Chk} \leftarrow (\text{Chk} \parallel 0^{\max(0,  M_*  - \kappa)}) \oplus M_*$ $M \leftarrow M \parallel M_*$ $\text{mlen} \leftarrow \text{encode}_r(m \parallel (C_* \neq \varepsilon))$ $T \leftarrow G_H(6, N_S, \text{Chk}, \text{mlen})$ <b>return</b> $(M, N_S, T)$

Figure 3.9: **Pseudocode for  $\Theta\text{CX}[\tilde{E}, G]$ .**

First, we use the “half-checksum method” of Inoue and Minematsu [70, §4] to reduce the size of the checksum to  $\kappa$  bits for full blocks, except for a final partial block which still needs a full checksum. Then, instead of encrypting this checksum with a TBC tweaked by the nonce, the block length, and whether there was a partial block, we AXU hash the checksum and secret nonce. We omit the public nonce from the AXU hash since we don’t need it for NAX (since Xth commits to it.) Notice that all the inputs to the AXU hash have constant length (in the message size.) Also,  $\Theta\text{CX}$  is optimal in terms of the number of TBC calls needed to encrypt the secret nonce and the message.

Why include the secret nonce in the AXU hash if the checksum already includes it? Otherwise, an adversary could force a checksum collision by querying  $(N_S, M) = (0, 1)$  and  $(N_S', M') = (1, 0)$  and, thereby, produce a tag collision across distinct secret nonces  $N_S \neq N_S'$ , breaking NAX security. Since we must hash  $N_S$  separately, it is tempting to omit adding its first  $\kappa$  bits into the checksum. But that only would work if  $\kappa > n_s$ ; we avoid having to handle different cases here by including it.

$\Theta\text{CH}[\tilde{E}, G, F].\text{Enc}(K, N_P, N_S, A, M)$	$\Theta\text{CH}[\tilde{E}, G, F].\text{Dec}(K, N_P, C^* \parallel T^*)$
$(\tilde{K}, H) \leftarrow F_K^k(3)$ <b>if</b> $ N_S  +  M  < n$ <b>then</b> $\text{// } \Theta\text{CH-tiny}$ $P \leftarrow N_S \parallel M \text{ } \text{// }  P  < n$ $T \leftarrow G_H(4, P)$ $T^* \leftarrow F_K^{2\kappa}(5, N_P, \text{ixor}(A, T))$ $N_0 \leftarrow N_P \parallel 0^{n_s}$ $C^* \leftarrow \tilde{E}_{\tilde{K}}^{N_0, 1, \mathcal{S}}(T^* \parallel 0^{n-2\kappa})[1 :  P ] \oplus P$ <b>return</b> $C^* \parallel T^*$ <b>endif</b> $\text{// } \Theta\text{CH-core} = \text{XtH}[F, \Theta\text{CX}[\tilde{E}, G]].\text{Enc}$ $C^* \parallel T \leftarrow \Theta\text{CX}[\tilde{E}, G].\text{Enc}((\tilde{K}, H), N_P, N_S, M)$ $T^* \leftarrow F_K^{2\kappa}(2, N_P, \text{ixor}(A, T))$ <b>return</b> $C^* \parallel T^*$	$(\tilde{K}, H) \leftarrow F_K^k(3)$ <b>if</b> $ C^*  < n$ <b>then</b> $\text{// } \Theta\text{CH-tiny}$ $N_0 \leftarrow N_P \parallel 0^{n_s}$ $P \leftarrow \tilde{E}_{\tilde{K}}^{N_0, 1, \mathcal{S}}(T^* \parallel 0^{n-2\kappa})[1 :  C^* ] \oplus C^*$ $N_S \leftarrow P[1 : n_s]; M \leftarrow P[n_s :  P ]$ $T \leftarrow G_H(4, P)$ $T_d^* \leftarrow F_K^{2\kappa}(5, N_P, \text{ixor}(A, T))$ <b>if</b> $T_d^* = T^*$ <b>then return</b> $(M, N_S)$ <b>else return</b> $\perp$ <b>endif</b> $\text{// } \Theta\text{CH-core} = \text{XtH}[F, \Theta\text{CX}[\tilde{E}, G]].\text{Dec}$ $(M, N_S, T) \leftarrow \Theta\text{CX}[\tilde{E}, G].\text{ParDec}((\tilde{K}, H), N_P, C^*)$ $T_d^* \leftarrow F_K^{2\kappa}(2, N_P, \text{ixor}(A, T))$ <b>if</b> $T_d^* = T^*$ <b>then return</b> $(M, N_S)$ <b>else return</b> $\perp$

Figure 3.10: **Pseudocode for  $\Theta\text{CH}[\tilde{E}, G, F]$ .**

(sanketh: OCT SPEC START)

**OCT: an optimized TBC.** We now turn to the underlying TBC. While one could use any TBC that supports the tweakspace  $\mathcal{T}_{\text{ECH}}$  and security levels, our scheme’s OCB-like [108, 81] structure has most tweaks that are “increments” of prior tweaks. In addition, we would like to support precomputation given just the secret key (since in many applications keys are used across many messages) and have quick startup for the special case of nonces that are counters (since they are widespread). We therefore specify a custom permutation-based TBC called OCT.

OCT is at its core a tweakable Even-Mansour cipher [41], though we use tweaks that are less than  $w$  bits as also used in [6]. Let  $\Pi : \{0, 1\}^w \rightarrow \{0, 1\}^w$  be a permutation with width  $w \geq 2\kappa$ . Then  $\text{OCT}[\Pi] :: \{0, 1\}^{2\kappa} \times \mathcal{T}_{\text{ECH}} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$  is defined by

(sanketh: Edit this)

$$\text{OCT}[\Pi](\tilde{K}, (N, i, b), X) := \Pi(\Delta \oplus X) \oplus \Delta$$

for  $\Delta \leftarrow H_{\tilde{K}}(N, i, b)$  and where  $H :: \{0, 1\}^{2\kappa} \times \mathcal{T} \rightarrow \{0, 1\}^w$ . While any AXU hash  $H$  would suffice, we build a custom one called  $\text{OH}[\Pi]$  that allows high performance. Its pseudocode appears in Figure 3.11. Looking ahead to the next section, one complexity will be analyzing security of OCT given that we use an  $H$  that is also built from  $\Pi$ .

This construction borrows some techniques from OCB [81]. It uses the finite field  $\text{GF}(2^{2\kappa})$  with the default minimal weight irreducible polynomial. It derives a subkey  $L$  by applying  $\Pi$  to  $\tilde{K}$  (the notation  $0^*$  means sufficiently many bits to get  $n$ ) and taking the first  $2\kappa$  bits from the result. Then from  $i$  and  $b$  in the tweak, we use a  $2\kappa$ -bit Gray code  $\gamma$  to generate an offset  $\lambda_i^b \leftarrow 4\gamma_i + b$ . Multiplication and addition are in  $\text{GF}(2^{2\kappa})$  and we interpret  $b \in \{0, 1, 2, 3\}$  as the numerical value associated to each symbol. Since



OH[Π]( $\tilde{K}, (N, i, b)$ )
$L \leftarrow \Pi(\tilde{K} \parallel 0^*)[1 : 2\kappa]$
$\text{Top} \leftarrow N[1 : \ell - 6]$
$\text{Bottom} \leftarrow \text{str2num}_6(N[\ell - 5 : \ell])$
$\text{KTop} \leftarrow \Pi((\tilde{K} \oplus \text{Top}) \parallel 0^*) \oplus \tilde{K}$
$R_N \leftarrow (\text{KTop} \ll \text{Bottom})[1 : 2\kappa]$
$\lambda_i^b \leftarrow 4\gamma_i + b$
<b>return</b> $((\lambda_i^b \cdot L) \oplus R_N) \parallel 0^{n-2\kappa}$

Figure 3.11: **Definition of the AXU hash OH[Π] used within OCT[Π].** Here  $\gamma_i$  is the standard Gray-code encoding of  $i$ . (**sanketh:** remove this figure and replace with the one from appendix)

Gray code values are bounded  $0 \leq \gamma_i \leq 2i$  and  $\text{maxBlocks} \leq 2^{\kappa-3}$ , all these offsets

$$\Lambda = \{\lambda_i^b := 4\gamma_i + b : 1 \leq i \leq 2^{\kappa-3}, 0 \leq b \leq 3\}$$

are unique and lie between  $1 \leq \lambda_i^b < 2^\kappa$ .

(**sanketh:** Edit) Then we build a *Stretch-then-Shift* [81] inner hash (described in the next paragraph), keyed with the TBC key  $\tilde{K}$  to map the nonce  $N$  to an inner offset  $R_N$ . Finally, we multiply the first inner offset  $\lambda_i^b$  with the subkey  $L$  (in  $\text{GF}(2^{2\kappa})$ ), xor the result with the second inner offset  $R_N$ , and zero-pad everything to from  $2\kappa$  bits to  $n$  bits to get the offset  $\Delta \leftarrow ((\lambda_i^b \cdot L) \oplus R_N) \parallel 0^{n-2\kappa}$ .

The *Stretch-then-Shift* [81] construction assumes that the nonce length  $\ell > 6$ . First, we take the top  $(\ell - 6)$  bits of  $N$ , call it  $\text{Top}$  and encipher it using Even-Mansour [56]:  $\text{KTop} \leftarrow \Pi(\tilde{K} \oplus (\text{Top} \parallel 0^{n-\ell+6})) \oplus \tilde{K}$ . Then, we take the bottom 6 bits of the nonce, call it  $\text{Bottom}$  and use it to shift  $\text{KTop}$  to get

$$R_N \leftarrow (\text{KTop} \ll \text{Bottom})[1 : 2\kappa].$$

Note that  $\text{KTop}$  is the most expensive part of this computation since it requires a permutation call. But since  $\text{KTop}$  only depends on the top  $(\ell - 6)$  bits of  $N$ , it can be

cached across messages. Specifically, if we use counter nonces, then we only need to compute KTop once every 64 successive encryptions.

**Fast incremental computation.** We specifically designed OCT to have fast incremental computation. Suppose we precomputed two offsets  $L_*$  and  $L_\$$ , and a table L of  $Lsize \leftarrow \log_2(\text{maxBlocks})$  offsets, all by doubling the subkey  $L$ . Then, given the current offset  $\Delta_{(N,i)}$  we can compute any of its three possible increments  $\Delta_{(N,i,*)}$ ,  $\Delta_{(N,i,\$)}$ , and  $\Delta_{(N,i+1)}$  by xor-ing the current offset with one of these precomputed values. Note that this precomputation only depends on the TBC key  $\tilde{K}$ , and doubling in  $\text{GF}(2^{2\kappa})$  can be computed with shifts and xor-s. We discuss this further in ??.

(sanketh: OCT SPEC END)

### 3.5 Security Analysis of OCH

We now turn to analyzing the security of OCH. We first analyze context commitment security and then NAE security.

**CMT Security of the TBC-based scheme.** For non-tiny messages,  $\Theta\text{CH}$  is context committing assuming  $F$  is CR—this is a corollary of the commitment analysis of  $X\text{tH}$  (Theorem 9 in Section 3.3). For tiny messages, context commitment stems from the use of  $F$  to commit to the key and public nonce, which in turn fixes the secret nonce and plaintext to a single value.

**Theorem 12.** *Let  $\tilde{E} :: \{0, 1\}^{\tilde{k}} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be an  $n$ -bit TBC with tweak space  $\mathcal{T} = \{\mathcal{N} \times \mathbb{Z} \times \{\varepsilon, \star, \$\}\}$ . Let  $G :: \{0, 1\}^h \times \{0, 1\}^{**} \rightarrow \{0, 1\}^\tau$  be an  $\epsilon$ -AXU hash function. Let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$  where  $k = \tilde{k} + h$ . Then for any CMT adversary  $\mathcal{A}$  we can construct a CR adversary  $\mathcal{B}$  such that*

$$\text{Adv}_{\Theta\text{CH}[\tilde{E}, G, F]}^{\text{cmt}}(\mathcal{A}) \leq \text{Adv}_F^{\text{cr}}(\mathcal{B}),$$

*and  $\mathcal{B}$  emulates  $\mathcal{A}$  and makes up to two additional calls to  $F$ ,  $\tilde{E}$ ,  $G$ , and  $\Theta\text{CX}[\tilde{E}, G].\text{ParDec}$  each.*

*Proof.* We construct  $\mathcal{B}$  such that it wins the CR game whenever  $\mathcal{A}$  wins the CMT game.  $\mathcal{B}$  runs  $\mathcal{A}$  and then uses  $F$ ,  $\tilde{E}$ , and  $\Theta\text{CX}[\tilde{E}, G].\text{ParDec}$  on the output to reconstruct the inputs to the underlying PRF  $F$ . For long messages, it emulates the CR adversary from Theorem 9. We define  $\mathcal{B}$  formally in Figure 3.12 and analyze its success probability next. Let's assume without loss of generality that  $\mathcal{A}$  either wins and produces a valid output or loses and returns  $\perp$ . Then, we need to prove that  $\mathcal{B}$  wins whenever  $\mathcal{A}$  wins and produces a valid output.

<p><b>Adversary <math>\mathcal{B}^F</math></b></p> <hr/> <p><math>X \leftarrow \mathcal{A}</math>  <b>if</b> <math>X = \perp</math> <b>then return</b> <math>\perp</math>  <math>((K_1, N_{P_1}, A_1), (K_2, N_{P_2}, A_2), C^* \parallel T^*) \leftarrow X</math>  <math>(\tilde{K}_1, H_1) \leftarrow F_{K_1}^k(3); (\tilde{K}_2, H_2) \leftarrow F_{K_2}^k(3)</math>  <b>if</b> <math> C^*  \geq n</math> <b>then</b>              <math>(M_1, N_{S_1}, T_1) \leftarrow \Theta\text{CX}[\tilde{E}, G].\text{ParDec}((\tilde{K}_1, H_1), N_{P_1}, \varepsilon, C^*)</math>              <math>(M_2, N_{S_2}, T_2) \leftarrow \Theta\text{CX}[\tilde{E}, G].\text{ParDec}((\tilde{K}_2, H_2), N_{P_2}, \varepsilon, C^*)</math>              <b>return</b> <math>(2\kappa, (K_1, 2, N_{P_1}, \text{ixor}(A_1, T_1)), (K_2, 2, N_{P_2}, \text{ixor}(A_2, T_2)))</math>  <b>else</b>              <math>N_{01} \leftarrow N_{P_1} \parallel 0^{n_s}; N_{02} \leftarrow N_{P_2} \parallel 0^{n_s}</math>              <math>P_1 \leftarrow \tilde{E}_{\tilde{K}_1}^{N_{01}, 1, \\$}(T^* \parallel 0^{n-2\kappa})[1 :  C^* ] \oplus C^*; P_2 \leftarrow \tilde{E}_{\tilde{K}_2}^{N_{02}, 1, \\$}(T^* \parallel 0^{n-2\kappa})[1 :  C^* ] \oplus C^*</math>              <math>N_{S_1} \leftarrow P_1[1 : n_s]; N_{S_2} \leftarrow P_2[1 : n_s]</math>              <math>T_1 \leftarrow G_{H_1}(5, P_1); T_2 \leftarrow G_{H_2}(5, P_2)</math>              <b>return</b> <math>(2\kappa, (K_1, 5, N_{P_1}, \text{ixor}(A_1, T_1)), (K_2, 5, N_{P_2}, \text{ixor}(A_2, T_2)))</math></p>
---

Figure 3.12: Definition of CR adversary  $\mathcal{B}$ , used in proof of Theorem 12.

Assume towards contradiction that  $\mathcal{A}$  wins and produces a valid output with ciphertext core  $C^*$ , but  $\mathcal{B}$  loses. If  $|C^*| \geq n$ , then  $\mathcal{A}$  wins the CMT game against XtH and the statement follows from Theorem 9 (since  $\mathcal{B}$  emulates ?? in this case.) Now, suppose  $|C^*| < n$ . Then, by definition, for  $\mathcal{B}$  to lose, either the outputs don't collide or the inputs are equal. First, if the outputs don't collide

$$F_{K_1}^{2\kappa}(3, N_{P_1}, \text{ixor}(A_1, T_1)) \neq F_{K_2}^{2\kappa}(3, N_{P_2}, \text{ixor}(A_2, T_2)),$$

then the contexts produce different outer tags  $T^*$  implying that one of the decryptions fails. Thus,  $\mathcal{A}$  could not have won the CMT game. Second, if the inputs are equal, then

$$(K_1, 3, N_{P_1}, \text{ixor}(A_1, T_1)) = (K_2, 3, N_{P_2}, \text{ixor}(A_2, T_2)),$$

which implies that both contexts have the same key  $K_1 = K_2$  and public nonce  $N_{P_1} = N_{P_2}$ . This means that the only way  $\mathcal{A}$  could have won the CMT game was if  $A_1 \neq A_2$ . But same key and public nonce implies  $P_1 = P_2$  which implies that the inner tags  $T_1 = T_2$ . And  $T_1 = T_2$  and  $\text{ixor}(A_1, T_1) = \text{ixor}(A_2, T_2)$  implies that  $A_1 = A_2$ . Thus,  $\mathcal{A}$  could not have won the CMT game.  $\square$

**NAE Security of the TBC-based scheme.** Analysis of the NAE security of  $\Theta\text{CH}$  is involved. It is immediately more subtle than that for  $\text{OCB3}$ , even for just achieving confidentiality, because  $\text{OCH}$  targets nonce-hiding while  $\text{OCB3}$  does not. In particular,  $\text{OCB3}$  achieves a rather direct result by reducing to the underlying TBC’s security—each tweak uses the nonce and since adversaries use unique nonces, then each TBC invocation is on a distinct tweak. From there, the proof can promptly conclude confidentiality.

This approach does not work for  $\text{OCH}$  because our nonce-hiding construction means that necessarily some of the tweaks used by the TBC cannot rely on the full nonce. Thus, the TBC ends up being used on the same tweak across multiple encryption queries. We then must separately, after replacing  $F$  and  $\tilde{E}$  with their ideal counterparts, perform an analysis showing that security holds despite reuse of a tweak.

Our analysis takes advantage of the modularity. Due to tweak separation for  $\tilde{E}$  and domain separation for  $F$ , we can first transition to the information theoretic setting in a standard way. At this point, queries are handled by either  $\text{OCH-tiny}$  and  $\text{OCH-core}$ , depending on length, and, importantly, using completely independent underlying permutations and functions. Thus, we can reduce to the information theoretic security of the two underlying schemes. See Appendix ??.

**Theorem 13.** *Let  $\tilde{E} :: \{0, 1\}^{\tilde{k}} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be an  $n$ -bit TBC with tweak space  $\mathcal{T} = \{\mathcal{N} \times \mathbb{Z} \times \{\epsilon, \star, \$\}\}$ . Let  $G :: \{0, 1\}^h \times \{0, 1\}^{**} \rightarrow \{0, 1\}^\tau$  be an  $\epsilon$ -AXU hash function. Let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$  where  $k = \tilde{k} + h$ . Let  $\mathcal{A}$  be an NAE adversary making at most  $q < 2^{2\kappa-1}$  queries totalling  $\sigma$  blocks, with at most  $q_v$  verification queries. Assume  $\mathcal{A}$  queries  $u$  users, making at most  $B$  queries (encryption and verification) per user totalling  $s$  blocks per user. Then, we can construct a PRF adversary  $\mathcal{B}$  and TPRP adversary  $\mathcal{C}$  such*

that

$$\begin{aligned} \text{Adv}_{\Theta\text{CH}[\tilde{E},\text{G},\text{F}]}^{\text{nae}}(\mathcal{A}) \leq & + \text{Adv}_{\text{F}}^{\text{prf}}(\mathcal{B}) + \text{Adv}_{\tilde{E}}^{\text{tprp}}(\mathcal{C}) + \frac{2q_v}{2^\kappa} + \frac{q(3B+1)}{2^{2\kappa}} \\ & + \frac{q \cdot \min(B, 2^{n_s})}{2^w} + 3q \min(B, 2^{n_s}) \cdot \epsilon(\ell), \end{aligned}$$

where  $\ell := 8 + n_s + w + r$  is the maximum length that  $\Theta\text{CX}$  calls the AXU hash  $\text{G}$  on. Adversary  $\mathcal{B}$  makes at most  $2q$  PRF queries and has running time similar to  $\mathcal{A}$ . Adversary  $\mathcal{C}$  makes at most  $u + \sigma$  TPRP queries and at most  $s$  TPRP queries per user, and has running time similar to  $\mathcal{A}$ .

(sanketh: OCT Security Start)

**Security of OCT.** Finally, we show multi-user strong TPRP security of OCT in the ideal permutation model. The latter means that the games  $\text{TPRP}^{\text{real}}$  and  $\text{TPRP}^{\text{rand}}$  in ?? (Section 3.2) are extended to draw  $\Pi \leftarrow \$ \text{Perms}(n)$  at random, and then  $\tilde{E} := \text{OCT}[\Pi]$  in replying to  $\text{TPRPENC}$  queries, and  $\tilde{D} := \text{OCT}^{-1}[\Pi, \Pi^{-1}]$  for  $\text{TPRPDEC}$  queries. Also, the adversary has oracles for  $\Pi, \Pi^{-1}$ . Let  $p$  be the number of ideal permutation queries made (to either  $\Pi$  or  $\Pi^{-1}$ ). We write the adversary advantage as  $\text{Adv}_{\text{OCT}}^{\text{tprp}}(\mathcal{A})$ . Note that  $\Pi$  does not enter this notation because now it is a random variable in the game rather than fixed. In ?? we show the following.

**Theorem 14.** *Let  $\mathcal{A}$  be a TPRP adversary against  $\text{OCT}[\Pi]$ , over an ideal permutation  $\Pi : \{0, 1\}^n \rightarrow \{0, 1\}^n$ . Suppose it makes at most  $q$  TPRP queries (to  $\text{TPRPKG}$ ,  $\text{TPRPENC}$ , and  $\text{TPRPDEC}$ ) and at most  $p$  queries to the ideal permutation (to  $\Pi$  and  $\Pi^{-1}$ ), across all users, then*

$$\text{Adv}_{\text{OCT}}^{\text{tprp}}(\mathcal{A}) \leq \frac{(5q + 2)(p + q) + 2q + 2}{2^{2\kappa}}.$$

This bound gives good multi-user security in the sense that the bound does not depend on the number of users. Rather, it depends only on the number of TPRP queries  $q$  made across all users.

We refer to ?? for the full proof, and discuss the ideas and nuances here. We first give two building blocks, as follows.

First, we consider  $\text{TEM}[\Pi, H]$ , which is the tweakable Even-Mansour cipher defined as  $\Pi((X \oplus H_{\tilde{K}}(T)) \parallel 0^*) \oplus H_{\tilde{K}}(T)$  and where the hash function  $H$  (unlike  $\text{OH}$ ) is independent of  $\Pi$ . That means that in the ideal permutation model,  $H$  does not make calls to the  $\Pi, \Pi^{-1}$  oracles. CLS [41] showed single-user strong PRP security in this setting assuming  $H$  is  $\epsilon$ -AXU and  $\delta$ -uniform. ?? in the Appendix gives a good bound for the

multi-user setting under the same conditions.

Second, we define a hash function  $\overline{\text{OH}} :: \text{Perms}(n) \times \mathcal{T} \rightarrow \{0, 1\}^n$  that we call the idealized version of OH. As the notation indicates, the key is a permutation  $\Gamma \in \text{Perms}(n)$ . This function will replace calls to computations of  $\Pi(K \oplus X)$  in OH with calls to  $\Gamma(X)$ , so that  $\overline{\text{OH}}$  is independent of  $\Pi$ . Lemma ?? computes  $\epsilon, \delta$  for which it is shown that  $\overline{\text{OH}}$  is  $\epsilon$ -AXU and  $\delta$ -uniform. Since  $\overline{\text{OH}}$  is independent of  $\Pi$ , Theorem ?? applies to show multi-user strong PRP security of the TBC  $\text{TEM}[\Pi, \overline{\text{OH}}]$  in the model where  $\Pi$  is ideal.

Now we want to conclude that  $\text{OCT}[\Pi]$  is secure. The difficulty is that the hash used  $H = \text{OH}[\Pi]$  is *not* independent of  $\Pi$ : it does make  $\Pi$  queries. We address this by using the multi-user strong SPRP security of padded-key Even-Mansour from ADMV [6] (Lemma ??) to move to a game where OH is replaced by  $\overline{\text{OH}}$ . (Actually we only need PRP security, not strong.) Once this is done, we have  $\text{TEM}[\Pi, \overline{\text{OH}}]$ , which as above can be analyzed via Theorem ?? and Lemma ??.

(sanketh: OCT Security End)



**Security of OCH.** The following theorem bounds the multi-user NAE security of  $\text{OCH}[\Pi, G, F] = \Theta\text{CH}[\text{OCT}[\Pi], G, F]$  in the ideal permutation model, it combines Theorem 13 and Theorem 19.

**Theorem 15.** *Let  $\Pi \leftarrow \$\text{Perms}(w)$  be an ideal  $w$ -bit permutation. Let  $G :: \{0, 1\}^h \times \{0, 1\}^{**} \rightarrow \{0, 1\}^\tau$  be an  $\epsilon$ -AXU hash function. Let  $F :: \{0, 1\}^{2\kappa} \times \mathcal{L} \times \{0, 1\}^{**} \rightarrow \{0, 1\}^*$  be a variable-output-length collision-resistant PRF with length space  $\mathcal{L} = \{k, 2\kappa\}$  where  $k = 2\kappa + h$ . Let  $\mathcal{A}$  be an NAE adversary making at most  $q < 2^{2\kappa-1}$  queries totalling  $\sigma$  blocks, with at most  $q_v$  verification queries. Assume  $\mathcal{A}$  queries  $u$  users, making at most  $B$  queries (encryption and verification) per user totalling  $s$  blocks per user. Assume  $\mathcal{A}$  makes  $p \geq \sigma$  queries to the ideal permutation (to  $\Pi$  and  $\Pi^{-1}$ ). Then, we can construct a PRF adversary  $\mathcal{B}$  such that*

$$\begin{aligned} \text{Adv}_{\text{OCH}[\Pi, G, F]}^{\text{nae}}(\mathcal{A}) \leq & \text{Adv}_F^{\text{prf}}(\mathcal{B}) + \frac{10p^2 + 3qB + 56pu + q}{2^{2\kappa}} + \frac{2q_v}{2^\kappa} \\ & + \frac{q \cdot \min(B, 2^{n_s})}{2^w} + 3q \min(B, 2^{n_s}) \cdot \epsilon(\ell), \end{aligned}$$

where  $\ell := 8 + n_s + w + r$  is the maximum length that  $\Theta\text{CX}$  calls the AXU hash  $G$  on. Adversary  $\mathcal{B}$  makes at most  $2q$  PRF queries and has running time similar to  $\mathcal{A}$ .

(**sanketh:** Define all the terms  $q$ ,  $q_v$ ,  $B$ , etc., explain this bound as a proof sketch, give matching attacks, and note that it is tight (up to a constant factor).)

(**sanketh:** Discuss the choice of constants,  $w = 256$ , etc.)

(**sanketh:** Put in the text what bounds we achieve with different amounts of :  $2^{32}$ ,  $2^{64}$ ,  $2^{80}$ . Focus on ROR security bounds.)

(**sanketh:** Cite Iwata for GCM, use their resource comparison?)

### 3.6 Performance

We evaluate the performance of both XtH as a general transform and OCH as a full construction. We start by describing our concrete instantiation and our evaluation methodology, then answer the following questions:

1. How does XtH compare to prior hash-based transforms?
2. What is the performance difference between public and hidden nonces?
3. Is OCH faster than other high-security schemes that also achieve 128-bit context commitment and AE security?
4. Is OCH competitive with widely deployed, but low security schemes like AES-GCM and ChaCha20/Poly1305?

Our experimental analysis targets devices with the widely available AES instructions, accounting for 98% of the devices that participated in the March 2025 Steam survey [42] and encompassing even lighter weight platforms such as the Raspberry Pi 5 [86].

**OGH instantiations.** In our experiments, we fix key length, total nonce length, and tag length to be  $k = \ell = \tau = 256$ . We implemented and experimented with a variety of OGH instantiation choices. This included various permutation choices for the TBC: Sparkle256 [50], Areion256 [72], Areion512 [72], and the reduced round version of the 1024-bit-wide blake2b permutation [109, 9, 62].

We also explored different choices for the CR-PRF including SHA256 and SHA3-256, a Sponge [32] hash using Sparkle512, and a Sponge hash using Areion512. For the Sponge instances, we used rate  $r = 256$  and capacity  $c = 256$ . One could use an overwrite Sponge [31] to reduce memory state by 256 bits; this won't otherwise

materially affect performance. We use padding to ensure that computing  $(\tilde{K}, H) \leftarrow F_K^k(3)$  takes only two permutation calls, and then Sponge computation continues from the intermediate state to finalize tag computation.

For AXU hash, we use two POLYVAL (the version in [67]) instances (making  $|H| = 256$ ) in parallel to provide ensure 128-bit security for internal tags. A discussion of this construction and its security appears in ???. This eschews a performance optimization possible when using only public nonces (where we could use a single POLYVAL instance); the performance benefit of customizing the hash to the instantiation is negligible given that we are only applying the AXU hash to a few blocks of data.

Experiments showed that for our target platforms, a choice of Areion256 for the TBC and Areion512 for the CR-PRF is highest performing. We call this instantiation AreionOCH, and below just refer to it as OCH. We let AreionOCH-P be OCH using 256-bit public nonces and no secret nonces. Similarly, AreionOCH-S denotes using no public nonce and a 256-bit secret nonce.

**Implementing OCH.** We report on our optimized implementation of OCH, which consists of about 1,000 lines of C. We also provide an unoptimized, simpler to understand reference implementation that is a bit more than 170 lines of C. Our implementation uses a modified version of the reference code for Areion [72] to achieve higher throughput by interleaving four concurrent permutation computations in order to saturate the processor’s CPU pipelining. We implemented in-place buffer management, using careful pointer arithmetic to handle secret nonces which must be placed before the ciphertext.

**Experimental setup.** We use three platforms to represent typical hardware: an Intel

NUC with an i7-1360P, a Raspberry Pi 5 with an ARM Cortex-A76, and an Apple Macbook Pro with an M2 Pro. The Intel NUC is representative of Intel servers and laptops, the Raspberry Pi is representative of phones and medium-to-high-end IoT devices, and the Macbook Pro is representative of Arm servers and laptops. All systems had greater than 1 GB of memory and `aes`, `clmul`, and `sha2` instructions.<sup>1</sup> In addition, the Intel platform has `vaes` and `vclmul` instructions, and the Apple platform has `sha3` instructions. Our implementations are optimized to use the fastest available instructions on each platform.

To run the benchmarks, we used BoringSSL’s speed tool [34]. It runs the function we want to benchmark in a loop for an estimated 100 milliseconds and returns the number of successful invocations and the actual time taken in microseconds. We extended this tool to also measure CPU cycles on the supported Intel NUC and the Raspberry Pi following Pornin [99]. From these measurements, we can compute throughput in millions of bytes-per-second as  $(\text{msgLen} \cdot \text{numCalls})/\text{micros}$ , and cycles-per-byte as  $\text{cycles}/(\text{msgLen} \cdot \text{numCalls})$ , where `msgLen` is in bytes. We report on the median of four executions of the benchmark.

**Performance of the XtH transform** Figure 3.13 compares the performance of XtH (Section 3.3) to the best prior transform in the literature CTY [39, 16] for different choices of the CR-PRF. Since these transforms only differ on the amount of data hashed, we fix the base AEAD to be AES-GCM and keep the message size constant at 16 bytes, and only vary the CR-PRF and the amount of associated data. We report on four CR-PRFs built from SHA256, Blake2b, Ascon, and SHA3-256.

Figure 3.13 (lower is better) shows how XtH outperforms CTY for some associated

---

<sup>1</sup>`clmul` corresponds to `pclmulqdq` on Intel and `pmull` on Arm, and `vclmul` corresponds to `vpclmulqdq` on Intel.

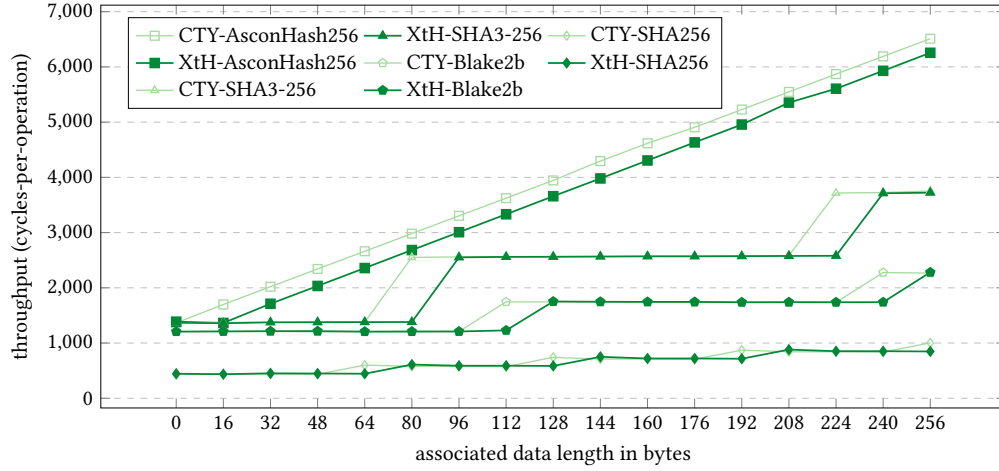


Figure 3.13: **Committing transform performance on desktop x86-64.** Throughput in CPU cycles-per-operation (y-axis, lower is faster) for encrypting a 16-byte message with varying amounts of associated data (x-axis, in bytes) on an Intel Raptor Lake processor. Solid lines indicate schemes that achieve 128-bit NAE and 128-bit CMT security, and dashed lines indicate schemes that do not.

data lengths by up to 1.8x (when using SHA3-256). In most cases, however, the performance difference is negligible: XtH saves on performance in situations where it avoids an extra underlying primitive call (compression function or permutation) that cannot be avoided by CTY. Ultimately we conclude that XtH is, as expected, the preferred transform for rendering a scheme context committing.

**Public versus secret nonces.** We next turn to evaluating the performance impact of nonce hiding. OCH was engineered to make nonce hiding cheap: when one uses a secret nonce instead of a public nonce you add an additional TBC call but avoid an extra block of data processing by the CR-PRF. The difference will be small, except perhaps for small messages.

Figure 3.14 reports performance of AreionOCH-P (solid green circles) and AreionOCH-S (unfilled green circles) across the three platforms in terms of throughput in millions of bytes per second for encryption of increasingly large messages with

a fixed associated data length of 13 bytes (the amount used in TLS 1.2). We focus on throughput since it demonstrates the performance trends; we provide the more standard cycles-per-byte graphs in ??.

Interestingly, we see across all three platforms that performance either does not much change depending on nonce type, or that OCH with public nonces performs slightly better. The latter effect is most visible for moderate sized messages on the two ARM systems (middle and bottom graphs). Further experimentation showed that this is due to buffer management effects, as in-place encryption for secret nonces requires writing into the next block. Nevertheless, the performance gap is minor, and using a separate buffer for the portion of the ciphertext containing the nonce might obviate it.

We also experimented with split public and secret nonces (each 128 bits). This performs slightly worse than fully secret nonces (and so public nonces), because we need to initialize OCT twice (once with the public nonce for encryption of the secret nonce, once for the remaining blocks of encryption).

**OCH versus high-security schemes.** For most message lengths, OCH matches or is faster than other high-security schemes. We now turn to comparing the performance of OCH with other schemes achieving similar levels of security, most relevantly 128-bit NAE and context commitment security. This rules out some schemes that achieve these security goals, but at lower security levels, e.g., CTX applied to AES-GCM doesn't achieve 128-bit NAE and Ascon [47] only enjoys 64 bits of context-commitment security. We discuss such comparisons below.

We are aware of only one family of schemes, the recently proposed context committing AEAD schemes from Daemen et al. [43] built from the hash functions SHAKE128 and TurboSHAKE128. Due to the current lack of availability of source code for




Figure 3.14: **AEAD performance.** Throughput in millions of bytes per second (y-axis, higher is better) for encrypting messages of various sizes (x-axis, in bytes) for 13 bytes of associated data on **(top)** Intel Raptor Lake, **(middle)** an Apple M2 Pro, and **(bottom)** ARM Cortex-A76. Solid lines indicate schemes that achieve 128-bit NAE and 128-bit CMT security, and dashed lines indicate schemes that do not.

	AreionOCH-S	AreionOCH-P	SHAKE128	TurboSHAKE128	AES128-GCM	AES256-GCM	ChaCha20/Poly1305	Ascon-AEAD128
	high-security				low-security			
Intel Raptor Lake	3458.8 MB/s	0.98×	<b>10.88×</b>	<b>5.58×</b>	0.62×	0.72×	<b>2.52×</b>	<b>11.40×</b>
Apple M2 Pro	5988.4 MB/s	0.99×	<b>7.18×</b>	<b>3.62×</b>	0.79×	0.91×	<b>3.50×</b>	<b>12.02×</b>
ARM Cortex-A76	1389.9 MB/s	0.99×	<b>3.76×</b>	<b>1.94×</b>	0.63×	0.76×	<b>1.75×</b>	<b>3.40×</b>

Figure 3.15: Throughput (millions of bytes per second) speedup of AreionOCH-S with respect to other AEAD schemes for 65,537 byte messages with 13 bytes of associated data and on Intel Raptor Lake, an Apple M2 Pro, and ARM Cortex-A76. AreionOCH-S is significantly faster than most schemes (bold entries with speedup >1), and competitive elsewhere.

these AEAD schemes, we instead measure timings of the underlying hash functions (SHAKE128, TurboSHAKE128) on the same number of bytes as would be hashed in the AEAD schemes. This provides a conservative comparison (the AEAD modes add overhead beyond hashing); we confirmed this fact via communication with the authors of [43].

As can be seen, both versions of OCH perform significantly better for moderate and longer messages. For very long messages, Figure 3.15 shows OCH achieving up to 10× the throughput of SHAKE128 and 5× of TurboSHAKE128. Even on Apple M2 Pro with sha3 instruction, OCH reaches 7× the throughput of SHAKE128 and 3.5× of TurboSHAKE128. For tiny messages hashing is slightly faster, though recall that this comparison ignores other overheads for implementing AEAD using the hash functions. The crossover point is 256 bytes for AreionOCH-P and 512 bytes for AreionOCH-S.

**OCH versus low-security schemes.** We now turn to an *unfair* comparison of OCH to schemes that only achieve much weaker security levels. Even here OCH can out-



perform many schemes for a range of message lengths, and comes close to matching the performance of the fastest low-security AEAD schemes. To establish this, we performed extensive experiments with other AEAD schemes, including AES128-GCM and AES256-GCM [53]; ChaCha20/Poly1305 [97] and XChaCha20/Poly1305 [8]; Blake2b-OPP-MEM [62]; Aegis [46]; Ascon [47]; and OCB3 using AES128 [81]. We plan to release a public report with the full gamut of performance results, and only discuss representative comparisons here for brevity. (sanketh: Link to <https://cryptography.run> after polishing it)

As seen in Figure 3.14, OCH outperforms ChaCha20/Poly1305 for medium and long messages on all three platforms. For longer messages of length  $2^{16} + 1$  bytes, OCH achieves 1.75× to 3.5× the throughput of ChaCha20/Poly1305. For the ARM architectures, ChaCha20/Poly1305 achieves slightly higher throughput for shorter messages, with crossover point around 1,024 bytes for AreionOCH-P and around 2,048 bytes for AreionOCH-S.

AES256-GCM is very fast on these platforms given the extensive hardware support tailored to it, and it outperforms OCH. We again emphasize that AES256-GCM is insufficiently secure for important applications as explained earlier in the paper. For longer messages, the overhead of OCH over AES-GCM can be as low as 10% (on Apple M2, against AES256-GCM), to as high as 61% (on Intel Raptor Lake, against AES128-GCM.) This shows that moving to OCH with its significantly better security does not cost much.

Figure 3.15 in the appendix contains more speed-up comparisons with other AEAD schemes.

(sanketh: Specify a “simplified version” of OCH that does not use an AXU hash,

instead applies CR-PRF to (checksum, secnonce, ad, pubnonce) An appendix on “simplified OCH”, don’t need full proofs. )

(**sanketh:** Add cold-start numbers with random keys.)

(**sanketh:** Add random nonces and sequential nonces.)

(**sanketh:** set up a website with all the numbers.)

## 3.7 (In Preparation) TPRP Security of OCT

### 3.7.1 Updated TBC Intro

(sanketh: adapted from prelims, move back)

(sanketh: make sure that we count all ideal permutation queries, even those inside EM and Offset oracles. In prelims, explicitly mention all the operations that we assume are unit cost.)

### 3.7.2 TPRP Security of OCT

In this section, we will prove the TPRP security of OCT in the ideal permutation model, as stated in Theorem 19. But first we consider the SPRP and TPRP security of Even-Mansour and tweakable Even-Mansour, respectively.

**SPRP security of Even-Mansour.** First, we consider the multi-user SPRP security of the padded-key Even-Mansour blockcipher (defined in Figure 3.16) in the ideal permutation model. This was shown by ADMA [6, Theorem 2 & 3] (see also [93, Theorem 1]) and we give a simplified version of their result below in Theorem 16.

**Theorem 16** ([6]). *Fix a block size  $w \geq 2\kappa$ , and let  $P \leftarrow_{\$} \text{Perm}(\{0, 1\}^w)$  be a  $w$ -bit ideal permutation. Let  $\mathcal{A}$  be an SPRP adversary against EM making  $q$  SPRP queries (across all users) and  $p$  ideal permutation queries. Then,*

$$\text{Adv}_{\text{EM}}^{\text{sprp}}(\mathcal{A}) \leq \frac{q^2 + 2qp}{2^{2\kappa}}.$$

$\text{EM}^{P^\pm}.\text{Enc}(K, X):$ $\triangleright \{0, 1\}^{2\kappa} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\Delta_K \leftarrow K \parallel 0^{w-2\kappa}$ $\text{return } P(\Delta_K \oplus X) \oplus \Delta_K$ $\text{EM}^{P^\pm}.\text{Dec}(K, Y):$ $\triangleright \{0, 1\}^{2\kappa} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\Delta_K \leftarrow K \parallel 0^{w-2\kappa}$ $\text{return } P^{-1}(\Delta_K \oplus Y) \oplus \Delta_K$ $\text{TEM}^{P^\pm}[\text{G}].\text{Enc}(K, T, X):$ $\triangleright \{0, 1\}^{2\kappa} \times \mathcal{T} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\Delta_{K,T} \leftarrow K \parallel 0^{w-2\kappa}$ $\text{return } P(\Delta_{K,T} \oplus X) \oplus \Delta_{K,T}$ $\text{TEM}^{P^\pm}[\text{G}].\text{Dec}(K, T, Y):$ $\triangleright \{0, 1\}^{2\kappa} \times \mathcal{T} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\Delta_{K,T} \leftarrow \text{G}(K, T)$ $\text{return } P^{-1}(\Delta_{K,T} \oplus Y) \oplus \Delta_{K,T}$	$\text{OH}^{P^\pm}(K, T):$ $\triangleright \{0, 1\}^{2\kappa} \times \mathcal{T}_{\text{OCH}} \rightarrow \{0, 1\}^{2\kappa}$ $L \leftarrow \text{EM}^{P^\pm}.\text{Enc}(K, 0^w)[1 : 2\kappa]$ $\text{match } T \quad // \text{ structural pattern matching}$ $\text{case } (N_P, j) \leftarrow T \text{ then}$ $R_{N_P} \leftarrow \text{EM}^{P^\pm}.\text{Enc}(K, N_P \parallel 0^{w-n_P-2} \parallel 10)[1 : 2\kappa]$ $\text{return } (((j+2) \cdot L) \oplus R_{N_P})$ $\text{case } (N_P, N_S, i, j) \leftarrow T \text{ then}$ $\text{Top} \leftarrow N_P \parallel N_S[1 : n_S - 6]$ $\text{KTop} \leftarrow \text{EM}^{P^\pm}.\text{Enc}(K, \text{Top} \parallel 0^{w-n_P-n_S+4} \parallel 01)$ $\text{Bottom} \leftarrow \text{str2int}(N_S[n_S - 5 : n_S])$ $R_{N_P, N_S} \leftarrow \text{SH}(\text{KTop}, \text{Bottom})$ $\text{return } (((4\gamma_i + j) \cdot L) \oplus R_{N_P, N_S})$ $\text{SH}(\text{KTop}, \text{Bottom}):$ $\triangleright \{0, 1\}^w \times [1..63] \rightarrow \{0, 1\}^{2\kappa}$ $\text{if } w < (2\kappa + 64) \text{ then}$ $\text{TKTop} \leftarrow \text{KTop}[1 : 2\kappa]$ $\text{KTopStretch} \leftarrow \text{TKTop} \parallel (\text{TKTop} \oplus (\text{TKTop} \ll c))$ $\text{elseif } w \geq (2\kappa + 64) \text{ then}$ $\text{KTopStretch} \leftarrow \text{KTop}$ $\text{return } (\text{KTopStretch} \ll \text{Bottom})[1 : 2\kappa]$	$\text{OCT}^{P^\pm}[\text{G}].\text{Enc}(K, T, X):$ $\triangleright \{0, 1\}^{2\kappa} \times \mathcal{T}_{\text{OCH}} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\text{return TEM}^{P^\pm}[\text{OH}^{P^\pm}].\text{Enc}(K, T, X)$ $\text{OCT}^{P^\pm}[\text{G}].\text{Dec}(K, T, Y):$ $\triangleright \{0, 1\}^{2\kappa} \times \mathcal{T}_{\text{OCH}} \times \{0, 1\}^w \rightarrow \{0, 1\}^w$ $\text{return TEM}^{P^\pm}[\text{OH}^{P^\pm}].\text{Dec}(K, T, Y)$
--	---	---

Figure 3.16: Definitions of (Left) the Even-Mansour blockcipher EM and the tweakable Even-Mansour TBC TEM, (Middle) the offset function OH, and (Right) the tweakable blockcipher OCT. All constructions are based on a permutation  $P : \{0, 1\}^w \rightarrow \{0, 1\}^w$  with blocksize  $w \geq 2\kappa$ . The choice of constant  $c$  in SH is in Lemma 18.

**TPRP security of tweakable Even-Mansour.** Next, we consider the multi-user TPRP security of tweakable Even-Mansour TBC TEM[G] (defined in Figure 3.16) in the ideal permutation model, with an the offset function G that is *independent* of the permutation  $P$ . This was previously considered by Zhang and Hu [119], but unfortunately their proof is flawed because of an incorrect combination of inequalities. Still, the proof can be easily fixed, and the updated bound is given in Theorem 17 below with the corrected proof is in ??.

**Theorem 17.** Fix a block size  $w \geq 2\kappa$ . Let  $P \leftarrow_s \text{Perm}(\{0, 1\}^w)$  be a  $w$ -bit ideal permutation, let  $\mathcal{T}$  be the tweakspace, let  $H : \{0, 1\}^{2\kappa} \times \mathcal{T} \rightarrow \{0, 1\}^w$  be an  $\epsilon$ -AXU and  $\delta$ -uniform offset function (that is independent of the ideal permutation  $P$ ). Let  $\mathcal{A}$  be a TPRP adversary making  $q$  TPRP queries and  $p$  ideal permutation queries. Assume  $\mathcal{A}$  queries  $u$  users,

$\kappa$	$c$																															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	32											
64	3	15	7	3	124	7	3	85	120	3	118	63	3	31	63	3	7	31	3	7	15											
128	255	3	31	15	3	7	250	3	248	247	3	7	7	3	242	85	3	239	7	3	3											
256	3	511	7	3	7	31	3	15	15	3	63	7	3	499	498	3	7	341	3	341	85											

Figure 3.17: `maxBottom` for various choices of constant  $c$ . The selected values of  $c$  are highlighted.

making at most  $B$  TPRP queries (encryption and verification) per user. Then,

$$\text{Adv}_{\text{TEM}}^{\text{tpRP}}(\mathcal{A}) \leq qB \cdot \epsilon + q(q + 2p) \cdot \delta.$$

**AXU security of the SH hash.** Now, we consider the  $\delta$ -uniform and  $\epsilon$ -AXU security of the hash  $\text{SH} : \{0, 1\}^w \times [1..63] \rightarrow \{0, 1\}^{2\kappa}$  (defined in Figure 3.16). SH generalizes the Stretch-then-Shift hash in OCB3 [79, §7.2], and our analysis is an extension of the analysis in OCB3 [79, §7.2]. We state the result below in Lemma 18.

**Lemma 18.** Fix a security parameter  $\kappa \in \{64, 128, 256\}$  and a block size  $w \geq 2\kappa$ . Then  $\text{SH} : \{0, 1\}^w \times [1..63] \rightarrow \{0, 1\}^{2\kappa}$  (defined in Figure 3.16) is  $\delta$ -uniform and  $\epsilon$ -AXU, with  $\delta = \epsilon = 2^{-2\kappa}$ . That is, for any input  $\text{Bottom} \in [1..63]$ , any distinct inputs  $\text{Bottom}_1, \text{Bottom}_2 \in [1..63]$ , and any output  $\Delta \in \{0, 1\}^{2\kappa}$ , it holds that

$$\Pr_{\text{KTop} \leftarrow \{0, 1\}^w} [\text{SH}(\text{KTop}, \text{Bottom}) = \Delta] \leq 2^{-2\kappa},$$

$$\Pr_{\text{KTop} \leftarrow \{0, 1\}^w} [\text{SH}(\text{KTop}, \text{Bottom}_1) \oplus \text{SH}(\text{KTop}, \text{Bottom}_2) = \Delta] \leq 2^{-2\kappa}.$$

*Proof.* Let's first consider the simpler case when blocksize  $w \geq 2\kappa + 64$ . Then  $\text{SH}(\text{KTop}, \text{Bottom}) = (\text{KTop} \ll \text{Bottom})[1 : 2\kappa]$  selects  $2\kappa$  bits from the first  $2\kappa + 64$  bits of  $\text{KTop}$ . So for any input  $\text{Bottom} \in [1..63]$  and any output  $\Delta \in \{0, 1\}^{2\kappa}$ , it holds that

$$\Pr_{\text{KTop} \leftarrow \{0, 1\}^w} [\text{SH}(\text{KTop}, \text{Bottom}) = \Delta \mid w \geq 2\kappa + 64] \leq 2^{-2\kappa}.$$

Now, fix some distinct inputs  $\text{Bottom}_1 < \text{Bottom}_2$  and some output  $\Delta$ . Then the statement  $\text{SH}(\text{KTop}, \text{Bottom}_1) \oplus \text{SH}(\text{KTop}, \text{Bottom}_2) = \Delta$  implies that  $\text{KTop}[i] \oplus \text{KTop}[i + (\text{Bottom}_2 - \text{Bottom}_1)]$  for every  $i \in \{\text{Bottom}_1 + 1, \dots, \text{Bottom}_1 + 2\kappa\}$ . This further implies that  $\text{KTop}[\text{Bottom}_2 + 1 : \text{Bottom}_2 + 2\kappa]$  is uniquely determined from  $\text{KTop}[\text{Bottom}_1 + 1 : \text{Bottom}_2]$  and  $\Delta$ , which over the choice of  $\text{KTop}$  happens with probability at most  $2^{-2\kappa}$ . Thus,

$$\Pr_{\text{KTop} \leftarrow \$\{0,1\}^w} [\text{SH}(\text{KTop}, \text{Bottom}_1) \oplus \text{SH}(\text{KTop}, \text{Bottom}_2) = \Delta \mid w \geq 2\kappa + 64] \leq 2^{-2\kappa}.$$

Let's now analyze the more complex case when blocksize  $w \in \{2\kappa, \dots, 2\kappa + 64\}$  with security parameter  $\kappa \in \{64, 128, 256\}$ . Our analysis follows that of the Stretch-then-Shift hash in OCB3 [79, §7.2]. For any fixed value of the constant  $c$ , following OCB3 [79, §7.2], we construct  $2\kappa \times 2\kappa$  matrices  $A_i$  for  $i$  such that  $\text{SH}(\text{KTop}, i) = A_i \text{TKTop}$ , where  $\text{TKTop} = \text{KTop}[1 : 2\kappa]$  is interpreted as a column vector of size  $|\text{KTop}| \in \{128, 256, 512\}$  and multiplication is in  $\text{GF}(2)$ . Then,  $2^{-2\kappa}$ -uniformity is equivalent to the matrices  $A_i$  being invertible for all  $i$ , and  $2^{-2\kappa}$ -AXU is equivalent to the matrix sums  $(A_i + A_j)$  being invertible for all distinct  $i, j$ . Following this, we wrote a SageMath [110] program to compute the largest value  $\text{maxBottom}$  where  $A_i$  is invertible for all  $i < \{0, \dots, \text{maxBottom} - 1\}$  and  $(A_i + A_j)$  is invertible for all distinct  $i, j \in \{0, \dots, \text{maxBottom} - 1\}$ . We show our results for various choices of  $c$  in Figure 3.17. Thankfully, for each  $\kappa \in \{64, 128, 256\}$ , there are many small choices of  $c$  with  $\text{maxBottom} \geq 64$ . For simplicity, like OCB3 [79, §7.2], we select multiples of eight; that is  $(\kappa, c) \in \{(64, 8), (128, 16), (256, 32)\}$ .  $\square$

**Proof of Theorem 19.** (sanketh: move Theorem 19 back into the body of the paper.)

**Theorem 19.** Fix a security parameter  $\kappa \in \{64, 128, 256\}$ , a block size  $w \geq 2\kappa$ , and let  $P \leftarrow \$\text{Perm}(\{0, 1\}^w)$  be a  $w$ -bit ideal permutation. Let  $\mathcal{A}$  be a TPRP adversary making

```

OFF $\Gamma_i$ ( $T$ ):
 $\triangleright \mathcal{T}_{\text{OCH}} \rightarrow \{0, 1\}^{2\kappa}$ 
 $L \leftarrow \Gamma_i(0^w)[1 : 2\kappa]$  // cached for each user
match  $T$  // structural pattern matching
case ( $N_p, j$ )  $\leftarrow T$  then
   $R_{N_p} \leftarrow \Gamma_i(N_p \parallel 0^{w-n_p-2} \parallel 10)[1 : 2\kappa]$ 
  return  $((j+2) \cdot L) \oplus R_{N_p}$ 
case ( $N_p, N_s, i, j$ )  $\leftarrow T$  then
   $\text{Top} \leftarrow N_p \parallel N_s[1 : n_s - 6]$ 
   $\text{KTop} \leftarrow \Gamma_i(\text{Top} \parallel 0^{w-n_p-n_s+4} \parallel 01)$ 
   $\text{Bottom} \leftarrow \text{str2int}(N_s[n_s - 5 : n_s])$ 
   $R_{N_p, N_s} \leftarrow \text{SH}(\text{KTop}, \text{Bottom})$  // def in Fig 3.16
  return  $((4j_i + j) \cdot L) \oplus R_{N_p, N_s}$ 
endmatch

```

Figure 3.18: Procedure OFF used in proof of Theorem 19. It calls subprocedure SH defined in Figure 3.16.

<p><b>Game <math>G_0(\mathcal{A})</math>:</b></p> <p><math>P \leftarrow \text{Perm}(\mathcal{M})</math>  <math>b \leftarrow \mathcal{A}^{\text{TPRP}[\text{KG}, \text{ENC}, \text{DEC}], P^\pm}</math>  <b>return</b> <math>b</math></p> <p><b>TPRPKG</b>():  <math>u \leftarrow u + 1</math>  <math>K_u \leftarrow \mathcal{K}</math></p> <p><b>TPRPENC</b>(<math>i, T, X</math>):  <math>\Delta_{i,T} \leftarrow \text{OH}^{P^\pm}(K_i, T)</math>  <b>return</b> <math>P(\Delta_{i,T} \oplus X) \oplus \Delta_{i,T}</math></p> <p><b>TPRPDEC</b>(<math>i, T, Y</math>):  <math>\Delta_{i,T} \leftarrow \text{OH}^{P^\pm}(K_i, T)</math>  <b>return</b> <math>P^{-1}(\Delta_{i,T} \oplus Y) \oplus \Delta_{i,T}</math></p>	<p><b>Game <math>G_2(\mathcal{A})</math>:</b></p> <p><math>P \leftarrow \text{Perm}(\mathcal{M})</math>  <math>b \leftarrow \mathcal{A}^{\text{TPRP}[\text{KG}, \text{ENC}, \text{DEC}], P^\pm}</math>  <b>return</b> <math>b</math></p> <p><b>TPRPKG</b>():  <math>u \leftarrow u + 1</math>  <math>\tilde{\pi}_u \leftarrow \text{Perm}(\mathcal{T}, \mathcal{M})</math></p> <p><b>TPRPENC</b>(<math>i, T, X</math>):  <b>return</b> <math>\tilde{\pi}_i(T, X)</math></p> <p><b>TPRPDEC</b>(<math>i, T, Y</math>):  <b>return</b> <math>\tilde{\pi}_i^{-1}(T, Y)</math></p>	<p><b>Game <math>G_1(\mathcal{A})</math>:</b></p> <p><math>P \leftarrow \text{Perm}(\mathcal{M})</math>  <math>b \leftarrow \mathcal{A}^{\text{TPRP}[\text{KG}, \text{ENC}, \text{DEC}], P^\pm}</math>  <b>return</b> <math>b</math></p> <p><b>TPRPKG</b>():  <math>u \leftarrow u + 1</math>  <math>\Gamma_u \leftarrow \text{Perm}(\{0, 1\}^w)</math></p> <p><b>TPRPENC</b>(<math>i, T, X</math>):  <math>\Delta_{i,T} \leftarrow \text{OFF}^{\Gamma_i}(T)</math>  <b>return</b> <math>P(\Delta_{i,T} \oplus X) \oplus \Delta_{i,T}</math></p> <p><b>TPRPDEC</b>(<math>i, T, Y</math>):  <math>\Delta_{i,T} \leftarrow \text{OFF}^{\Gamma_i}(T)</math>  <b>return</b> <math>P^{-1}(\Delta_{i,T} \oplus Y) \oplus \Delta_{i,T}</math></p>
---	---	--

Figure 3.19: Games  $G_0$ - $G_2$  used in proof of Theorem 19.  $G_1$  calls procedure OFF defined in Figure 3.18

$q$  TPRP queries and  $p$  ideal permutation queries. Assume  $\mathcal{A}$  queries  $u$  users, making at most  $B$  TPRP queries (encryption and verification) per user. Then, we can construct an PRP adversary  $\mathcal{B}$  and a TPRP adversary  $\mathcal{C}$  such that

$$\text{Adv}_{\text{OCT}}^{\text{tprp}}(\mathcal{A}) \leq \frac{2.5q^2 + 2qp + 1.5qB + 3p}{2^{2\kappa}}.$$

Adversary  $\mathcal{B}$  makes at most  $2u + q$  PRP queries to  $u$  users, with at most  $B + 1$  queries per

user. Adversary  $\mathcal{C}$  makes at most  $u + q$  TPRP queries to  $u$  users, with at most  $B$  queries per user. Both adversaries  $\mathcal{B}$  and  $\mathcal{C}$  make  $p$  ideal permutation queries each and have running time similar to  $\mathcal{A}$ .

*Proof.* We will transition from  $\text{TPRP}^{\text{real}}$  to  $\text{TPRP}^{\text{rand}}$  using a sequence of games defined in Figure 3.19.

Game  $G_0$  is the real game  $\text{TPRP}^{\text{real}}$ , and game  $G_1$  identical except two changes: (1) for each user  $i$ , it samples a uniformly random permutation  $\Gamma_i \leftarrow \text{Perm}(\{0, 1\}^w)$ , and (2) it replaces each call to the real offset function  $\text{OH}^{P^\pm}(K_i, T)$  with a call to the procedure  $\text{OFF}^{\Gamma_i}(T)$  (defined in Figure 3.19 and described next). To construct  $\text{OFF}^{\Gamma_i}(T)$ , we leverage the fact that  $\text{OH}^{P^\pm}(K_i, T)$  (defined in Figure 3.16) only uses its key  $K_i$  input and permutation  $P^\pm$  oracle to invoke  $\text{EM}^{P^\pm}.\text{Enc}(K_i, \cdot)$ . It makes three such calls, once to compute the subkey  $L$ , once to compute  $R_{N_p}$ , and once to compute  $\text{KTop}$ . Now, in  $\text{OFF}^{\Gamma_i}(T)$  we replace these calls to  $\text{EM}^{P^\pm}.\text{Enc}(K_i, \cdot)$  with calls to the uniformly random permutation  $\Gamma_i$ . After this replacement,  $\text{OFF}^{\Gamma_i}(T)$  no longer depends on the key  $K_i$  or the permutation  $P$ , so we can safely shed them. We bound the gap between games  $G_1$  and  $G_0$  by constructing a PRP adversary  $\mathcal{B}$  for EM in the ideal permutation model.  $\mathcal{B}$  emulates  $G_1(\mathcal{A})$  but runs  $\text{PRPKG}$  for each user  $i$  and replaces calls to the permutation  $\Gamma_i(X)$  with calls to the PRP game's  $\text{PRPENC}(i, X)$ .  $\mathcal{B}$  also forwards all ideal permutation queries of  $G_1(\mathcal{A})$  to its ideal permutation oracle. Then applying Theorem 16 we get

$$\Pr[G_0(\mathcal{A})] - \Pr[G_1(\mathcal{A})] \leq \text{Adv}_{\text{EM}}^{\text{prp}}(\mathcal{B}) \leq \frac{q^2 + 2qp}{2^{2\kappa}}.$$

Adversary  $\mathcal{B}$  queries  $\text{PRPKG}$  once for each user,  $\text{PRPENC}$  once for each user to compute the subkey  $L$  in  $\text{OFF}$  (which it caches for future calls by the same user),  $\text{PRPENC}$  once for each  $\text{TPRPENC}$  and  $\text{TPRPDEC}$  call inside  $\text{OFF}$  to compute either  $R_{N_p}$  or  $\text{KTop}$ . In sum,  $\mathcal{B}$  makes at most  $q + 2u$  PRP queries, forwards the  $p$  ideal permutation queries, and



has time complexity similar to  $\mathcal{A}$ .

Game  $G_2$  is the random game  $\text{TPRP}^{\text{rand}}$ , and we want to bound the gap between  $G_1$  and  $G_2$  by constructing a TPRP adversary  $C$  for  $\text{TEM}[\text{OFF}]$  in the ideal permutation model. Since  $\text{OFF}$  does not depend on the ideal permutation, we are in the setting of Theorem 17. Adversary  $C$  emulates  $G_1(\mathcal{A})$  except it forwards calls to  $\text{TPRPKG}$ ,  $\text{TPRPENC}$ , and  $\text{TPRPDEC}$  to its TPRP oracles, and forwards all ideal permutation queries to its ideal permutation oracle. Suppose  $\text{OFF}$  is  $\epsilon$ -AXU and  $\delta$ -uniform for some  $\epsilon$  and  $\delta$  to be determined later. Then, applying Theorem 17, we get

$$\Pr[G_1(\mathcal{A})] - \Pr[G_2(\mathcal{A})] \leq \text{Adv}_{\text{TEM}[\text{OFF}]}^{\text{tprp}}(C) \leq qB \cdot \epsilon + q(q + 2p) \cdot \delta.$$

Adversary  $C$  has the same query complexity as  $\mathcal{A}$  and similar running time.

At last, it remains to compute the values  $\epsilon$  and  $\delta$  such that  $\text{OFF}$  is  $\delta$ -uniform and  $\epsilon$ -AXU. That is, for any input  $T \in \mathcal{T}_{\text{ECH}}$ , any distinct inputs  $T_1, T_2 \in \mathcal{T}_{\text{ECH}}$ , and any output  $\Delta \in \{0, 1\}^{2\kappa}$ , it should hold that

$$\begin{aligned} \Pr_{\Gamma \leftarrow \text{Perm}(\{0, 1\}^w)} [\text{OFF}^\Gamma(T) = \Delta] &\leq \delta, \\ \Pr_{\Gamma \leftarrow \text{Perm}(\{0, 1\}^w)} [\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta] &\leq \epsilon. \end{aligned}$$

Let's start with  $\delta$ -uniformity. Fix a tweak  $T \in \mathcal{T}_{\text{ECH}}$  and output  $\Delta \in \{0, 1\}^{2\kappa}$ , and pick  $\pi \leftarrow \text{Perm}(\{0, 1\}^w)$ . Let  $L = \pi(0^w)[1 : 2\kappa]$ . First let's consider the case when  $(N_p, j) \leftarrow T$ . Then, there's a constant  $s$  such that  $\text{OFF}^\Gamma(T) = ((s \cdot L) \oplus R_{N_p})$  where  $R_{N_p} = \pi(N_p \parallel 0^{w-n_p-2} \parallel 10)[1 : 2\kappa]$ . Given  $\pi(0^w)$ , the value  $\pi(N_p \parallel 0^{w-n_p-2} \parallel 10)$  is uniformly distributed over a set of  $2^w - 1$  elements. Of these there are at most  $2^{w-2\kappa}$  choices that lead to  $R_{N_p} = (s \cdot L) \oplus \Delta$ . Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0, 1\}^w)} [\text{OFF}^\Gamma(T) = \Delta \mid (N_p, j) \leftarrow T] \leq \frac{2^{w-2\kappa}}{2^w - 1} = \frac{1 + (2^w - 1)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

Now, let's consider the case when  $(N_p, N_s, i, j) \leftarrow T$ . Then, there's a constant  $s$  such that  $\text{OFF}^\Gamma(T) = ((s \cdot L) \oplus R_{N_p, N_s})$  where  $R_{N_p, N_s} = \text{SH}(\text{KTop}, \text{Bottom})$ ,  $\text{KTop} = \pi(\text{Top} \parallel 0^{w-n_p-n_s+4} \parallel 01)$ ,  $\text{Top} = N_p \parallel N_s[1 : n_s - 6]$ , and  $\text{Bottom} = \text{str2int}(N_s[n_s - 5 : n_s])$ . Given  $\pi(0^w)$ , the random variable  $\text{KTop}$  is uniformly distributed over a set of  $2^w - 1$  elements. Using the fact from Lemma 18 that  $\text{SH}$  is  $2^{-2\kappa}$ -uniform, there are at most  $2^{w-2\kappa}$  choices of  $\text{KTop}$  that result in  $\text{SH}(\text{KTop}, \text{Bottom}) = (\Delta \oplus (s \cdot L))$ . Thus, in this case too,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)}[\text{OFF}^\Gamma(T) = \Delta \mid (N_p, N_s, i, j) \leftarrow T] \leq \frac{2^{w-2\kappa}}{2^w - 1} = \frac{1 + (2^w - 1)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

Now, let's consider  $\epsilon$ -AXU. Fix distinct tweaks  $T_1, T_2 \in \mathcal{T}_{\text{CH}}$  and output  $\Delta \in \{0, 1\}^{2\kappa}$ , and pick  $\pi \leftarrow \text{Perm}(\{0, 1\}^w)$ . Let  $L = \pi(0^w)[1 : 2\kappa]$ , and we proceed by case analysis.

*Case 1:*  $(N_{P_1}, j_1) \leftarrow T_1$  and  $(N_{P_2}, j_2) \leftarrow T_2$ . Then, there's a constant  $s$  such that  $\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = ((s \cdot L) \oplus R_{N_{P_1}} \oplus R_{N_{P_2}})$  where  $R_{N_{P_1}} = \pi(N_{P_1} \parallel 0^* \parallel 10)[1 : 2\kappa]$  and  $R_{N_{P_2}} = \pi(N_{P_2} \parallel 0^* \parallel 10)[1 : 2\kappa]$ . Given  $\pi(0^w)$  and  $\pi(N_{P_1} \parallel 0^* \parallel 10)$ , the value  $\pi(N_{P_2} \parallel 0^* \parallel 10)$  is uniformly distributed over a set of  $2^w - 2$  elements. Of these there are at most  $2^{w-2\kappa}$  choices that lead to  $R_{N_{P_2}} = (\Delta \oplus (s \cdot L) \oplus R_{N_{P_1}})$ . Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)}[\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta \mid \text{Case 1}] \leq \frac{2^{w-2\kappa}}{2^w - 2} = \frac{1 + (2^w - 2)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

*Case 2:*  $(N_{P_1}, j_1) \leftarrow T_1$  and  $(N_{P_2}, N_{S_2}, i_2, j_2) \leftarrow T_2$ . Then, there's a constant  $s$  such that  $\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = ((s \cdot L) \oplus R_{N_{P_1}} \oplus R_{N_{P_2}, N_{S_2}})$  where  $R_{N_{P_1}} = \pi(N_{P_1} \parallel 0^* \parallel 10)[1 : 2\kappa]$ ,  $R_{N_{P_2}, N_{S_2}} = \text{SH}(\text{KTop}_2, \text{Bottom}_2)$ ,  $\text{KTop}_2 = \pi(\text{Top}_2 \parallel 0^* \parallel 01)$ ,  $\text{Top}_2 = N_{P_2} \parallel N_{S_2}[1 : n_s - 6]$ , and  $\text{Bottom}_2 = \text{str2int}(N_{S_2}[n_s - 5 : n_s])$ . Given  $\pi(0^w)$  and  $\pi(N_{P_1} \parallel 0^* \parallel 10)$ , the random variable  $\text{KTop}$  is uniformly distributed over a set of  $2^w - 2$  elements. Since we know that  $\text{SH}$  is  $2^{-2\kappa}$ -uniform from Lemma 18, there are at most  $2^{w-2\kappa}$  choices of  $\text{KTop}$  that

result in  $\text{SH}(\text{KTop}, \text{Bottom}) = (\Delta \oplus (s \cdot L) \oplus R_{N_{P_1}})$ . Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)} [\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta \mid \text{Case 2}] \leq \frac{2^{w-2\kappa}}{2^w - 2} = \frac{1 + (2^w - 2)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

*Case 3:*  $(N_{P_1}, N_{S_1}, i_1, j_1) \leftarrow T_1$  and  $(N_{P_2}, N_{S_2}, i_2, j_2) \leftarrow T_2$ . Then, there's a constant  $s$  such that  $\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = ((s \cdot L) \oplus R_{N_{P_1}, N_{S_1}} \oplus R_{N_{P_2}, N_{S_2}})$  where  $R_{N_{P_1}, N_{S_1}} = \text{SH}(\text{KTop}_1, \text{Bottom}_1)$ ,  $\text{KTop}_1 = \pi(\text{Top}_1 \parallel 0^* \parallel 01)$ ,  $R_{N_{P_2}, N_{S_2}} = \text{SH}(\text{KTop}_2, \text{Bottom}_2)$ , and  $\text{KTop}_2 = \pi(\text{Top}_2 \parallel 0^* \parallel 01)$ . Let's proceed by subcase analysis.

*Case 3.1:*  $\text{Top}_1 \neq \text{Top}_2$  where  $\text{Top}_1 = N_{P_1} \parallel N_{S_1}[1 : n_s - 6]$  and  $\text{Top}_2 = N_{P_2} \parallel N_{S_2}[1 : n_s - 6]$ . Given  $\pi(0^w)$  and  $\text{KTop}_1$ , the random variable  $\text{KTop}_2$  is uniformly distributed over a set of  $2^w - 2$  elements. Since we know that  $\text{SH}$  is  $2^{-2\kappa}$ -uniform from Lemma 18, there are at most  $2^{w-2\kappa}$  choices of  $\text{KTop}_2$  that result in  $\text{SH}(\text{KTop}_2, \text{Bottom}_2) = (\Delta \oplus (s \cdot L) \oplus R_{N_{P_1}, N_{S_1}})$ . Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)} [\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta \mid \text{Case 3.1}] \leq \frac{2^{w-2\kappa}}{2^w - 2} = \frac{1 + (2^w - 2)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

*Case 3.2:*  $\text{Top}_1 = \text{Top}_2$  but  $\text{Bottom}_1 \neq \text{Bottom}_2$  where  $\text{Bottom}_1 = \text{str2int}(N_{S_1}[n_s - 5 : n_s])$ , and  $\text{Bottom}_2 = \text{str2int}(N_{S_2}[n_s - 5 : n_s])$ . Given  $\pi(0^w)$ , the random variable  $\text{KTop} := \text{KTop}_1 = \text{KTop}_2$  is uniformly distributed over a set of  $2^w - 1$  elements. Since  $\text{SH}$  is  $2^{-2\kappa}$ -AXU from Lemma 18, we have that there are at most  $2^{w-2\kappa}$  choices of  $\text{KTop}$  that result in  $\text{SH}(\text{KTop}, \text{Bottom}_1) \oplus \text{SH}(\text{KTop}, \text{Bottom}_2) = (\Delta \oplus (s \cdot L))$ . Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)} [\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta \mid \text{Case 3.2}] \leq \frac{2^{w-2\kappa}}{2^w - 1} = \frac{1 + (2^w - 1)^{-1}}{2^{2\kappa}} \leq \frac{1.5}{2^{2\kappa}}$$

*Case 3.3:*  $\text{Top}_1 = \text{Top}_2$  and  $\text{Bottom}_1 \neq \text{Bottom}_2$ , but  $(i_1, j_1) \neq (i_2, j_2)$ . Then  $\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = (s \cdot L)$  where the constant  $s = (4\gamma_{i_1} + j_1) \oplus (4\gamma_{i_2} + j_2)$ . Since the  $2\kappa$ -bit Gray code  $\gamma$  is a permutation on  $\mathbb{Z}_{2^n}$ ,  $i_1, i_2 \leq 2^{\kappa-3}$  and  $j_1, j_2 \in \{0, 1, 2, 3\}$ , it follows that the

constant  $s$  is non-zero. Thus, in this case,

$$\Pr_{\Gamma \leftarrow \text{Perm}(\{0,1\}^w)} [\text{OFF}^\Gamma(T_1) \oplus \text{OFF}^\Gamma(T_2) = \Delta \mid \text{Case 3.3}] \leq \frac{2^{w-2\kappa}}{2^w} = \frac{1}{2^{2\kappa}}.$$

Maximizing over all cases, we get  $\text{OFF}$  is  $\delta$ -uniform and  $\epsilon$ -AXU with  $\delta = \epsilon = 1.5 \cdot 2^{-2\kappa}$ . Subtracting differences completes the proof.  $\square$

## CHAPTER 4

### FLEXIBLE AUTHENTICATED ENCRYPTION\*

In this work we propose a new evolution of symmetric encryption, what we refer to as flexible AEAD. We think this will facilitate and enable upcoming standardization efforts. We give definitions for flexible AEAD as well as a concrete realization in the form of a scheme we simply call ARF. First we set the stage.

**Emerging goals for AEAD.** Recall that in a scheme for AEAD (Authenticated Encryption with Associated Data) [103], encryption takes key, nonce, associated data and message to deterministically return a ciphertext. The classical security requirement is unique-nonce AE (UNAE) security. This means privacy of the message, and authenticity of the message and associated data, assuming encryption never reuses a nonce. The NIST standard AES-GCM is an example of an AEAD scheme.

Since the standardization of AES-GCM, developers and researchers have identified a number of further, desirable security and operational attributes for AEAD. Security attributes include committing security [63, 51, 58, 2, 14, 39], misuse resistance [106] and AE2-security (also called nonce-hiding) [19]; operational attributes include parallelizability, robustness in the sense of [HoangKR15] and support for arbitrary-length nonces and keys. These and other attributes are part of a comprehensive list in the IETF draft on properties of AEAD algorithms [irtf-cfrg-aead-properties-01]. Let us now expand on (some of) these attributes and why they are desirable.

Committing security. A recent line of work has demonstrated the need for key com-

---

\*This chapter is joint work with Julia Len, Viet Tung Hoang, Mihir Bellare, and Thomas Ristenpart. Earlier versions of this work were presented at Third NIST Workshop on Block Cipher Modes of Operation [91] and at Real World Crypto 2024 [28]. This chapter is lightly edited from the draft we are preparing for submission.

mitment for AEAD schemes. Key commitment asks that it be hard to find a ciphertext that decrypts correctly under two (or more) different adversarially-chosen keys [58, 63]. Non-key-committing AEAD was first shown to be a problem for moderation in encrypted messaging [51, 63], and later in password-based encryption [83] and envelope encryption [2], among others. These findings have pushed the cryptography community to begin proposing [2] and deploying new key-committing AEAD schemes [2, 10]. Indeed the recent IETF draft on properties of AEAD algorithms explicitly cites key commitment as a security goal [**irtf-cfrg-aead-properties-01**].

Key commitment definitions don't model attacks in which adversaries exploit lack of commitment to the associated data or nonce. Bellare and Hoang [14] introduced the notion of *context commitment*, and recent work by Menda et al. [**MendaLGR23**] highlights that many AEAD schemes, including some that achieve key commitment, do not achieve context commitment. This motivates the need to expand our goal to context commitment.

Unfortunately, no currently standardized schemes achieve context commitment. This suggests we need to define and standardize new AEAD schemes.

Misuse resistance. Stipulating non-repeating nonces is easy in theory but harder to ensure in practice, where errors and misconfigurations have been reported to lead to repeating nonces. For many widely used and standardized UNAE AEAD schemes, this has led to damaging attacks [**BockZDSJ16**]. Misuse-resistant AE (MRAE) [106] mitigates this by providing UNAE-security when nonces do not repeat, plus as good as possible security if they do. Standardization of an MRAE scheme is a desirable goal.

AE2 security. Embedded in the syntax and usage of AEAD is a weakness relating to the way nonces are handled. Namely, since the nonce is needed for decryption, it is sent

in the clear unless the receiver already has it. But, as Bellare, Ng, and Tackmann [19] point out, some choices of nonces compromise message privacy (for both UNAE and MRAE) and others (like counters) compromise sender anonymity. AE2 (which comes in two forms, UNAE2 strengthening UNAE, and MRAE2 strengthening MRAE) hides the nonce, increasing privacy. This is important for anonymous AE [37]. AE2 emerges as another desirable attribute for a standard.

Robustness. The ciphertext expansion of a scheme is defined as the difference between ciphertext length and plaintext length. UNAE security requires some ciphertext expansion. (Usually 128 bits for AES-GCM.) Some applications cannot permit this. An answer is robust AE [HoangKR15], where one gets the best security possible with a given constraint on the ciphertext expansion.

Key and nonce lengths. The 96-bit nonce-length of AES-GCM is viewed as a limitation because with random nonces it permits at most  $2^{48}$  encryptions before a key change is needed. Schemes would ideally have either no maximum nonce length or a very large one. Similarly, different users or applications want different security levels and thus different key sizes. The need for post-quantum security is also pushing key sizes for symmetric cryptography up. Ideally, a scheme should handle keys of different and arbitrary lengths.

**The challenge for standardization.** One approach for standardization would be to standardize a different scheme for each of some choice of goals. However, the dimensions indicated above give rise to rather a lot of goals. Indeed, there are two choices for AE security (UNAE or MRAE), then a choice for committing security (yes or no) and another for AE2 (yes or no), already  $2^3 = 8$  goals, with further possible choices for robustness, parallelizability, streaming support, and key and nonce lengths.

An alternative is to standardize a scheme for the strongest goal (for example, MRAE, committing and AE2) but this will mean that applications needing less stringent attributes will pay unnecessarily in cost. Indeed, one can't unfortunately achieve all the goals simultaneously while providing best-possible speed for each individual goal. Misuse resistance requires a full pass over the plaintext before the first bits of ciphertext can be produced, and known robust AE construction approaches require building a length-preserving enciphering scheme and then combining with the encode-then-encipher paradigm [20]. In either case you cannot have a scheme that is streaming, let alone parallelizable.

With flexibility, we suggest a different route.

**Flexible AEAD.** At a high level, our idea is to reformulate AEAD so that it takes as run-time input a configuration. The configuration, denoted `cfg` throughout, specifies an operating mode. But the same secret key can be safely used across different configurations, even on a message-by-message basis. So one can encrypt one message under a secret key using a fully parallelizable configuration, and another message under an MRAE configuration, with the same secret key. Protocol developers can use this flexibility as needed for their applications. In most cases they will presumably use a single configuration; in this case flexible AEAD provides some defense-in-depth with respect to misconfigurations for the same secret key.

We provide new definitions to capture the syntax, semantics, and security of flexible AEAD. Flexible AEAD will start with some mandatory (always provided) security and operational attributes. Then through choices of the above-discussed configuration, it will further provide other attributes as options. Yet, this will be implemented in a unified way with what is essentially a single scheme.



As mandatory, we ask for classic UNAE and context commitment. The reason for making the latter mandatory is to avoid errors arising from developers not knowing that they need to turn it on. (Indeed, the errors and attacks we have seen are arising from applications assuming implicitly that commitment is present.) We also always ask for the ability to handle arbitrary-length keys and nonces. As optional, if requested in `cfg`, we support providing MRAE, AE2 and streaming capability. Other configurations can be added as needed.

**The ARF scheme.** Towards realizing flexible AEAD, we propose a clean-slate approach that reimagines AEAD scheme construction to allow it to provide modern security and performance for a broad range of application domains. Our ARF scheme conceptually has a straightforward structure: use a key derivation function (KDF) to derive sub-keys that can be safely used with different underlying schemes. Thus constructively, ARF, brings into the AEAD what used to be external, namely, a KDF function.

This flexibility would perhaps seem to require a complex AEAD scheme, depending on many underlying mechanisms. But we build a “wide waist” strategy: one scheme has a variety of different configurations all using the same underlying primitives. In particular, we build all of ARF’s various configurations from a single underlying cryptographic permutation. For instance, we use a realization of tweaked Even-Mansour [CogliatiLS15] built from a permutation. While ARF works for any sufficiently wide permutation, we have mostly focused on instantiations using Simpira [GueronM16].

We propose three primary ARF configurations. The OCH configuration provides an OCB-like mode of operation that is fully parallelizable, but not MR nor robust. It is the most performant. We also provide CIV (committing synthetic IV) an

SIV [RogawayS07] variant that provides MR security, and is as parallelizable as possible. Finally, we aim to provide robust AE based on prior constructions [Halevi04, HoangKR15] using our underlying tweakable block cipher.

**Organization.** The rest of the paper is organized as follows. In Section ?? we define our new cryptographic primitive flexible AEAD (F-AEAD). In Section ?? we provide an overview of our F-AEAD construction called ARF. This includes an overview of the three primary ARF configurations. Finally, in Section ?? we describe preliminary benchmarks on the OCH configuration compared with other AEAD constructions.

(tom: Some thoughts / questions below)

1. Nonce-agnostic means we give up on counter-specific optimizations, such as the special XOR-universal hash used in OCB3. I conjecture that if we want context commitment, those types of optimizations are at odds with security — the universal hash based approach of deriving the initial offset value in OCB3 would seem at odds with key commitment security. Be good to confirm this.
2. For an MR mode, we can combine a CR-PRF hash of the configuration, key, and associated data plus a non-CR PRF (i.e., something like PMAC) of the message to generate a subkey for use with CTR mode. Alternatively we could just use an OCB-style mode (sans the checksum) for the second-pass of encryption. CTR would seem likely to be faster, but OCB would seem more consistent with the non-MR mode possibly simplifying the design slightly. For now will go with CTR style; my guess is the speed and simplicity (compared to OCB) will outweigh the cross-configuration consistency issue.
3. So far I have optimized for short AD at expense of some simplicity and parallelism. In particular, we are serially processing via a special-purpose KDF the con-

fig, key, and associated data to derive a subkey that is then used with either the non-MR mode (OCB-like) or MR mode (SIV-like). For long AD and MR mode, we can gain some parallelism by ensuring we can start processing message in parallel with processing AD. So here we may want to separately process the config and key to generate a sub-key, while in parallel CR hashing the AD (keyed or not with subkey, doesn't seem to matter) and PRF-ing the message. I suspect for all this we'll also want to just jump to a "long-AD" configuration that handles vectorized AD's via tree hashing, probably one MR and one for non-MR.

4. The last point above suggests another simple approach which is to just embrace non-keyed CR hashing of AD. Then we can isolate the AD processing as totally independent of the rest of the context, effectively replacing AD with its CR hash when AD is long. This would seem to simplify all the above without much in security loss (basically, can move AD collision finding to be key-independent).

## 4.1 Background

**AEAD syntax.** We use a simplified version of the general AE3 syntax [25, 16], that matches the CAESAR call for submissions [36], dubbed AE5 in [96].<sup>1</sup> Thus we define an AEAD scheme as a tuple of algorithms  $\text{AEAD} = (\text{Enc}, \text{Dec})$ , defined over a key space  $\mathcal{K} \subseteq \{0, 1\}^*$ , public nonce space  $\mathcal{N}_p \subseteq \{0, 1\}^*$ , secret nonce space  $\mathcal{N}_s \subseteq \{0, 1\}^*$ , associated data space  $\mathcal{A} \subseteq \{0, 1\}^*$ , message space  $\mathcal{M} \subseteq \{0, 1\}^*$ , and ciphertext space  $\mathcal{C} \subseteq \{0, 1\}^*$ , where:

1.  $\text{Enc} :: \mathcal{K} \times \mathcal{N}_p \times \mathcal{N}_s \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$  is a deterministic algorithm that takes a 5-tuple of a key  $K$ , public nonce  $N_p$ , secret nonce  $N_s$ , associated data  $A$ , and message  $M$ ,

---

<sup>1</sup>In ?? we show that these syntaxes are equivalent.

and returns a ciphertext  $C$ .

2.  $\text{Dec} :: \mathcal{K} \times \mathcal{N}_p \times \mathcal{A} \times \mathcal{C} \rightarrow (\mathcal{M} \times \mathcal{N}_s) \cup \{\perp\}$  is a deterministic algorithm that takes a 4-tuple of a key  $K$ , public nonce  $N_p$ , associated data  $A$ , and ciphertext  $C$ , and returns a message  $M$  and a secret nonce  $N_s$  or an error  $\perp$ .

We call the non-message inputs to  $\text{Enc}$ —the key, public nonce, secret nonce, and associated data—the *encryption context* and the non-ciphertext inputs to  $\text{Dec}$ —the key, public nonce, and associated data—the *decryption context*.

We impose correctness and *tidyness* [95] requirements on AEAD which requires  $\text{Enc}$  and  $\text{Dec}$  to be inverses of each other. Namely for all  $(K, N_p, N_s, A, M, C) \in \mathcal{K} \times \mathcal{N}_p \times \mathcal{N}_s \times \mathcal{A} \times \mathcal{M} \times \mathcal{C}$ :

$$\text{Enc}(K, N_p, N_s, A, M) = C \implies \text{Dec}(K, N_p, A, C) = (M, N_s),$$

$$\text{Dec}(K, N_p, A, C) = (M, N_s) \neq \perp \implies \text{Enc}(K, N_p, N_s, A, M) = C.$$

We also require that the lengths of keys  $|K| =: k$ , public nonces  $|N_p| =: n_p$ , and secret nonces  $|N_s| =: n_s$  be public constants, and the length of the ciphertext  $|C| =: \text{clen}(|M|)$  to be computable by a deterministic public function  $\text{clen}$  of the length of the message  $|M|$ . And define the *ciphertext overhead* to be the function that on input  $m$  returns the difference  $\text{clen}(m) - m$ .

**AEAD security.** We define this in the multi-user setting following Bellare and Hoang [16]. We use the pair of games  $\text{NAE}^{\text{real}}$  and  $\text{NAE}^{\text{rand}}$  which are defined in Figure 4.1, and then set

$$\text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{A}) := \Pr[\text{NAE}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{NAE}^{\text{rand}}(\mathcal{A}) \Rightarrow 1].$$

To prevent trivial victories, we require adversaries to be *valid*—it may not query

Game $\text{NAE}_{\text{AEAD}}^{\text{real}}(\mathcal{A})$	Game $\text{NAE}_{\text{AEAD}}^{\text{rand}}(\mathcal{A})$
$b \leftarrow \mathcal{A}^{\text{NAEKg, NAEenc, NAEver}}$	$b \leftarrow \mathcal{A}^{\text{NAEKg, NAEenc, NAEver}}$
<b>return</b> $b$	<b>return</b> $b$
$\text{NAEKg}()$	$\text{NAEKg}()$
$u \leftarrow u + 1; K_u \leftarrow \mathcal{K}$	$u \leftarrow u + 1$
$\text{NAEenc}(i, N_P, N_S, A, M)$	$\text{NAEenc}(i, N_P, N_S, A, M)$
if $(i, N_P, N_S) \in \mathcal{Q}_N$ <b>then return</b> $\perp$	if $(i, N_P, N_S) \in \mathcal{Q}_N$ <b>then return</b> $\perp$
$C \leftarrow \text{AEAD.Enc}(K_i, N_P, N_S, A, M)$	$C \leftarrow \{0, 1\}^{\text{clen}( M )}$
$\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_P, A, C)\}$	$\mathcal{Q}_D \leftarrow \mathcal{Q}_D \cup \{(i, N_P, A, C)\}$
$\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_P, N_S)\}$	$\mathcal{Q}_N \leftarrow \mathcal{Q}_N \cup \{(i, N_P, N_S)\}$
<b>return</b> $C$	<b>return</b> $C$
$\text{NAEVER}(i, N_P, A, C)$	$\text{NAEVER}(i, N_P, A, C)$
if $(i, N_P, A, C) \in \mathcal{Q}_D$ <b>then return</b> $\perp$	if $(i, N_P, A, C) \in \mathcal{Q}_D$ <b>then return</b> $\perp$
$(M, N_S) \leftarrow \text{AEAD.Dec}(K_i, N_P, A, C)$	<b>return false</b>
if $(M, N_S) \neq \perp$ <b>then return true</b>	
<b>else return false</b>	

Figure 4.1: Games  $(\text{NAE}^{\text{real}}, \text{NAE}^{\text{rand}})$  for an AEAD scheme [16].

NAEENC with the same input twice, and it may not query NAEVER on an output of NAEENC. The latter condition is enforced with the set  $\mathcal{Q}_D$ . In addition, we require adversaries to be *nonce-respecting*—for a given user  $i$ , it may not query NAEENC with the same nonce  $(N_P, N_S)$ . This is enforced with the set  $\mathcal{Q}_N$ .

We say that an adversary is *orderly* [35, 19, 14, 16] if its NAEVER queries are made after all the NAEENC queries, all the NAEVER queries are made at once (i.e., NAEVER queries may depend on NAEENC queries but not on other NAEVER queries), and it returns **true** if any of the NAEVER queries returns **true**. We find it useful to restrict to orderly adversaries and the following lemma of [16] upper bounds the advantage loss of such a restriction.

**Lemma 20** (Lemma 3.2 in [16]). *Let AEAD be an AEAD scheme with constant ciphertext overhead  $\tau$ . Then for any NAE adversary  $\mathcal{A}$  making  $q_b$  NAEVER queries, we can construct an orderly adversary  $\mathcal{B}$  such that  $\text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{A}) \leq \text{Adv}_{\text{AEAD}}^{\text{nae}}(\mathcal{B}) + \frac{q_b}{2^\tau}$ . Adversary  $\mathcal{B}$  has the same query complexity as  $\mathcal{A}$  and similar time complexity.*

<div style="border-bottom: 1px solid black; margin-bottom: 5px;">Game CMT(<math>\mathcal{A}</math>)</div> <div style="padding: 5px;"> <math>((K_1, N_{P_1}, A_1), (K_2, N_{P_2}, A_2), C) \leftarrow \mathcal{A}</math>  <math>(M_1, N_{S_1}) \leftarrow \text{AEAD.Dec}(K_1, N_{P_1}, A_1, C)</math>  <math>(M_2, N_{S_2}) \leftarrow \text{AEAD.Dec}(K_2, N_{P_2}, A_2, C)</math>  <b>if</b> <math>(M_1, N_{S_1}) = \perp</math> <b>or</b> <math>(M_2, N_{S_2}) = \perp</math> <b>then return false</b>  <b>return</b> <math>(K_1, N_{P_1}, A_1) \neq (K_2, N_{P_2}, A_2)</math> </div>
---

Figure 4.2: Game CMT for an AEAD scheme [14, 39].

**Committing security.** We focus on the strongest commitment notion *context commitment* [14, 39] which requires the ciphertext to commit to the entire decryption context. In Figure 4.2 we define security. The outputs  $(K_1, N_{P_1}, A_1), (K_2, N_{P_2}, A_2)$  of  $\mathcal{A}$  are required to be in  $\mathcal{K} \times \mathcal{N}_{\mathcal{P}} \times \mathcal{A}$ . We define the following advantage measure

$$\text{Adv}_{\text{AEAD}}^{\text{cmt}}(\mathcal{A}) := \Pr[\text{CMT}(\mathcal{A}) \Rightarrow \text{true}].$$

## 4.2 The Flex Framework

We start with new definitions for AEAD. These build off a variety of prior work, as we explain further below, but weave in new aspects that we think will benefit secure, flexible deployability and usage.

**Elements of F-AEAD.** Recall that in classical AEAD [108, 103], encryption  $C \leftarrow \text{Enc}(K, N, A, M)$  takes key, nonce, AD and message to return a ciphertext, while decryption  $M \leftarrow \text{Dec}(K, N, A, C)$  takes key, nonce, AD and ciphertext to return the message. Security can be basic (unique-nonce) [108, 103], requiring that encryption under a key not repeat a nonce, or advanced (misuse resistant) [106], requiring only that it not repeat nonce, AD and message.

We provide a slightly different and more general formalization called flexible AEAD

(F-AEAD). A novel element of our framework is the introduction of an extra initialization algorithm `Init` which chooses the appropriate AEAD algorithm based on the parameters specified in a configuration input `cfg`. For instance, amongst other things, `cfg` includes a flag `cfg.mr` that says whether or not misuse resistance is requested. This allows a user to dynamically make this choice in different settings without changing the key.

(**sanketh:** I got rid of allowable key lengths in `FAE.Kg`, it was unnecessary complexity.)

(**sanketh:** In this definition, like previous ones, I am assuming that a configuration maps to an explicit AEAD scheme. This is necessary for analysis. For ARF, I will add a new function that takes constraints and outputs a configuration.)

(**sanketh:** Configurations now have associated spaces, not associated lengths. This resolves julia's remark about how some ciphertexts are tuples, for example a key label and a ciphertext.)

(**sanketh:** This definition supports schemes with key identification, but that only works for subkeys. To identify mainkeys, we need to extend this syntax. I wanna discuss this before adding.)

(**sanketh:** This definition uses the AE5 syntax that we used in OCH.)

(**sanketh:** I updated the games to add adaptive corruptions similar to [15], I haven't updated the proofs yet.)

**F-AEAD syntax.** Formally, a flexible AEAD (F-AEAD) scheme `FAE` specifies several algorithms and associated quantities, as follows:

- $K_* \leftarrow \$ \text{FAE.Kg}()$ : The randomized key generation algorithm takes no inputs and outputs a main key  $K_* \in \text{FAE.K}_*$ . The scheme specifies a main key space  $\text{FAE.K}_*$  (later we will set  $\text{FAE.K}_* = \{0, 1\}^{2\kappa}$  where  $\kappa$  is the security parameter.)
- $L = (\text{cfg}, \text{key}) \leftarrow \text{FAE.Init}(\text{cfg}, K_*, S)$ : The deterministic initialization algorithm initializes a session. It takes as input a configuration  $\text{cfg} \in \text{FAE.Cfgs}$ , a main key  $K_* \in \text{FAE.K}_*$ , and some session associated data  $S \in \text{FAE.S}$ . The scheme specifies a set  $\text{FAE.Cfgs}$  of supported configurations  $\text{cfg}$  with an associated key space  $\text{cfg.K} \subseteq \{0, 1\}^*$ , public nonce space  $\text{cfg.N}_p \subseteq \{0, 1\}^*$ , secret nonce space  $\text{cfg.N}_s \subseteq \{0, 1\}^*$ , associated data space  $\text{cfg.A} \subseteq \{0, 1\}^*$ , message space  $\text{cfg.M} \subseteq \{0, 1\}^*$ , and ciphertext space  $\text{cfg.C} \subseteq \{0, 1\}^*$ . The scheme specifies an associated data space  $\text{FAE.S}$ . The algorithm outputs a (secret) subkey  $L = (\text{cfg}, \text{key})$ , where the first component  $L.\text{cfg} \in \text{FAE.Cfgs}$  is the provided configuration and the second component  $L.\text{key} \in \text{cfg.K}$  is a (secret) key for  $L.\text{cfg}$ , or it outputs the distinguished error symbol  $\perp$ .
- $C \leftarrow \text{FAE.Enc}(L, N_p, N_s, A, M)$ : The deterministic encryption algorithm takes a subkey  $L = (\text{cfg}, \text{key})$  with  $L.\text{cfg} \in \text{FAE.Cfgs}$  and  $L.\text{key} \in L.\text{cfg.K}$ , public nonce  $N \in \text{cfg.N}_p$ , secret nonce  $N_s \in \text{cfg.N}_s$ , associated data  $A \in \text{cfg.A}$ , and message  $M \in \text{cfg.M}$ , and returns a ciphertext  $C \in \text{cfg.C}$ .
- $(N_s, M) \leftarrow \text{FAE.Dec}(L, N_p, A, C)$ : The decryption algorithm takes a subkey  $L = (\text{cfg}, \text{key})$  with  $L.\text{cfg} \in \text{FAE.Cfgs}$  and  $L.\text{key} \in L.\text{cfg.K}$ , public nonce  $N_p \in \text{cfg.N}_p$ , associated data  $A \in \text{cfg.A}$ , and ciphertext  $C \in \text{cfg.C}$ , and returns a secret nonce and message  $(N_s, M) \in \text{cfg.N}_s \times \text{cfg.M}$ , or an error  $\perp$ .

We proceed to some remarks, explanations, and choices.



UNAE.InvalidEnc( $(N_P, N_S, A, M), Q_E$ )	MRAE.InvalidEnc( $(N_P, N_S, A, M), Q_E$ )
<b>return</b> $(N_P, N_S, *, *) \notin Q_E$	<b>return</b> $(N_P, N_S, A, M) \notin Q_E$
UNAE.InvalidDec( $(N_P, A, C), Q_D$ )	MRAE.InvalidDec( $(N_P, A, C), Q_D$ )
<b>return</b> $(N_P, A, C) \in Q_D$	<b>return</b> $(N_P, A, C) \in Q_D$

Figure 4.3: The functions InvalidEnc and InvalidDec for UNAE and MRAE schemes. We use  $*$  to denote matching to anything.

**Configurations.** Configurations specify an operating mode, and the list of supported configurations FAE.Cfgs can be extended. We assume configurations are uniquely identifiable and require every configuration  $\text{cfg}$  to specify the following spaces: a key space  $\text{cfg}.\mathcal{K} \subseteq \{0, 1\}^*$ , a public nonce space  $\text{cfg}.\mathcal{N}_P \subseteq \{0, 1\}^*$ , a secret nonce space  $\text{cfg}.\mathcal{N}_S \subseteq \{0, 1\}^*$ , an associated data space  $\text{cfg}.\mathcal{A} \subseteq \{0, 1\}^*$ , a message space  $\text{cfg}.\mathcal{M} \subseteq \{0, 1\}^*$ , and a ciphertext space  $\text{cfg}.\mathcal{C} \subseteq \{0, 1\}^*$ . We also require a configuration  $\text{cfg}$  to specify two predicates  $\text{cfg}.\text{InvalidEnc}$  and  $\text{cfg}.\text{InvalidDec}$  as follows.

- $\text{cfg}.\text{InvalidEnc} : (\text{cfg}.\mathcal{N}_P \times \text{cfg}.\mathcal{N}_S \times \text{cfg}.\mathcal{A} \times \text{cfg}.\mathcal{M}) \times (\text{cfg}.\mathcal{N}_P \times \text{cfg}.\mathcal{N}_S \times \text{cfg}.\mathcal{A} \times \text{cfg}.\mathcal{M})^* \rightarrow \{\text{true}, \text{false}\}$  takes an encryption input and the list of previous encryption inputs to this configuration with the same key, returns whether the input is invalid.
- $\text{cfg}.\text{InvalidDec} : (\text{cfg}.\mathcal{N}_P \times \text{cfg}.\mathcal{A} \times \text{cfg}.\mathcal{C}) \times (\text{cfg}.\mathcal{N}_P \times \text{cfg}.\mathcal{A} \times \text{cfg}.\mathcal{C})^* \rightarrow \{\text{true}, \text{false}\}$ : takes a decryption input and the list of previous encryption outputs to this configuration with the same key, returns whether the input is invalid.

These predicates capture constraints on the adversary like not decrypting encryption outputs and nonce-uniqueness, and we define them for unique-nonce (UNAE) and nonce-misuse resistant (MRAE) schemes in Figure 4.3.

**Validity of inputs.** All specified algorithms return  $\perp$  if any of the inputs are not valid. For example, if  $\text{FAE.Init}$  is called with an unsupported configuration  $\text{cfg} \notin \text{FAE.Cfgs}$ , if  $\text{FAE.Enc}$  is called with an inconsistent subkey  $L = (\text{cfg}, \text{key})$  that has  $L.\text{key} \notin L.\text{cfg}.\mathcal{K}$ , or if  $\text{FAE.Dec}$  is called with an unsupported secret nonce  $N_S \notin L.\text{cfg}.\mathcal{N}_S$ .

**Correctness and tidyness.** We also extend and impose correctness and *tidiness* [95] requirements on F-AEAD. These require that  $\text{FAE.Enc}$  and  $\text{FAE.Dec}$  to be inverses of each other for *honest subkeys*. In detail, for any  $K_*$ ,  $\text{cfg}$ ,  $S$ ,  $N_P$ ,  $N_S$ ,  $A$ ,  $M$ , and  $C$ : let  $L = (\text{cfg}, \text{key}) \leftarrow \text{FAE.Init}(\text{cfg}, K_*, S)$  then

$$\text{FAE.Enc}(L, N_P, N_S, A, M) = C \implies \text{FAE.Dec}(L, N_P, A, C) = (N_S, M),$$

$$\text{FAE.Dec}(L, N_P, A, C) = (N_S, M) \neq \perp \implies \text{FAE.Enc}(L, N_P, N_S, A, M) = C$$

**F-AEAD confidentiality and integrity.** We formalize an all-in-one, multi-user notion of security for F-AEAD that captures confidentiality and ciphertext integrity. A notable aspect of our definition is that it allows adversaries to obtain ciphertexts encrypted under the same key but different configurations, and thus in particular allows moving from non-misuse-resistant to misuse-resistant security dynamically and without key change.

We formalize this in the multi-user setting using a standard games-based approach with a pair of games  $\text{FNAE}^{\text{real}}$  and  $\text{FNAE}^{\text{rand}}$  which are defined in Figure 4.5, and then set

$$\text{Adv}_{\text{FAE}}^{\text{fnae}}(\mathcal{A}) := \Pr[\text{FNAE}^{\text{real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{FNAE}^{\text{rand}}(\mathcal{A}) \Rightarrow 1].$$

**F-AEAD context commitment.**

We formalize a notion of context commitment for Flexible AEAD, extending context commitment [14, 39] to incorporate configurations, using a game FCMT defined in Figure 4.6 as

$$\text{Adv}_{\text{FAE}}^{\text{fcmt}}(\mathcal{A}) := \Pr[\text{FCMT}(\mathcal{A}) \Rightarrow \text{true}].$$

<p><b>Game</b> <math>\text{FNAE}_{\text{FAE}}^{\text{real}}(\mathcal{A})</math>:</p> <p><math>u \leftarrow 0</math></p> <p><math>\tau_L \leftarrow \{\}</math></p> <p><math>b \leftarrow \\$ \mathcal{A}^{\text{FNAEKg}, \text{FNAEEnc}, \text{FNAEVER}}</math></p> <p><b>return</b> <math>b</math></p> <p><u>FNAEKg():</u></p> <p><u>FNAEINIT(<math>i, \text{cfg}, S</math>):</u></p> <p><u>FNAEEnc(<math>i, j, N_P, N_S, A, M</math>):</u></p> <p><u>FNAEVER(<math>i, j, N_P, A, C</math>):</u></p>	<p><b>Game</b> <math>\text{FNAE}_{\text{FAE}}^{\text{rand}}(\mathcal{A})</math>:</p> <p><u>FNAEKg():</u></p> <p><u>FNAEINIT(<math>i, \text{cfg}, S</math>):</u></p> <p><u>FNAEEnc(<math>i, j, N_P, N_S, A, M</math>):</u></p> <p><u>FNAEVER(<math>i, j, N_P, A, C</math>):</u></p>
---	--

Figure 4.4: Games  $(\text{FNAE}^{\text{real}}, \text{FNAE}^{\text{rand}})$  for a flexible AEAD.

Game $\text{FNAE}_{\text{FAE}}^{\text{real}}(\mathcal{A})$	Game $\text{FNAE}_{\text{FAE}}^{\text{rand}}(\mathcal{A})$
$u \leftarrow 0; \tau_L \leftarrow \{\}; \tau_E \leftarrow \{\}; \tau_C \leftarrow \{\}$ $b \leftarrow \mathcal{A}^{\text{FNAEKg, FNAEEnc, FNAEVer}}$ <b>return</b> $b$ <hr/> $\text{FNAEKg}()$	$u \leftarrow 0; \tau_L \leftarrow \{\}; \tau_E \leftarrow \{\}; \tau_C \leftarrow \{\}$ $b \leftarrow \mathcal{A}^{\text{FNAEKg, FNAEEnc, FNAEVer}}$ <b>return</b> $b$ <hr/> $\text{FNAEKg}()$
$u \leftarrow u + 1; v_u \leftarrow 0; K_u \leftarrow \text{FAE.K}_s$ $\tau_L[u] \leftarrow \{\}; \tau_E[u] \leftarrow \{\}; \tau_C[u] \leftarrow \{\}$ <hr/> $\text{FNAEInit}(i, \text{cfg}, S)$	$u \leftarrow u + 1; v_u \leftarrow 0;$ $\tau_L[u] \leftarrow \{\}; \tau_E[u] \leftarrow \{\}; \tau_C[u] \leftarrow \{\}$ <hr/> $\text{FNAEInit}(i, \text{cfg}, S)$
$v_i \leftarrow v_i + 1$ $\tau_L[i][v_i] \leftarrow \text{FAE.Init}(\text{cfg}, K_i, S)$ <hr/> $\text{FNAEEnc}(i, j, N_p, N_s, A, M)$	$v_i \leftarrow v_i + 1; L \leftarrow \text{cfg.K}$ $\tau_L[i][v_i] \leftarrow (\text{cfg}, L)$ <hr/> $\text{FNAEEnc}(i, N_p, N_s, A, M)$
<b>if</b> $\tau_L[i][j].\text{cfg.InvalidEnc}((N_p, N_s, A, M), Q_E[i][j])$ <b>return</b> $\perp$ $C \leftarrow \text{FAE.Enc}(\tau_L[i][j], N_p, N_s, A, M)$ $Q_E[i][j] \leftarrow Q_E[i][j] \cup (N_p, N_s, A, M)$ $Q_D[i][j] \leftarrow Q_D[i][j] \cup (N_p, A, C)$ <b>return</b> $C$ <hr/> $\text{FNAEVer}(i, j, N_p, A, C)$	<b>if</b> $\tau_L[i][j].\text{cfg.InvalidEnc}((N_p, N_s, A, M), Q_E[i][j])$ <b>return</b> $\perp$ $C \leftarrow \tau_L[i][j].\text{cfg.C}$ $Q_E[i][j] \leftarrow Q_E[i][j] \cup (N_p, N_s, A, M)$ $Q_D[i][j] \leftarrow Q_D[i][j] \cup (N_p, A, C)$ <b>return</b> $C$ <hr/> $\text{FNAEVer}(i, j, N_p, A, C)$
<b>if</b> $\tau_L[i][j].\text{cfg.InvalidDec}((N_p, A, C), Q_D[i][j])$ <b>return</b> $\perp$ $(M, N_s) \leftarrow \text{AEAD.Dec}(K_i, N_p, A, C)$ <b>if</b> $(M, N_s) \neq \perp$ <b>then return true</b> <b>else return false</b> <hr/> $\text{FNAEExp}(i, j)$	<b>if</b> $\tau_L[i][j].\text{cfg.InvalidDec}((N_p, A, C), Q_D[i][j])$ <b>return</b> $\perp$ <b>return false</b> <hr/> $\text{FNAEExp}(i, j)$
<b>if</b> $j \in \tau_C[i]$ <b>then return</b> $\perp$ $\tau_E[i] \leftarrow \tau_E[i] \cup \{j\}$ <b>return</b> $\tau_L[i][j]$ <hr/> $\text{FNAECha}(i, j)$	<b>if</b> $j \in \tau_C[i]$ <b>then return</b> $\perp$ $\tau_E[i] \leftarrow \tau_E[i] \cup \{j\}$ <b>return</b> $\tau_L[i][j]$ <hr/> $\text{FNAECha}(i, j)$
<b>if</b> $j \in \tau_E[i]$ <b>then return</b> $\perp$ $\tau_C[i] \leftarrow \tau_C[i] \cup \{j\}$	<b>if</b> $j \in \tau_E[i]$ <b>then return</b> $\perp$ $\tau_C[i] \leftarrow \tau_C[i] \cup \{j\}$

Figure 4.5: Games  $(\text{FNAE}^{\text{real}}, \text{FNAE}^{\text{rand}})$  for a flexible AEAD.

Game $\text{FCMT}(\mathcal{A})$
$((K_1, \text{cfg}_1, S_1, N_{p1}, A_1), (K_2, \text{cfg}_2, S_2, N_{p2}, A_2), C) \leftarrow \mathcal{A}$ $L_1 \leftarrow \text{FAE.Init}(\text{cfg}_1, K_1, S_1)$ $L_2 \leftarrow \text{FAE.Init}(\text{cfg}_2, K_2, S_2)$ $(M_1, N_{s1}) \leftarrow \text{FAE.Dec}(L_1, N_{p1}, A_1, C)$ $(M_2, N_{s2}) \leftarrow \text{FAE.Dec}(L_2, N_{p2}, A_2, C)$ <b>if</b> $(M_1, N_{s1}) = \perp$ <b>or</b> $(M_2, N_{s2}) = \perp$ <b>then return false</b> <b>return</b> $(K_1, \text{cfg}_1, S_1, N_{p1}, A_1) \neq (K_2, \text{cfg}_2, S_2, N_{p2}, A_2)$

Figure 4.6: Game FCMT for a flexible AEAD.

### 4.3 ARF: An Instantiation of Flexible AEAD

We propose an F-AEAD called ARF.

- ARF *automates* the selection AEADs. It bakes in what is traditionally left to users of AEAD, like the choice of AEAD class (like misuse-resistance versus not) and the choice of AEAD scheme (like ChaCha20/Poly1305 or AES-GCM).
- ARF *reuses* components across different AEAD choices. This modular design minimizes code and analysis complexity.
- ARF *flexibly* adapts to changing application requirements. It allows changing the AEAD scheme, while retaining the API and without need for key rotation.

**Overview.** We build ARF modularly. At a high level there are four main components all built from the same underlying permutation  $\Pi$ :

- A key derivation function (KDF) that achieves CR-PRF security. It is used to derive a subkey from the secret subkey, associated data, and configuration. By default this is a simple sponge-style hash construction, but we also give modes that allow parallelization and precomputation of static elements of the input, such as the associated data .
- A non-misuse resistant encryption mode called OCS that is an OCB-like mode that provides context committing security and nonce-hiding while being fully parallelizable.
- A misuse resistant encryption mode called CIV that is an SIV-like mode that provides context committing security and nonce-hiding while being maximally parallelizable.

- A robust encryption mode called CEC (committing encipher with counter) that is an HCTR2-like mode that provides context committing security and robustness while being maximally parallelizable.

All components are new to this work, but of course use some design elements seen in prior work as we will highlight where relevant. We can then mix-and-match the different KDF modes with the different encryption modes to provide various designs, which enables a wide variety of AEAD schemes. At this level our individual constructions follow the Hash-then-Encrypt mechanism from Bellare and Hoang [14] that transforms a key-committing AEAD into a context-committing AEAD. But we carefully arrange our constructions, down to the level of using a fast-to-rekey specially constructed Even-Mansour [56] cipher, to avoid inefficiencies.

In Section ?? we describe how we build an API for ARF that will allow safe use of this wide variety of cryptographic designs. Indeed our security analysis will show that ARF enjoys F-AEAD security, and the modularity of our construction makes showing it rather straightforward despite reuse of underlying primitives (such as  $\Pi$ ). A key technical lemma (Lemma ??) will show that security composes across different configurations. We refer to this as the configuration separation lemma, as it allows us to reduce to separately analyzing the security of different configurations and, in turn, different components.

### 4.3.1 ARF Overview

We provide a top-down description of ARF. It supports up to 256 different configurations, allowing `cfg` to be represented by one byte. We have defined a number of initial configurations (see Figure 4.7), and expect one might extend ARF with further

Description	RN or IN?	MR?	Notes
Short, Stateful	IN	Y	96-bit stateful nonces, $ M  < n - 96$
Short, Randomized	RN	Y	256-bit random nonces, $ M  < n$
Long Messages, Stateful, NMR	IN	N	96-bit stateful nonces, $ M  \geq n$
Long Messages, Randomized, NMR	RN	N	256-bit randomized nonces, $ M  \geq n$
Long Messages, Randomized, MR	RN	N	256-bit randomized nonces, $ M  \geq n$

Figure 4.7: Starting configurations for ARF. All configurations using stateful, incremental nonces are nonce-hiding; configurations with randomized nonces dispense with nonce hiding. Short messages can always be handled using a specialized MR mode. We allow for future extensions without security impacting other extensions, via our configuration separation lemma (Lemma ??).

configurations as one desires new modes.

Let  $\kappa$  be the security parameter, later we will set this to 128.

(sanketh: stopped here, define keyspace, etc.)

We define ARF as follows:

- $K_* \leftarrow \text{ARF.Kg}()$ : The randomized key generation algorithm takes no inputs and outputs a main key  $K_* \in \text{FAE}.\mathcal{K}_*$ . The scheme specifies a main key space  $\text{FAE}.\mathcal{K}_*$  (later we will set  $\text{FAE}.\mathcal{K}_* = \{0, 1\}^{2\kappa}$  where  $\kappa$  is the security parameter.)
- $L = (\text{cfg}, \text{key}) \leftarrow \text{FAE.Init}(\text{cfg}, K_*, S)$ : The deterministic initialization algorithm initializes a session. It takes as input a configuration  $\text{cfg} \in \text{FAE.Cfgs}$ , a main key  $K_* \in \text{FAE}.\mathcal{K}_*$ , and some session associated data  $S \in \text{FAE}.\mathcal{S}$ . The scheme specifies a set  $\text{FAE.Cfgs}$  of supported configurations  $\text{cfg}$  with an associated key space  $\text{cfg}.\mathcal{K} \subseteq \{0, 1\}^*$ , public nonce space  $\text{cfg}.\mathcal{N}_p \subseteq \{0, 1\}^*$ , secret nonce space  $\text{cfg}.\mathcal{N}_s \subseteq \{0, 1\}^*$ , associated data space  $\text{cfg}.\mathcal{A} \subseteq \{0, 1\}^*$ , message space  $\text{cfg}.\mathcal{M} \subseteq \{0, 1\}^*$ , and ciphertext space  $\text{cfg}.\mathcal{C} \subseteq \{0, 1\}^*$ . The scheme specifies an as-



sociated data space  $\text{FAE.S}$ . The algorithm outputs a (secret) subkey  $L = (\text{cfg}, \text{key})$ , where the first component  $L.\text{cfg} \in \text{FAE.Cfgs}$  is the provided configuration and the second component  $L.\text{key} \in \text{cfg.K}$  is a (secret) key for  $L.\text{cfg}$ , or it outputs the distinguished error symbol  $\perp$ .

- $C \leftarrow \text{FAE.Enc}(L, N_p, N_s, A, M)$ : The deterministic encryption algorithm takes a subkey  $L = (\text{cfg}, \text{key})$  with  $L.\text{cfg} \in \text{FAE.Cfgs}$  and  $L.\text{key} \in L.\text{cfg.K}$ , public nonce  $N \in \text{cfg.N}_p$ , secret nonce  $N_s \in \text{cfg.N}_s$ , associated data  $A \in \text{cfg.A}$ , and message  $M \in \text{cfg.M}$ , and returns a ciphertext  $C \in \text{cfg.C}$ .
- $(N_s, M) \leftarrow \text{FAE.Dec}(L, N_p, A, C)$ : The decryption algorithm takes a subkey  $L = (\text{cfg}, \text{key})$  with  $L.\text{cfg} \in \text{FAE.Cfgs}$  and  $L.\text{key} \in L.\text{cfg.K}$ , public nonce  $N_p \in \text{cfg.N}_p$ , associated data  $A \in \text{cfg.A}$ , and ciphertext  $C \in \text{cfg.C}$ , and returns a secret nonce and message  $(N_s, M) \in \text{cfg.N}_s \times \text{cfg.M}$ , or an error  $\perp$ .

#### 4.4 Security Analysis of ARF

#### 4.5 Performance

## BIBLIOGRAPHY

- [1] Michel Abdalla, Mihir Bellare, and Gregory Neven. “Robust Encryption”. In: *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*. Ed. by Daniele Micciancio. Vol. 5978. Lecture Notes in Computer Science. Springer, 2010, pp. 480–497. doi: [10.1007/978-3-642-11799-2\\_28](https://doi.org/10.1007/978-3-642-11799-2_28).
- [2] Ange Albertini et al. *How to Abuse and Fix Authenticated Encryption Without Key Commitment*. USENIX Security 2022. 2022. URL: <https://ia.cr/2020/1456>.
- [3] Ange Albertini et al. “How to Abuse and Fix Authenticated Encryption Without Key Commitment”. In: *USENIX Security 2022*. Ed. by Kevin R. B. Butler and Kurt Thomas. USENIX Association, Aug. 2022, pp. 3291–3308. URL: <https://www.usenix.org/conference/usenixsecurity22/presentation/albertini>.
- [4] Moreno Ambrosin et al. *Tink*. 2021. URL: <https://github.com/google/tink/releases/tag/v1.6.1>.
- [5] Moreno Ambrosin et al. *Tink EAX Key Manager*. 2021. URL: [https://github.com/google/tink/blob/v1.6.1/java\\_src/src/main/java/com/google/crypto/tink/ae/AesEaxKeyManager.java#L115-L116](https://github.com/google/tink/blob/v1.6.1/java_src/src/main/java/com/google/crypto/tink/ae/AesEaxKeyManager.java#L115-L116).
- [6] Elena Andreeva et al. “Security of Keyed Sponge Constructions Using a Modular Proof Approach”. In: *FSE 2015*. Ed. by Gregor Leander. Vol. 9054. LNCS. Springer, Berlin, Heidelberg, Mar. 2015, pp. 364–384. doi: [10.1007/978-3-662-48116-5\\_18](https://doi.org/10.1007/978-3-662-48116-5_18).
- [7] Scott Arciszewski. *Galois Extended Mode*. NIST Workshop on the Requirements for an Accordion Cipher Mode. 2024. URL: <https://csrc.nist.gov/>

[csrc/media/Events/2024/accordion-cipher-mode-workshop-2024/documents/papers/galois-extended-mode.pdf](https://csrc.media/Events/2024/accordion-cipher-mode-workshop-2024/documents/papers/galois-extended-mode.pdf).

- [8] Scott Arciszewski. *XChaCha: eXtended-nonce ChaCha and AEAD\_XChaCha20\_Poly1305*. Internet-Draft draft-irtf-cfrg-xchacha-03. Work in Progress. Internet Engineering Task Force, Jan. 2020. 18 pp. URL: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-xchacha/03/>.
- [9] Jean-Philippe Aumasson et al. “BLAKE2: Simpler, Smaller, Fast as MD5”. In: *ACNS 2013*. Ed. by Michael J. Jacobson Jr. et al. Vol. 7954. LNCS. Springer, Berlin, Heidelberg, June 2013, pp. 119–135. DOI: [10.1007/978-3-642-38980-1\\_8](https://doi.org/10.1007/978-3-642-38980-1_8).
- [10] AWS. *Supported algorithm suites in the AWS Encryption SDK*. Accessed 2021-09-23. 2021. URL: <https://docs.aws.amazon.com/encryption-sdk/latest/developer-guide/supported-algorithms.html>.
- [11] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. “Pseudorandom Functions Revisited: The Cascade Construction and Its Concrete Security”. In: *37th FOCS*. IEEE Computer Society Press, Oct. 1996, pp. 514–523. DOI: [10.1109/SFCS.1996.548510](https://doi.org/10.1109/SFCS.1996.548510).
- [12] Mihir Bellare, Hannah Davis, and Felix Günther. “Separate Your Domains: NIST PQC KEMs, Oracle Cloning and Read-Only Indifferentiability”. In: *EUROCRYPT 2020, Part II*. Ed. by Anne Canteaut and Yuval Ishai. Vol. 12106. LNCS. Springer, Cham, May 2020, pp. 3–32. DOI: [10.1007/978-3-030-45724-2\\_1](https://doi.org/10.1007/978-3-030-45724-2_1).
- [13] Mihir Bellare and Viet Tung Hoang. “Efficient Schemes for Committing Authenticated Encryption”. In: *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part II*. Ed. by Orr Dunkelman and Stefan Dziembowski. Vol. 13276. Lecture Notes in Com-

- puter Science. Springer, 2022, pp. 845–875. doi: [10.1007/978-3-031-07085-3\\_29](https://doi.org/10.1007/978-3-031-07085-3_29).
- [14] Mihir Bellare and Viet Tung Hoang. “Efficient Schemes for Committing Authenticated Encryption”. In: *EUROCRYPT 2022, Part II*. Ed. by Orr Dunkelman and Stefan Dziembowski. Vol. 13276. LNCS. Springer, Cham, May 2022, pp. 845–875. doi: [10.1007/978-3-031-07085-3\\_29](https://doi.org/10.1007/978-3-031-07085-3_29).
- [15] Mihir Bellare and Viet Tung Hoang. “Identity-Based Format-Preserving Encryption”. In: *ACM CCS 2017*. Ed. by Bhavani M. Thuraisingham et al. ACM Press, Oct. 2017, pp. 1515–1532. doi: [10.1145/3133956.3133995](https://doi.org/10.1145/3133956.3133995).
- [16] Mihir Bellare and Viet Tung Hoang. “Succinctly-Committing Authenticated Encryption”. In: *CRYPTO 2024, Part IV*. Ed. by Leonid Reyzin and Douglas Stebila. Vol. 14923. LNCS. Springer, Cham, Aug. 2024, pp. 305–339. doi: [10.1007/978-3-031-68385-5\\_10](https://doi.org/10.1007/978-3-031-68385-5_10).
- [17] Mihir Bellare, Viet Tung Hoang, and Cong Wu. *The Landscape of Committing Authenticated Encryption*. The Third NIST Workshop on Block Cipher Modes of Operation. 2023. URL: <https://csrc.nist.gov/Presentations/2023/landscape-of-committing-authenticated-encryption>.
- [18] Mihir Bellare and Chanathip Namprempre. “Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm”. In: *J. Cryptol.* 21.4 (2008), pp. 469–491. doi: [10.1007/s00145-008-9026-x](https://doi.org/10.1007/s00145-008-9026-x).
- [19] Mihir Bellare, Ruth Ng, and Björn Tackmann. “Nonces Are Noticed: AEAD Revisited”. In: *CRYPTO 2019, Part I*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Vol. 11692. LNCS. Springer, Cham, Aug. 2019, pp. 235–265. doi: [10.1007/978-3-030-26948-7\\_9](https://doi.org/10.1007/978-3-030-26948-7_9).

- [20] Mihir Bellare and Phillip Rogaway. “Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography”. In: *Advances in Cryptology - ASIACRYPT 2000, 6th International Conference on the Theory and Application of Cryptology and Information Security, Kyoto, Japan, December 3-7, 2000, Proceedings*. Ed. by Tatsuaki Okamoto. Vol. 1976. Lecture Notes in Computer Science. Springer, 2000, pp. 317–330. DOI: [10.1007/3-540-44448-3\\_24](https://doi.org/10.1007/3-540-44448-3_24). URL: <https://cseweb.ucsd.edu/~mihir/papers/ee.pdf>.
- [21] Mihir Bellare and Phillip Rogaway. *Introduction to Modern Cryptography*. 2005. URL: <https://web.cs.ucdavis.edu/~rogaway/classes/227/spring05/book/main.pdf>.
- [22] Mihir Bellare and Phillip Rogaway. “The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs”. In: *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*. Ed. by Serge Vaudenay. Vol. 4004. Lecture Notes in Computer Science. Springer, 2006, pp. 409–426. DOI: [10.1007/11761679\\_25](https://doi.org/10.1007/11761679_25).
- [23] Mihir Bellare and Phillip Rogaway. “The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs”. In: *EUROCRYPT 2006*. Ed. by Serge Vaudenay. Vol. 4004. LNCS. Springer, Berlin, Heidelberg, May 2006, pp. 409–426. DOI: [10.1007/11761679\\_25](https://doi.org/10.1007/11761679_25).
- [24] Mihir Bellare, Phillip Rogaway, and David A. Wagner. “The EAX Mode of Operation”. In: *Fast Software Encryption, 11th International Workshop, FSE 2004, Delhi, India, February 5-7, 2004, Revised Papers*. Ed. by Bimal K. Roy and Willi Meier. Vol. 3017. Lecture Notes in Computer Science. Springer, 2004, pp. 389–407. DOI: [10.1007/978-3-540-25937-4\\_25](https://doi.org/10.1007/978-3-540-25937-4_25). URL: <https://web.cs.ucdavis.edu/~rogaway/papers/eax.pdf>.

- [25] Mihir Bellare and Laura Shea. “Flexible Password-Based Encryption: Securing Cloud Storage and Provably Resisting Partitioning-Oracle Attacks”. In: *CT-RSA 2023*. Ed. by Mike Rosulek. Vol. 13871. LNCS. Springer, Cham, Apr. 2023, pp. 594–621. DOI: [10.1007/978-3-031-30872-7\\_23](https://doi.org/10.1007/978-3-031-30872-7_23).
- [26] Mihir Bellare and Björn Tackmann. “The Multi-user Security of Authenticated Encryption: AES-GCM in TLS 1.3”. In: *CRYPTO 2016, Part I*. Ed. by Matthew Robshaw and Jonathan Katz. Vol. 9814. LNCS. Springer, Berlin, Heidelberg, Aug. 2016, pp. 247–276. DOI: [10.1007/978-3-662-53018-4\\_10](https://doi.org/10.1007/978-3-662-53018-4_10).
- [27] Mihir Bellare et al. “Ask Your Cryptographer if Context-Committing AEAD Is Right for You”. In: *Presented at Real World Crypto Symposium*. 2023.
- [28] Mihir Bellare et al. *Building the Next Generation of AEAD*. Presented at Real World Crypto 2024. 2024.
- [29] Mihir Bellare et al. “Format-Preserving Encryption”. In: *Selected Areas in Cryptography, 16th Annual International Workshop, SAC 2009, Calgary, Alberta, Canada, August 13-14, 2009, Revised Selected Papers*. Ed. by Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini. Vol. 5867. Lecture Notes in Computer Science. Springer, 2009, pp. 295–312. DOI: [10.1007/978-3-642-05445-7\\_19](https://doi.org/10.1007/978-3-642-05445-7_19).
- [30] Guido Bertoni et al. “Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications”. In: *Selected Areas in Cryptography - 18th International Workshop, SAC 2011, Toronto, ON, Canada, August 11-12, 2011, Revised Selected Papers*. Ed. by Ali Miri and Serge Vaudenay. Vol. 7118. Lecture Notes in Computer Science. Springer, 2011, pp. 320–337. DOI: [10.1007/978-3-642-28496-0\\_19](https://doi.org/10.1007/978-3-642-28496-0_19).

- [31] Guido Bertoni et al. “Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications”. In: *SAC 2011*. Ed. by Ali Miri and Serge Vaude-  
nay. Vol. 7118. LNCS. Springer, Berlin, Heidelberg, Aug. 2012, pp. 320–337. doi:  
[10.1007/978-3-642-28496-0\\_19](https://doi.org/10.1007/978-3-642-28496-0_19).
- [32] Guido Bertoni et al. “Sponge functions”. In: *ECRYPT hash workshop*. Vol. 2007.  
2007.
- [33] John Black, Phillip Rogaway, and Thomas Shrimpton. “Black-Box Analysis of  
the Block-Cipher-Based Hash-Function Constructions from PGV”. In: *Advances  
in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference,  
Santa Barbara, California, USA, August 18-22, 2002, Proceedings*. Ed. by Moti  
Yung. Vol. 2442. Lecture Notes in Computer Science. Springer, 2002, pp. 320–  
335. DOI: [10.1007/3-540-45708-9\\_21](https://doi.org/10.1007/3-540-45708-9_21).
- [34] BoringSSL Authors. *speed.cc*. 2025. URL: [https://github.com/google/  
boringssl/blob/e056f59c7dfdcf891af03bc7900c946ac485c78f/tool/  
speed.cc](https://github.com/google/boringssl/blob/e056f59c7dfdcf891af03bc7900c946ac485c78f/tool/speed.cc).
- [35] Priyanka Bose, Viet Tung Hoang, and Stefano Tessaro. “Revisiting AES-GCM-  
SIV: Multi-user Security, Faster Key Derivation, and Better Bounds”. In: *EURO-  
CRYPT 2018, Part I*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Vol. 10820.  
LNCS. Springer, Cham, Apr. 2018, pp. 468–499. DOI: [10.1007/978-3-319-  
78381-9\\_18](https://doi.org/10.1007/978-3-319-78381-9_18).
- [36] CAESAR committee. *CAESAR call for submissions, draft 1*. 2013. URL: [http:  
//competitions.cr.yp.to/caesar-call-1.html](http://competitions.cr.yp.to/caesar-call-1.html).
- [37] John Chan and Phillip Rogaway. “Anonymous AE”. In: *ASIACRYPT 2019, Part II*.  
Ed. by Steven D. Galbraith and Shiho Moriai. Vol. 11922. LNCS. Springer, Cham,  
Dec. 2019, pp. 183–208. DOI: [10.1007/978-3-030-34621-8\\_7](https://doi.org/10.1007/978-3-030-34621-8_7).

- [38] John Chan and Phillip Rogaway. “On Committing Authenticated-Encryption”. In: *European Symposium on Research in Computer Security*. Springer. 2022, pp. 275–294. URL: <https://ia.cr/2022/1260>.
- [39] John Chan and Phillip Rogaway. “On Committing Authenticated-Encryption”. In: *ESORICS 2022, Part II*. Ed. by Vijayalakshmi Atluri et al. Vol. 13555. LNCS. Springer, Cham, Sept. 2022, pp. 275–294. DOI: [10.1007/978-3-031-17146-8\\_14](https://doi.org/10.1007/978-3-031-17146-8_14).
- [40] Yu Long Chen et al. *Key Committing Security of AEZ*. The Third NIST Workshop on Block Cipher Modes of Operation. 2023. URL: <https://csrc.nist.gov/Presentations/2023/key-committing-security-of-aez>.
- [41] Benoit Cogliati, Rodolphe Lampe, and Yannick Seurin. “Tweaking Even-Mansour Ciphers”. In: *CRYPTO 2015, Part I*. Ed. by Rosario Gennaro and Matthew J. B. Robshaw. Vol. 9215. LNCS. Springer, Berlin, Heidelberg, Aug. 2015, pp. 189–208. DOI: [10.1007/978-3-662-47989-6\\_9](https://doi.org/10.1007/978-3-662-47989-6_9).
- [42] Valve Corporation. *Steam Hardware & Software Survey: March 2025*. steampowered.com. 2025. URL: <https://web.archive.org/web/20250402020128/https://store.steampowered.com/hwsurvey/Steam-Hardware-Software-Survey-Welcome-to-Steam>.
- [43] Joan Daemen et al. *Shaking up authenticated encryption*. Cryptology ePrint Archive, Report 2024/1618. 2024. URL: <https://eprint.iacr.org/2024/1618>.
- [44] Jean Paul Degabriele et al. “The Security of ChaCha20-Poly1305 in the Multi-User Setting”. In: *ACM CCS 2021*. Ed. by Giovanni Vigna and Elaine Shi. ACM Press, Nov. 2021, pp. 1981–2003. DOI: [10.1145/3460120.3484814](https://doi.org/10.1145/3460120.3484814).



- [45] Frank Denis et al. *AEAD constructions, Robustness*. 2022. URL: [https://doc.libsodium.org/secret-key%5C\\_cryptography/aead#robustness](https://doc.libsodium.org/secret-key%5C_cryptography/aead#robustness).
- [46] Frank Denis and Samuel Lucas. *The AEGIS Family of Authenticated Encryption Algorithms*. Internet-Draft draft-irtf-cfrg-aegis-aead-16. Work in Progress. Internet Engineering Task Force, Feb. 2025. 73 pp. URL: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-aegis-aead/16/>.
- [47] Christoph Dobraunig et al. *Ascon v1.2*. Submission to NIST. 2021. URL: <https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/finalist-round/updated-spec-doc/ascon-spec-final.pdf>.
- [48] Christoph Dobraunig et al. “Ascon v1.2: Lightweight Authenticated Encryption and Hashing”. In: *Journal of Cryptology* 34.3 (July 2021), p. 33. DOI: [10.1007/s00145-021-09398-9](https://doi.org/10.1007/s00145-021-09398-9).
- [49] Christoph Dobraunig et al. *Ascon v1.2: Submission to NIST*. May 2021. URL: <https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/finalist-round/updated-spec-doc/ascon-spec-final.pdf>.
- [50] Christoph Dobraunig et al. *Schwaemm and Esch: Lightweight Authenticated Encryption and Hashing using the Sparkle Permutation Family v1.1*. Submission to NIST. 2019. URL: <https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/round-2/spec-doc-rnd2/sparkle-spec-round2.pdf>.
- [51] Yevgeniy Dodis et al. “Fast Message Franking: From Invisible Salamanders to Encryptment”. In: *CRYPTO 2018, Part I*. Ed. by Hovav Shacham and Alexandra Boldyreva. Vol. 10991. LNCS. Springer, Cham, Aug. 2018, pp. 155–186. DOI: [10.1007/978-3-319-96884-1\\_6](https://doi.org/10.1007/978-3-319-96884-1_6).

- [52] Thai “thaidn” Duong. *AWS: In-band protocol negotiation and robustness weaknesses in AWS KMS and Encryption SDKs*. CVE-2020-8897. 2020. URL: <https://github.com/google/security-research/security/advisories/GHSA-wqgp-vphw-hphf>.
- [53] Morris Dworkin. *Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC*. Tech. rep. NIST Special Publication (SP) NIST SP 800-38D. Gaithersburg, MD: National Institute of Standards and Technology, 2007. DOI: [10.6028/NIST.SP.800-38D](https://doi.org/10.6028/NIST.SP.800-38D).
- [54] Morris Dworkin. *Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality*. NIST Special Publication 800-38C. 2004. DOI: [10.6028/NIST.SP.800-38C](https://doi.org/10.6028/NIST.SP.800-38C).
- [55] Morris Dworkin. *Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication*. NIST Special Publication 800-38B. 2005. DOI: [10.6028/NIST.SP.800-38B](https://doi.org/10.6028/NIST.SP.800-38B).
- [56] Shimon Even and Yishay Mansour. “A Construction of a Cipher From a Single Pseudorandom Permutation”. In: *ASIACRYPT’91*. Ed. by Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto. Vol. 739. LNCS. Springer, Berlin, Heidelberg, Nov. 1993, pp. 210–224. DOI: [10.1007/3-540-57332-1\\_17](https://doi.org/10.1007/3-540-57332-1_17).
- [57] Facebook. *Messenger Secret Conversations: Technical Whitepaper*. 2017. URL: <https://about.fb.com/wp-content/uploads/2016/07/messenger-secret-conversations-technical-whitepaper.pdf>.
- [58] Pooya Farshim, Claudio Orlandi, and Razvan Rosie. “Security of Symmetric Primitives under Incorrect Usage of Keys”. In: *IACR Trans. Symmetric Cryptol.* 2017.1 (2017), pp. 449–473. DOI: [10.13154/tosc.v2017.i1.449-473](https://doi.org/10.13154/tosc.v2017.i1.449-473).

- [59] Pooya Farshim, Claudio Orlandi, and Răzvan Roşie. “Security of Symmetric Primitives under Incorrect Usage of Keys”. In: *IACR Trans. Symm. Cryptol.* 2017.1 (2017), pp. 449–473. ISSN: 2519-173X. DOI: [10.13154/tosc.v2017.i1.449-473](https://doi.org/10.13154/tosc.v2017.i1.449-473).
- [60] Pooya Farshim et al. “Robust Encryption, Revisited”. In: *Public-Key Cryptography - PKC 2013 - 16th International Conference on Practice and Theory in Public-Key Cryptography, Nara, Japan, February 26 - March 1, 2013. Proceedings*. Ed. by Kaoru Kurosawa and Goichiro Hanaoka. Vol. 7778. Lecture Notes in Computer Science. Springer, 2013, pp. 352–368. DOI: [10.1007/978-3-642-36362-7\\_22](https://doi.org/10.1007/978-3-642-36362-7_22).
- [61] Michel Goemans. *Chernoff bounds, and some applications*. 2015. URL: <https://math.mit.edu/~goemans/18310S15/chernoff-notes.pdf>.
- [62] Robert Granger et al. “Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption”. In: *EUROCRYPT 2016, Part I*. Ed. by Marc Fischlin and Jean-Sébastien Coron. Vol. 9665. LNCS. Springer, Berlin, Heidelberg, May 2016, pp. 263–293. DOI: [10.1007/978-3-662-49890-3\\_11](https://doi.org/10.1007/978-3-662-49890-3_11).
- [63] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. “Message Franking via Committing Authenticated Encryption”. In: *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10403. Lecture Notes in Computer Science. Springer, 2017, pp. 66–97. DOI: [10.1007/978-3-319-63697-9\\_3](https://doi.org/10.1007/978-3-319-63697-9_3).
- [64] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. “Message Franking via Committing Authenticated Encryption”. In: *CRYPTO 2017, Part III*. Ed. by Jonathan Katz and Hovav Shacham. Vol. 10403. LNCS. Springer, Cham, Aug. 2017, pp. 66–97. DOI: [10.1007/978-3-319-63697-9\\_3](https://doi.org/10.1007/978-3-319-63697-9_3).

- [65] Shay Gueron. *Double Nonce Derive Key AES-GCM (DNDK-GCM)*. Internet-Draft draft-gueron-cfrg-dndkgcm-02. Work in Progress. Internet Engineering Task Force, Mar. 2025. 40 pp. URL: <https://datatracker.ietf.org/doc/draft-gueron-cfrg-dndkgcm/02/>.
- [66] Shay Gueron. *Key Committing AEADs*. Cryptology ePrint Archive, Report 2020/1153. 2020. URL: <https://eprint.iacr.org/2020/1153>.
- [67] Shay Gueron, Adam Langley, and Yehuda Lindell. *AES-GCM-SIV: Specification and Analysis*. Cryptology ePrint Archive, Report 2017/168. 2017. URL: <https://eprint.iacr.org/2017/168>.
- [68] Dan Harkins. *Synthetic Initialization Vector (SIV) Authenticated Encryption Using the Advanced Encryption Standard (AES)*. Request for Comments - Informational. RFC. 2008. URL: <https://datatracker.ietf.org/doc/rfc5297/>.
- [69] Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. “Robust Authenticated-Encryption AEZ and the Problem That It Solves”. In: *EUROCRYPT 2015, Part I*. Ed. by Elisabeth Oswald and Marc Fischlin. Vol. 9056. LNCS. Springer, Berlin, Heidelberg, Apr. 2015, pp. 15–44. DOI: [10.1007/978-3-662-46800-5\\_2](https://doi.org/10.1007/978-3-662-46800-5_2).
- [70] Akiko Inoue and Kazuhiko Minematsu. “Parallelizable Authenticated Encryption with Small State Size”. In: *SAC 2019*. Ed. by Kenneth G. Paterson and Douglas Stebila. Vol. 11959. LNCS. Springer, Cham, Aug. 2019, pp. 618–644. DOI: [10.1007/978-3-030-38471-5\\_25](https://doi.org/10.1007/978-3-030-38471-5_25).
- [71] Takanori Isobe and Mostafizar Rahman. *Key Committing Security Analysis of AEGIS*. Cryptology ePrint Archive, Report 2023/1495. 2023. URL: <https://eprint.iacr.org/2023/1495>.

- [72] Takanori Isobe et al. “Areion: Highly-Efficient Permutations and Its Applications to Hash Functions for Short Input”. In: *IACR TCHES* 2023.2 (2023), pp. 115–154. DOI: [10.46586/tches.v2023.i2.115-154](https://doi.org/10.46586/tches.v2023.i2.115-154).
- [73] Jérémy Jean et al. “The Deoxys AEAD Family”. In: *Journal of Cryptology* 34.3 (July 2021), p. 31. DOI: [10.1007/s00145-021-09397-w](https://doi.org/10.1007/s00145-021-09397-w).
- [74] Philipp Jovanovic et al. “Beyond Conventional Security in Sponge-Based Authenticated Encryption Modes”. In: *Journal of Cryptology* 32.3 (July 2019), pp. 895–940. DOI: [10.1007/s00145-018-9299-7](https://doi.org/10.1007/s00145-018-9299-7).
- [75] Charanjit S. Jutla. “Encryption Modes with Almost Free Message Integrity”. In: *EUROCRYPT 2001*. Ed. by Birgit Pfitzmann. Vol. 2045. LNCS. Springer, Berlin, Heidelberg, May 2001, pp. 529–544. DOI: [10.1007/3-540-44987-6\\_32](https://doi.org/10.1007/3-540-44987-6_32).
- [76] Panos Kampanakis et al. *Practical Challenges with AES-GCM and the need for a new cipher*. Third NIST Workshop on Block Cipher Modes of Operation. 2023. URL: <https://csrc.nist.gov/csrc/media/Events/2023/third-workshop-on-block-cipher-modes-of-operation/documents/accepted-papers/Practical%20Challenges%20with%20AES-GCM.pdf>.
- [77] Joe Kilian and Phillip Rogaway. “How to Protect DES Against Exhaustive Key Search (an Analysis of DESX)”. In: *J. Cryptol.* 14.1 (2001), pp. 17–35. DOI: [10.1007/s001450010015](https://doi.org/10.1007/s001450010015).
- [78] Hugo Krawczyk. *The OPAQUE Asymmetric PAKE Protocol*. Tech. rep. Oct. 2019. URL: <https://datatracker.ietf.org/doc/draft-krawczyk-cfrg-opaque/03/>.
- [79] Ted Krovetz and Phillip Rogaway. “The Design and Evolution of OCB”. In: *Journal of Cryptology* 34.4 (Oct. 2021), p. 36. DOI: [10.1007/s00145-021-09399-8](https://doi.org/10.1007/s00145-021-09399-8).

- [80] Ted Krovetz and Phillip Rogaway. *The OCB Authenticated-Encryption Algorithm*. Request for Comments - Informational. RFC. 2014. URL: <https://datatracker.ietf.org/doc/rfc7253/>.
- [81] Ted Krovetz and Phillip Rogaway. “The Software Performance of Authenticated-Encryption Modes”. In: *FSE 2011*. Ed. by Antoine Joux. Vol. 6733. LNCS. Springer, Berlin, Heidelberg, Feb. 2011, pp. 306–327. DOI: [10.1007/978-3-642-21702-9\\_18](https://doi.org/10.1007/978-3-642-21702-9_18).
- [82] Kaoru Kurosawa. *Power of a Public Random Permutation and its Application to Authenticated-Encryption*. Cryptology ePrint Archive, Report 2002/127. 2002. URL: <https://eprint.iacr.org/2002/127>.
- [83] Julia Len, Paul Grubbs, and Thomas Ristenpart. “Partitioning Oracle Attacks”. In: *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*. Ed. by Michael Bailey and Rachel Greenstadt. USENIX Association, 2021, pp. 195–212. URL: <https://ia.cr/2020/1491>.
- [84] Julia Len, Paul Grubbs, and Thomas Ristenpart. “Partitioning Oracle Attacks”. In: *USENIX Security 2021*. Ed. by Michael Bailey and Rachel Greenstadt. USENIX Association, Aug. 2021, pp. 195–212. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/len>.
- [85] Moses Liskov, Ronald L. Rivest, and David Wagner. “Tweakable Block Ciphers”. In: *CRYPTO 2002*. Ed. by Moti Yung. Vol. 2442. LNCS. Springer, Berlin, Heidelberg, Aug. 2002, pp. 31–46. DOI: [10.1007/3-540-45708-9\\_3](https://doi.org/10.1007/3-540-45708-9_3).
- [86] Raspberry Pi Ltd. *Raspberry Pi 5 Datasheet*. raspberrypi.com. 2025. URL: <https://datasheets.raspberrypi.com/rpi5/raspberry-pi-5-product-brief.pdf>.

- [87] Atul Luykx, Bart Mennink, and Kenneth G. Paterson. “Analyzing Multi-key Security Degradation”. In: *ASIACRYPT 2017, Part II*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Vol. 10625. LNCS. Springer, Cham, Dec. 2017, pp. 575–605. doi: [10.1007/978-3-319-70697-9\\_20](https://doi.org/10.1007/978-3-319-70697-9_20).
- [88] David A. McGrew and John Viega. “The Security and Performance of the Galois/Counter Mode (GCM) of Operation”. In: *INDOCRYPT 2004*. Ed. by Anne Canteaut and Kapalee Viswanathan. Vol. 3348. LNCS. Springer, Berlin, Heidelberg, Dec. 2004, pp. 343–355. doi: [10.1007/978-3-540-30556-9\\_27](https://doi.org/10.1007/978-3-540-30556-9_27).
- [89] Sanketh Menda et al. “Context Discovery and Commitment Attacks - How to Break CCM, EAX, SIV, and More”. In: *EUROCRYPT 2023, Part IV*. Ed. by Carmit Hazay and Martijn Stam. Vol. 14007. LNCS. Springer, Cham, Apr. 2023, pp. 379–407. doi: [10.1007/978-3-031-30634-1\\_13](https://doi.org/10.1007/978-3-031-30634-1_13).
- [90] Sanketh Menda et al. *Context Discovery and Commitment Attacks: How to Break CCM, EAX, SIV, and More*. Cryptology ePrint Archive, Report 2023/526. 2023. URL: <https://eprint.iacr.org/2023/526>.
- [91] Sanketh Menda et al. *Flexible Authenticated Encryption*. Third NIST Workshop on Block Cipher Modes of Operation. 2023. URL: <https://csrc.nist.gov/csrc/media/Events/2023/third-workshop-on-block-cipher-modes-of-operation/documents/accepted-papers/Flexible%20Authenticated%20Encryption.pdf>.
- [92] Sanketh Menda et al. *The OCH Authenticated Encryption Scheme*. To appear at ACM CCS 2025. 2025.
- [93] Nicky Mouha and Atul Luykx. “Multi-key Security: The Even-Mansour Construction Revisited”. In: *CRYPTO 2015, Part I*. Ed. by Rosario Gennaro and

- Matthew J. B. Robshaw. Vol. 9215. LNCS. Springer, Berlin, Heidelberg, Aug. 2015, pp. 209–223. DOI: [10.1007/978-3-662-47989-6\\_10](https://doi.org/10.1007/978-3-662-47989-6_10).
- [94] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. “Reconsidering Generic Composition”. In: *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. Lecture Notes in Computer Science. Springer, 2014, pp. 257–274. DOI: [10.1007/978-3-642-55220-5\\_15](https://doi.org/10.1007/978-3-642-55220-5_15).
  - [95] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. “Reconsidering Generic Composition”. In: *EUROCRYPT 2014*. Ed. by Phong Q. Nguyen and Elisabeth Oswald. Vol. 8441. LNCS. Springer, Berlin, Heidelberg, May 2014, pp. 257–274. DOI: [10.1007/978-3-642-55220-5\\_15](https://doi.org/10.1007/978-3-642-55220-5_15).
  - [96] Chanathip Namprempre, Phillip Rogaway, and Tom Shrimpton. *AE5 Security Notions: Definitions Implicit in the CAESAR Call*. Cryptology ePrint Archive, Report 2013/242. 2013. URL: <https://eprint.iacr.org/2013/242>.
  - [97] Yoav Nir and Adam Langley. *ChaCha20 and Poly1305 for IETF Protocols*. RFC 7539. May 2015. DOI: [10.17487/RFC7539](https://doi.org/10.17487/RFC7539). URL: <https://www.rfc-editor.org/info/rfc7539>.
  - [98] Yoav Nir and Adam Langley. *ChaCha20 and Poly1305 for IETF Protocols*. RFC 8439. June 2018. DOI: [10.17487/RFC8439](https://doi.org/10.17487/RFC8439). URL: <https://www.rfc-editor.org/info/rfc8439>.
  - [99] Thomas Pornin. *CPU Cycle Counter*. GitHub. 2025. URL: <https://github.com/pornin/cycle-counter>.
  - [100] John PreußMattsson, Ben Smeets, and Erik Thormarker. *Proposals for Standardization of Encryption Schemes*. The Third NIST Workshop on Block Cipher



Modes of Operation. 2023. URL: <https://csrc.nist.gov/csrc/media/Events/2023/third-workshop-on-block-cipher-modes-of-operation/documents/accepted-papers/Proposals%20for%20Standardization%20of%20Encryption%20Schemes%20Final.pdf>.

- [101] Eric Rescorla. *The Transport Layer Security (TLS) Protocol Version 1.3*. RFC 8446. Aug. 2018. DOI: [10.17487/RFC8446](https://doi.org/10.17487/RFC8446). URL: <https://www.rfc-editor.org/info/rfc8446>.
- [102] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. “Careful with Composition: Limitations of Indifferentiability and Universal Composability”. In: *IACR Cryptol. ePrint Arch.* (2011), p. 339. URL: <http://ia.cr/2011/339>.
- [103] Phillip Rogaway. “Authenticated-Encryption With Associated-Data”. In: *ACM CCS 2002*. Ed. by Vijayalakshmi Atluri. ACM Press, Nov. 2002, pp. 98–107. DOI: [10.1145/586110.586125](https://doi.org/10.1145/586110.586125).
- [104] Phillip Rogaway. “Authenticated-encryption with associated-data”. In: *Proceedings of the 9th ACM Conference on Computer and Communications Security, CCS 2002, Washington, DC, USA, November 18-22, 2002*. Ed. by Vijayalakshmi Atluri. ACM, 2002, pp. 98–107. DOI: [10.1145/586110.586125](https://doi.org/10.1145/586110.586125).
- [105] Phillip Rogaway and Thomas Shrimpton. “A Provable-Security Treatment of the Key-Wrap Problem”. In: *Lecture Notes in Computer Science 4004* (2006). Ed. by Serge Vaudenay, pp. 373–390. DOI: [10.1007/11761679\\_23](https://doi.org/10.1007/11761679_23).
- [106] Phillip Rogaway and Thomas Shrimpton. “A Provable-Security Treatment of the Key-Wrap Problem”. In: *EUROCRYPT 2006*. Ed. by Serge Vaudenay. Vol. 4004. LNCS. Springer, Berlin, Heidelberg, May 2006, pp. 373–390. DOI: [10.1007/11761679\\_23](https://doi.org/10.1007/11761679_23).

- [107] Phillip Rogaway and Thomas Shrimpton. *The SIV Mode of Operation for Deterministic Authenticated-Encryption (Key Wrap) and Misuse-Resistant Nonce-Based Authenticated-Encryption*. Draft 0.32. 2007. URL: <https://web.cs.ucdavis.edu/~rogaway/papers/siv.pdf>.
- [108] Phillip Rogaway et al. “OCB: A Block-Cipher Mode of Operation for Efficient Authenticated Encryption”. In: *ACM CCS 2001*. Ed. by Michael K. Reiter and Pierangela Samarati. ACM Press, Nov. 2001, pp. 196–205. DOI: [10.1145/501983.502011](https://doi.org/10.1145/501983.502011).
- [109] Markku-Juhani O. Saarinen and Jean-Philippe Aumasson. *The BLAKE2 Cryptographic Hash and Message Authentication Code (MAC)*. RFC 7693. Nov. 2015. DOI: [10.17487/RFC7693](https://doi.org/10.17487/RFC7693). URL: <https://www.rfc-editor.org/info/rfc7693>.
- [110] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 10.6)*. 2025. URL: <https://www.sagemath.org>.
- [111] Sophie Schmieg. *Invisible Salamanders in AES-GCM-SIV*. 2020. URL: <https://keymaterial.net/2020/09/07/invisible-salamanders-in-aes-gcm-siv/>.
- [112] Sophie, indistinguishable from random noise (@SchmiegSophie). via Twitter. 2020. URL: <https://web.archive.org/web/20200909134511/https://twitter.com/SchmiegSophie/status/1303690812933382148>.
- [113] NIST Lightweight Cryptography Team. *NIST announces the selection of the Ascon family for lightweight cryptography standardization*. Feb. 2023. URL: <https://www.nist.gov/news-events/news/2023/02/lightweight-cryptography-standardization-process-nist-selects-ascon>.
- [114] Meltem Sönmez Turan et al. *Ascon-Based Lightweight Cryptography Standards for Constrained Devices*. Tech. rep. NIST Special Publication (SP) NIST SP 800-

232 ipd. Gaithersburg, MD: National Institute of Standards and Technology, 2024. DOI: [10.6028/NIST.SP.800-232.ipd](https://doi.org/10.6028/NIST.SP.800-232.ipd).

- [115] Filippo Valsorda. *The XAES-256-GCM extended-nonce AEAD*. Community Cryptography Specification Project. 2024. URL: <https://c2sp.org/XAES-256-GCM>.
- [116] David A. Wagner. “A Generalized Birthday Problem”. In: *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings*. Ed. by Moti Yung. Vol. 2442. Lecture Notes in Computer Science. Springer, 2002, pp. 288–303. URL: <https://people.eecs.berkeley.edu/~daw/papers/genbday.html>.
- [117] Michael J. Wiener. “The Full Cost of Cryptanalytic Attacks”. In: *J. Cryptol.* 17.2 (2004), pp. 105–124. DOI: [10.1007/s00145-003-0213-5](https://doi.org/10.1007/s00145-003-0213-5).
- [118] Hongjun Wu and Bart Preneel. “AEGIS: A Fast Authenticated Encryption Algorithm”. In: *SAC 2013*. Ed. by Tanja Lange, Kristin Lauter, and Petr Lisonek. Vol. 8282. LNCS. Springer, Berlin, Heidelberg, Aug. 2014, pp. 185–201. DOI: [10.1007/978-3-662-43414-7\\_10](https://doi.org/10.1007/978-3-662-43414-7_10).
- [119] Ping Zhang and Honggang Hu. *On the Provable Security of the Tweakable Even-Mansour Cipher Against Multi-Key and Related-Key Attacks*. Cryptology ePrint Archive, Report 2016/1172. 2016. URL: <https://eprint.iacr.org/2016/1172>.