IACR Transactions on Symmetric Cryptology
ISSN 2519-173X, Vol. 2024, No. 2, pp. 348–370. DOI:10.46586/tosc.v2024.i2.348-370

Context-Committing Security of Leveled Leakage-Resilient AEAD

Chandranan Dhar[1], Jordan Ethan[2], Ravindra Jejurikar[3], Mustafa Khairallah[4], Eik List[5] and Sougata Mandal[6,7]

[1] Indian Statistical Institute, Kolkata, India
chandranandhar(at)gmail.com
[2] CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
jordan.ethan(at)cispa.de
[3] Cryptography Research Centre, Technology Innovation Institute, Abu Dhabi, UAE
Ravindra.Jejurikar(at)tii.ae
[4] Dept. of Electrical and Information Technology, Lund University, Lund, Sweden
khairallah(at)ieee.org
[5] Nanyang Technological University, Singapore, Singapore
elist(at)posteo.de
[6] Institute for Advancing Intelligence, TCG CREST, Kolkata, India
[7] Ramakrishna Mission Vivekananda Educational and Research Institute, Kolkata, India
sougata.mandal(at)tcgcrest.org

Abstract. During recent years, research on authenticated encryption has been thriving through two highly active and practically motivated research directions: provable leakage resilience and key- or context-commitment security. However, the intersection of both fields had been overlooked until very recently. In ToSC 1/2024, Struck and Weishäupl studied generic compositions of encryption schemes and message authentication codes for building committing leakage-resilient schemes. They showed that, in general, Encrypt-then-MAC (EtM) and MAC-then-Encrypt (MtE) are not committing while Encrypt-and-MAC (EaM) is, under plausible and weak assumptions on the components. However, real-world schemes are rarely strict black-box constructions. Instead, while various leakage-resilient schemes follow blueprints inspired by generic compositions, they often tweak them for security or efficiency.
In this paper, we study two blueprints, the first one based on EtM for one of the strongest possible levels of leakage resilience. The second one is a single-pass framework based on leveled implementations. We show that, with a careful selection of the underlying primitives such as with identical encryption and authentication keys and a collision-resistant PRF as the MAC, these blueprints are committing. Our results do not contradict the results by Struck and Weishäupl since we pose more, but practically-motivated, requirements on the components. We demonstrate the practical relevance of our results by showing that our results on those blueprints allow us to easily derive proofs that several state-of-the-art leakage-resilient schemes are indeed committing, including TEDT and its descendants TEDT2 and Romulus-T, as well as the single-pass scheme Triplex.

Keywords: Authenticated encryption · provable security · leakage resilience · committing encryption · PRF · authentication · tweakable block cipher

1 Introduction

Authenticated Encryption with Associated Data (AEAD) has become a fundamental component in modern security applications, providing both confidentiality and authenticity.

Received: 2024-03-01 Accepted: 2024-05-01 Published: 2024-06-18

The development of efficient AEAD schemes has resulted in widespread constructions like AES-GCM [SCM08], Ascon [DEMS21], Deoxys [JNPS21], and AES-GCM-SIV [GLL19] that address diverse security goals including but not limited to Nonce-based AEAD (nAE) [Rog04], Misuse-Resistant AEAD (MRAE) [RS06], or Deterministic AEAD (DAE) [RS06].

However, as AEAD schemes and their analyses mature, attackers continuously seek new ways to exploit their security. Additionally, new applications introduce fresh security challenges. Consequently, two areas of research have gained prominence:

1. *Leakage-resilient AEAD*: This area focuses on security notions and schemes where the adversary can observe different forms of auxiliary leakage that may depend on sensitive or secret information. The objective is to construct schemes that maintain confidentiality and authenticity even in the presence of certain leakage.

2. *Context-committing AEAD*: This area addresses scenarios where the adversary has access to, and can manipulate secret keys. For example, it deals with situations where the ciphertext allows for correct decryption under multiple contexts, where a context consists of the key K, the nonce N, and associated data A.

Leakage-resilient AEAD. This area of research has blossomed for almost two decades. In this work, we focus on recent developments, and on the schemes discussed by Bellizia et al. in [BBC+20] in particular. Therein, the authors categorized modern leakage-resilient AEAD schemes into four grades, with a focus on so-called leveled implementations. In such schemes, a few functions are assumed to be either leak-free or heavily protected, while the rest of the construction can leak a lot of information. In this work, we will focus on two relevant types of schemes:

- Grade-3 schemes: These schemes usually follow the pattern of Encrypt-then-MAC (EtM), using a hash function and two calls to a heavily protected Tweakable Block Cipher (TBC) implementation. They target Ciphertext Integrity with Misuse and Decryption Leakage (CIML2) security and indistinguishability against Chosen-Ciphertext Adversaries with misuse resilience and decryption Leakage (CCAmL2) security. Examples of schemes in this category are TEDT [BGP+20] and ISAP [DEM+17].

- Grade-2 schemes: These schemes usually employ a single-pass AEAD scheme, a hash function, and two heavily protected TBC calls. They target CIML2 security and indistinguishability against Chosen-Ciphertext Adversaries with misuse resilience and encryption Leakage (CCAmL1) security. An example is Triplex [SPS+22].

Because they cover many practical schemes, those two grades serve as the basis for the blueprints we study in this paper.

Context-Committing AEAD. In recent years, a series of attacks such as the Facebook message-franking attack [DGRW18], and the partitioning-oracle attack [LGR21] have shown vulnerabilities in the usage of conventionally secure AEAD schemes. Those works shared a common root cause: the existence of ciphertexts that can be decrypted correctly under multiple keys, which was out of the scope of conventional AEAD security but is necessary for security in the respective uses in practice.

To address this gap, Bellare and Hoang introduced commitment security in [BH22], which requires each ciphertext to commit to the key (CMT-1) or to the entire context (CMT-4) that produced it. Among the notions Bellare and Hoang proposed, CMT-4 represents the strongest and therefore most desirable form for designers. It is formalized through the following game. Given an AEAD scheme Π with an encryption function $\mathcal{E}$, an adversary has the task of providing two contexts, i.e. tuples (K, N, A, M) and (K', N', A', M'), consisting of a key, nonce, associated data, and message each. The adversary wins the game if the

contexts differ, i.e. $(K, N, A, M) \neq (K', N', A', M')$ but they both encrypt to the same ciphertexts: $\Pi.\mathcal{E}(K, N, A, M) = \Pi.\mathcal{E}(K', N', A', M')$.

Connecting Both Areas. At first glance, the overlap between leakage-resilient and context-committing AEAD is unclear as there has been little exploration of their potential synergies. With their recent work, Struck and Weishäupl [SW24] began to shed light on relations to investigating the generic compositions of Encryption and Message Authentication Code (MAC) schemes to develop schemes that are both leakage-resilient and committing. Their study revealed that EtM and MAC-then-Encrypt (MtE) are not committing in general. They also demonstrated that Encrypt-and-MAC (EaM) can achieve committing properties under weak assumptions on the underlying schemes. Additionally, they presented a transformation that converts an AEAD scheme into a leakage-resilient and context-committing scheme. In a separate work, Krämer, Struck and Weishäupl [KSW23] have shown that the Grade-3 scheme ISAP is committing.

Contribution. While black-box compositions such as EtM, EaM, or MtE are valuable for studying generic constructions and inspiring instantiations, real-world schemes often deviate from them. In particular, many leakage-resilient schemes are based on blueprints that take inspiration from generic compositions but incorporate small changes tailored to specific security goals or higher efficiency. In this paper, we explore two such blueprints. The first blueprint is based on EtM and aims to achieve the highest level of leakage resilience. The second blueprint targets leveled single-pass implementations. We demonstrate that, with a careful selection of underlying primitives, both blueprints can be committing.

While our findings on the first blueprint may seem to contradict the negative result in [SW24] on EtM, two differences in the underlying assumptions help clarify. Firstly, in [SW24], the authors considered a black-box composition where the encryption function and the MAC use independent keys, whereas we require a strict dependency between both. Secondly, our result requires a certain type of MAC in the scheme, namely a collision-resistant Pseudo-Random Function (PRF).

Those additional requirements are not impractical. Our result on EtM generalizes [FOR17, Theorem 3] and [GLR17, Theorem 3], where the authors demonstrated similar restrictions on keys and MACs. The results on both blueprints allow us to easily derive that several leakage-resilient schemes are committing, including TEDT and its descendants TEDT2 and Romulus-T, as well as the single-pass scheme Triplex.

For schemes that follow our blueprints, showing their commitment security reduces to showing the collision resistance of their building blocks. For this purpose, we study the collision resistance of several leakage-resilient MACs used in EtM-based schemes, including Hash-then-BC (HBC), Hash-then-TBC (HTBC), and LRMAC1. For single-pass schemes, we examine instead the collision resistance of more components, including their functions for Key Derivation (KDF), Encryption (Enc), and Tag Generation (TGF).

Our analysis poses a few cryptographic assumptions on the used components. For keyed primitives, we operate in the ideal-cipher model, which is unavoidable in the chosen-key setting of committing security. For hash functions and compression functions, we require either collision and everywhere-preimage resistance, or collision resistance only. For LRMAC1, which requires only collision resistance, this matches the assumption on the hash function in the original MAC proof. For HBC, we require collision and everywhere-preimage resistance, which is still more concrete than the random-oracle model used in the original proof. For HTBC, we require collision and everywhere-preimage resistance, while the MAC proof requires collision and range-oriented preimage resistance. While our assumptions are slightly stronger, they are close to practice, as everywhere-preimage resistance can be seen as a worst-case analysis of range-oriented preimage resistance and is therefore, expected to lead to a similar bound for any secure standard hash function.

2 Preliminaries

In this section, we define some necessary security notions of functions in general and hash functions in particular, e.g., collision resistance, preimage resistance, and right collision resistance. Thereupon, we recall the definitions for the primitives we need for authenticated encryption.

General Notation. For a set $\mathcal{X}$, we write $X \xleftarrow{\$} \mathcal{X}$ to denote that a value X is sampled uniformly at random from $\mathcal{X}$ and independent from other values. An adversary $\mathcal{A}$ is a computationally bounded algorithm that shall win a security game against a challenger. We call $\mathcal{A}$ t-bounded if it runs in time at most t. We indicate that $\mathcal{A}$ outputs X by $\mathcal{A} \Rightarrow X$. W.l.o.g., we assume that adversaries never ask pointless queries, i.e. queries to which they can compute the answers themselves. In the following, we will introduce standard primitives and security notions. Throughout this section, we will denote non-empty sets and spaces by calligraphic uppercase variables and use $\mathcal{K}$ (or $\mathcal{K}_h$), $\mathcal{M}$, $\mathcal{C}$, $\mathcal{X}$, as spaces for keys, plaintexts, ciphertexts, and hash values, respectively.

2.1 Security Notions for Keyed Hash Functions

Collision Resistance (CR). Let $H : \mathcal{K}_h \times \mathcal{M} \to \mathcal{X}$ be a hash function and $K_h \xleftarrow{\$} \mathcal{K}_h$. H is called $(\epsilon_{\mathsf{cr}}, t)$-collision-resistant ($\mathsf{CR}$) if for every t-bounded adversary $\mathcal{A}$, the probability that $\mathcal{A}(K_h)$ outputs a pair of distinct values $(M_0, M_1) \in \mathcal{M}^2$, such that $M_0 \neq M_1$ and $H_{K_h}(M_0) = H_{K_h}(M_1)$ is bounded by ϵ_{cr}:

$$\Pr\left[K_h \xleftarrow{\$} \mathcal{K}_h, \mathcal{A}(K_h) \Rightarrow (M_0, M_1) \in \mathcal{M}^2 \text{ s. t. } M_0 \neq M_1, H_{K_h}(M_0) = H_{K_h}(M_1)\right] \leq \epsilon_{\mathsf{cr}}\,.$$

Right Collision Resistance (RCR). Let $\mathcal{X}_l \times \mathcal{X}_r$ be nonempty sets of spaces $\mathcal{X}_l$ and $\mathcal{X}_r$. Let $H : \mathcal{K}_h \times \mathcal{M} \to \mathcal{X}_l \times \mathcal{X}_r$ be a hash function and $K_h \xleftarrow{\$} \mathcal{K}_h$. H is called $(\epsilon_{\mathsf{rcr}}, t)$-right collision-resistant (RCR) if for every t-bounded adversary $\mathcal{A}$, the probability that $\mathcal{A}(K_h)$ outputs a pair of distinct values $(M_0, M_1) \in \mathcal{M}^2$, such that $(X_0, Y_0) = H_{K_h}(M_0)$, $(X_1, Y_1) = H_{K_h}(M_1)$, $Y_0 = Y_1$, and $M_0 \neq M_1$, is bounded by ϵ_{rcr}:

$$\Pr\Big[K_h \xleftarrow{\$} \mathcal{K}_h, \mathcal{A}(K_h) \Rightarrow (M_0, M_1) \in \mathcal{M}^2 \text{ s. t. } M_0 \neq M_1, (X_0, Y_0) = H_{K_h}(M_0), (X_1, Y_1) = H_{K_h}(M_1), Y_0 = Y_1\Big] \leq \epsilon_{\mathsf{rcr}}\,.$$

We can define Left Collision Resistance (LCR) analogously. In the following, we extend the definition to a hash function with multiple inputs, where the collision resistance property holds for a subset of the inputs.

Collision Resistance on a Subset of the Inputs (Partial CR). Let $\mathcal{M}_1$, $\mathcal{M}_2$, $\ldots \mathcal{M}_n$ denote nonempty sets or spaces and define $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2 \times \cdots \times \mathcal{M}_n$. Let $H : \mathcal{K}_h \times \mathcal{M} \to \mathcal{X}$ be a hash-function and $K_h \xleftarrow{\$} \mathcal{K}_h$. For a positive integer $i \in \{1, \ldots, n\}$, H is called $(\epsilon_{\mathsf{cr}}, t)$ $(\mathcal{M}_1 \ldots \mathcal{M}_i)$-collision-resistant, if the probability of finding a hash collision on two distinct messages $M = (M_1, \ldots, M_n), M' = (M'_1, \ldots, M'_n) \in \mathcal{M}^2$ s.t. $(M_1 \ldots, M_i) \neq (M'_1 \ldots, M'_i)$ is bounded by ϵ_{cr}:

$$\Pr\Big[K_h \xleftarrow{\$} \mathcal{K}_h, \mathcal{A}(K_h) \Rightarrow (M_1 \ldots, M_n), (M'_1 \ldots, M'_n) \in \mathcal{M}^2, \text{ s. t. } (M_1 \ldots, M_i) \neq (M'_1 \ldots, M'_i), H_{K_h}(M_1, \ldots, M_n) = H_{K_h}(M'_1 \ldots, M'_n)\Big] \leq \epsilon_{\mathsf{cr}}\,.$$

Note that if $i < n$, then $(M_{i+1}, \ldots, M_n)$ and $(M'_{i+1}, \ldots, M'_n)$ may or may not be equal. If $i = n$, the definition is equivalent to standard collision resistance. In general, a function can be collision-resistant on any subset of the inputs, and the inputs are explicitly given in the collision-resistance property. We will call a hash function "partial right collision-resistant" (partial RCR) if it achieves right collision resistance on a subset of inputs.

Everywhere Preimage Resistance. Let $H : \mathcal{K}_h \times \mathcal{M} \to \mathcal{X}$ be a hash function and $K_h \xleftarrow{\$} \mathcal{K}_h$. H is called $(\epsilon_{\mathsf{epre}}, t)$-everywhere-preimage-resistant (ePre) if for every t-bounded adversary $\mathcal{A}$,

$$\max_{X \in \mathcal{X}} \left\{ \Pr\left[K_h \xleftarrow{\$} \mathcal{K}_h, \mathcal{A}(K_h) \Rightarrow M, \text{ s. t. } H_{K_h}(M) = X \right] \right\} \leq \epsilon_{\mathsf{epre}} .$$

2.2 Primitives

Tweakable Block Cipher (TBC). A TBC is a mapping $\widetilde{E} : \mathcal{K} \times \mathcal{T}_w \times \{0,1\}^n \to \{0,1\}^n$ such that for any choice of $K \in \mathcal{K}$ and $T_w \in \mathcal{T}_w$, $\widetilde{E}(K, T_w, \cdot)$ is a permutation over $\{0,1\}^n$. If $\mathcal{T}_w = \emptyset$, then $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ is a Block Cipher (BC). We will sometimes write $\widetilde{E}^T_K(X)$ and $E_K(X)$ for $\widetilde{E}(K, T, X)$ and $E(K, X)$, respectively. In this paper, we will analyze constructions in the ideal-cipher model, where $\widetilde{E}$ or E will be selected randomly from the set of all possible cipher families with the same domain and range.

Pseudo-random Function (PRF) and Pseudo-Random Number Generator (PRNG). A PRF is a deterministic mapping $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$. Let $\mathsf{Func}(\mathcal{X}, \mathcal{Y})$ be the set of all functions with domain $\mathcal{X}$ and range $\mathcal{Y}$. In the PRF game, a challenger samples $K \xleftarrow{\$} \mathcal{K}$ and $\rho \xleftarrow{\$} \mathsf{Func}(\mathcal{X}, \mathcal{Y})$ and provides an adversary $\mathcal{A}$ with access to either F_K or ρ. The PRF advantage of $\mathcal{A}$ on F_K is defined as

$$\mathbf{Adv}^{\mathsf{PRF}}_{F_K}(\mathcal{A}) = \left| \Pr[\mathcal{A}^{F_K} \Rightarrow 1] - \Pr[\mathcal{A}^{\rho} \Rightarrow 1] \right| .$$

We call F an $(\epsilon_{\mathsf{PRF}}, t)$-secure PRF if for all t-bounded adversaries $\mathcal{A}$, $\mathbf{Adv}^{\mathsf{PRF}}_{F_K}(\mathcal{A}) \leq \epsilon_{\mathsf{PRF}}$.

A PRNG is a deterministic mapping $G : \{0,1\}^k \times \mathbb{N} \to \mathcal{Y}$ with $\mathcal{Y} \subseteq \{0,1\}^*$ that takes an input $K \in \{0,1\}^k$ and a positive integer ℓ as inputs and outputs $Y \in \{0,1\}^{k+\ell}$. In the PRNG game, a challenger samples $K \xleftarrow{\$} \mathcal{K}$ and $\rho \xleftarrow{\$} \mathsf{Func}(\mathcal{K} \times \mathbb{N}, \mathcal{Y})$ and, on input of a length ℓ, outputs either $G(K, \ell)$ or $\rho(K, \ell)$. Then, the PRNG advantage of an adversary $\mathcal{A}$ against G is defined as

$$\mathbf{Adv}^{\mathsf{PRNG}}_{G}(\mathcal{A}) = \left| \Pr[\mathcal{A}^{G} \Rightarrow 1] - \Pr[\mathcal{A}^{\rho} \Rightarrow 1] \right| .$$

We call G an $(\epsilon_{\mathsf{PRNG}}, t)$-secure PRNG if for all t-bounded adversaries $\mathcal{A}$, it holds that $\mathbf{Adv}^{\mathsf{PRNG}}_{G}(\mathcal{A}) \leq \epsilon_{\mathsf{PRNG}}$. Later, we will use a generalization with multiple output-length parameters ℓ_1, ℓ_2 where $\mathcal{Y} = \{0,1\}^{\ell_1} \times \{0,1\}^{\ell_2}$ and so on.

Hash-function-based Message Authentication Codes (MACs). Let $\mathsf{MAC} : \mathcal{K}_h \times \mathcal{K} \times \mathcal{M} \to \mathcal{T}$ be a keyed function that transforms an input $M \in \mathcal{M}$ to a tag $T \in \mathcal{T}$. Let $\mathsf{Ver}_{\mathsf{MAC}} : \mathcal{K}_h \times \mathcal{K} \times \mathcal{M} \times \mathcal{T} \to \{\bot, \top\}$ be the verification function that takes a message $M \in \mathcal{M}$ and a would-be tag $T \in \mathcal{T}$ and returns $\top$ if $\mathsf{MAC}_{s,K}(M) = T$ and $\bot$ otherwise. In other contexts, security notions such as unforgeability and/or pseudo-randomness are needed, but in the context of this paper, we are interested only in collision resistance, which we define for hash-function-based MACs as follows.

Collision Resistance for Hash-function-based MACs. Let $K_h \xleftarrow{\$} \mathcal{K}_h$. A MAC is called $(\epsilon_{\mathsf{maccr}}, t)$-collision-resistant (MAC-CR) if for every t-bounded adversary $\mathcal{A}$, the probability that $\mathcal{A}(K_h)$ outputs a pair of distinct inputs $(K_0, M_0), (K_1, M_1) \in (\mathcal{K} \times \mathcal{M})^2$, such that $\mathsf{MAC}_{K_h,K_0}(M_0) = \mathsf{MAC}_{K_h,K_1}(M_1)$ and $(K_0, M_0) \neq (K_1, M_1)$, is bounded by $\epsilon_{\mathsf{maccr}}$:

$$\Pr\Big[K_h \xleftarrow{\$} \mathcal{K}_h, \mathcal{A}(K_h) \Rightarrow ((K_0, M_0), (K_1, M_1)) \in (\mathcal{K} \times \mathcal{M})^2 \\ \text{s. t. } (K_0, M_0) \neq (K_1, M_1), \mathsf{MAC}_{K_h,K_0}(M_0) = \mathsf{MAC}_{K_h,K_1}(M_1)\Big] \leq \epsilon_{\mathsf{maccr}}\,.$$

In all collision games in the remainder, we will drop the hash key K_h that is released to the adversary and assume that it is given to the adversary at the beginning of the game.

Authenticated Encryption. A nonce-based Authenticated Encryption scheme supporting Associated Data (nAEAD) is a pair of functions $\Pi = (\mathcal{E}, \mathcal{D})$ with associated sets $\mathcal{K}, \mathcal{N}, \mathcal{A}, \mathcal{M}$, $\mathcal{C}$, denoting the key space, nonce space, associated data space, message space, and ciphertext space, respectively. The elements of $\mathcal{C}$ comprise of a pair (C, T), with $T \in \{0,1\}^\sigma$. The encryption algorithm and decryption algorithms $\mathcal{E}$ and $\mathcal{D}$ are deterministic functions input and $\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C} \to \mathcal{M} \cup \{\perp\}$, where the special symbol $\perp$ indicates that (C, T) was deemed invalid. We sometimes write $\mathcal{E}_K^{N,A}(M)$ and $\mathcal{D}_K^{N,A}(C, T)$ to denote $\mathcal{E}(K, N, A, M)$ and $\mathcal{D}(K, N, A, (C, T))$. For all correct schemes, encryption is injective from $\mathcal{M}$ to $\mathcal{C}$ under fixed K, N, A. The scheme is correct if $\mathcal{D}_K^{N,A}(\mathcal{E}_K^{N,A}(M)) = M$ for all K, N, A, M and tidy if $M = \mathcal{D}_K^{N,A}(C, T)$ and $M \neq \perp$, then $\mathcal{E}_K^{N,A}(M) = (C, T)$ for all $K, N, A, (C, T)$. All AEAD schemes considered in this work are assumed to be correct and tidy.

CMT-4 Security. The two prevalent notions of committing security in the literature are

- CMT-1 security: A commitment to only the key K.
- CMT-4 security: A commitment to the complete context (K, N, A, M).

Since we consider only CMT-4 security, we define it more formally here. Note that Bellare and Hoang [BH22] demonstrated that incorporating the message M into the context is unnecessary, as committing to (K, N, A) is equivalent to committing to (K, N, A, M).

In the CMT-4 game against an AEAD scheme Π, an adversary $\mathcal{A}$ outputs (K_1, N_1, A_1, M_1) and (K_2, N_2, A_2, M_2); $\mathcal{A}$ wins if and only if $(K_1, N_1, A_1, M_1) \neq (K_2, N_2, A_2, M_2)$ and $\Pi.\mathsf{Enc}(K_1, N_1, A_1, M_1) = \Pi.\mathsf{Enc}(K_2, N_2, A_2, M_2)$. We write $\mathbf{Adv}_{\Pi}^{\mathtt{cmt4}}(\mathcal{A})$ to denote the probability that $\mathcal{A}$ wins the CMT-4 game where $\mathcal{A}$ has access to the ideal primitives and hash keys used by Π.

3 Blueprints for Leakage Resilience and their Motivation as Committing Schemes

In this section, we describe a paradigm for designing leakage-resilient schemes based on so-called leveled implementations. In this paradigm, different parts of the scheme have different assumptions on how they are implemented and the associated leakage functions. Since our focus is on the relation to context commitment, we study only integrity and Ciphertext Integrity with Misuse and Leakage with decryption leakage (CIML2). We will also consider a widespread leakage model wherein the adversary can receive unlimited leakage, most of the scheme is unprotected, and only the protected parts are assumed to be leak-free.

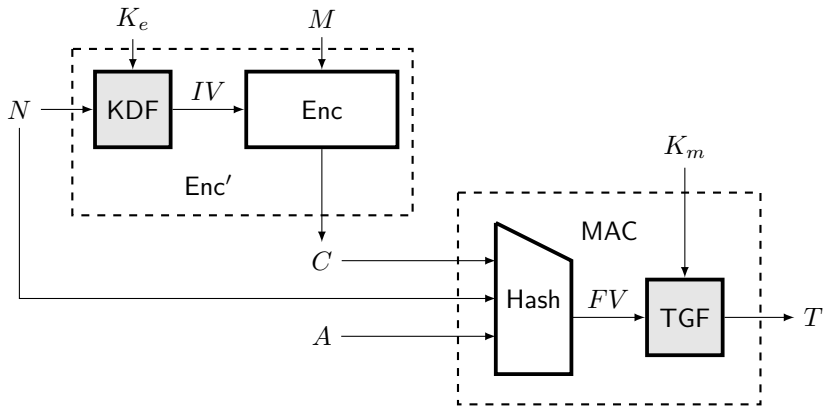

Figure 1: The EtM blueprint for leakage-resilient context-committing nAEAD. The gray components are assumed to be strongly protected.

3.1 The EtM-based Blueprint

The first blueprint we look at is based on EtM but tailored to leakage resilience. It is close to the FGHF′ construction by Degabriele et al. [DJS19], which itself is an instance of N2 [NRS14]. In the following, let $\mathcal{K}_e$, $\mathcal{K}_m$, $\mathcal{N}$, $\mathcal{A}$, $\mathcal{M}$, $\mathcal{C}$, $\mathcal{IV}$, $\mathcal{FV}$, and $\mathcal{T}$ be nonempty sets or spaces for encryption keys, MAC keys, nonces, associated data, plaintexts, ciphertexts, initial values, forward values, and tags, respectively. A leveled leakage-resilient EtM scheme requires two *leak-free* fixed-input-length primitives:

1. A key-derivation function $\mathsf{KDF} : \mathcal{K}_e \times \mathcal{N} \to \mathcal{IV}$, which takes the nonce and the encryption key and generates an initial value for the encryption phase.
2. A tag-generation function $\mathsf{TGF} : \mathcal{K}_m \times \mathcal{FV} \to \mathcal{T}$, which takes the MAC key and a fixed-length hash of the ciphertext, nonce, and associated data, and generates the verification tag using a PRF.

The scheme also uses an encryption scheme $\mathsf{Enc} : \mathcal{K} \times \mathcal{IV} \times \mathcal{M} \to \mathcal{C}$ and a collision-resistant hash function $\mathsf{Hash} : \mathcal{N} \times \mathcal{A} \times \mathcal{C} \to \mathcal{FV}$. However, these two primitives are assumed to have unlimited leakage when considering CIML2 security. The high-level blueprint is depicted in Figure 1 and the encryption of $\mathsf{EtM}[\mathsf{KDF}, \mathsf{Enc}, \mathsf{Hash}, \mathsf{TGF}](K, N, A, M)$ under $K = (K_e, K_m)$ is defined as

$$IV \leftarrow \mathsf{KDF}(K_e, N)\,, \quad C \leftarrow \mathsf{Enc}(IV, A, M)\,,$$
$$FV \leftarrow \mathsf{Hash}(N, A, C)\,, \quad T \leftarrow \mathsf{TGF}(K_m, FV)\,.$$

We observe three important properties. First, since the MAC, consisting of Hash and TGF, follows the Hash-then-PRF paradigm, it already binds the triplet (N, A, C) to any given key K_m. However, K_e is not part of the binding. Second, if the two parts of the key K_e and K_m are independent, the adversary can fix (K_m, C, T, N, A) and find two pairs (K_{e_1}, M_1) and (K_{e_2}, M_2) to break the commitment. This implies that even if the keys are dependent, we must ensure that the EtM scheme commits to (K_e, K_m) and the MAC is collision-resistant. In Section 4, we shall show that under these restrictions, EtM is indeed context-committing. Moreover, we shall show that three of the prominent leakage-resilient Hash-then-PRF MACs are indeed collision-resistant PRFs. Two of these MACs will require a stronger assumption on the hash function, where the hash function resists not only collision but also preimage attacks, while all three require a stronger assumption on the TGF, as the analysis has to be conducted in the ideal-cipher model. However, we will show that the KDF does not affect the committing security.

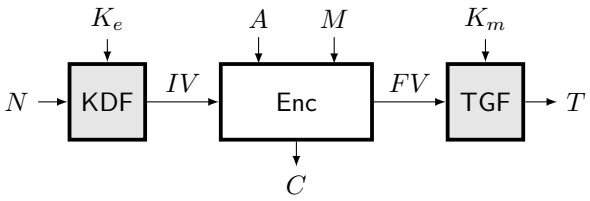

Figure 2: The KET blueprint for single-pass leveled leakage-resilient context-committing nAEAD. The gray components are assumed to be strongly protected.

3.2 The Single-pass Blueprint KET

The second blueprint we will consider is used for single-pass schemes. Similar to EtM, the scheme includes leak-free KDF and TGF functions and an encryption function Enc with unlimited leakage. The encryption function is responsible for generating both the ciphertext C as well as the hash value FV. This blueprint $\Pi[\mathsf{KDF}, \mathsf{Enc}, \mathsf{TGF}]$ denoted as KET is depicted in Figure 2 and its encryption of (N, A, M) under a key tuple (K_e, K_m) is defined as follows:

$$IV \leftarrow \mathsf{KDF}(K_e, N)\,, \quad (C, FV) \leftarrow \mathsf{Enc}(IV, A, M)\,, \quad T \leftarrow \mathsf{TGF}(K_m, FV)\,.$$

We can observe some requirements on the components. CIML2 security requires that the TGF is collision-resistant for a given key K. Similarly as EtM, this means that the scheme commits to (N, A, M) for K. Similarly as for EtM, further issues arise when considering the keys. If the TGF is not collision-resistant, then we can find $K_{m_1} \neq K_{m_2}$ such that $T_1 = T_2$ for the same (N, A, M). The commitment will break if K_m is independent of K_e. However, if K_m depends on K_e, the success of the attack depends on the properties of the KDF and the interaction between the KDF and the Enc function. Alternatively, it may be possible to relax the requirements on the KDF if the TGF is collision-resistant.

4 CMT-4 Security of EtM-based AEAD Schemes

The first blueprint we will study concerns EtM-type constructions with a MAC that follows the Hash-then-PRF design. This paradigm is used widely in leveled leakage-resilient schemes, including TEDT [BGP+20], TEDT2 [Lis21], Romulus-T [IKMP20], and ISAP [DEM+17]. Definition 1 describes the generalized EtM construction, including generating the encryption and MAC keys from a master key K. Therein, we abstract away some internal details. We aggregate the functions KDF and Enc into a function Enc′ and similarly, wrap Hash and TGF in a function MAC.

Definition 1. Let $\Pi[\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Mac}]$ be a nonce-based AEAD scheme such that

$$(K_e, K_m) \leftarrow \mathsf{KeyGen}(K), \quad C \leftarrow \mathsf{Enc}(K_e, N, M), \quad \text{and} \quad T \leftarrow \mathsf{Mac}(K_m, N, A, C)\,.$$

Then, we call Π an EtM scheme.

In [SW24], the EtM (or N2 [NRS14]) scheme is shown to be not context-committing in general. We show that the EtM scheme is nevertheless CMT-4-secure when KeyGen is right collision-resistant (which precludes independent keys K_e and K_m) and MAC is collision-resistant. Similar results were already shown for complete robustness (CROB) and binding security. Thus, Theorem 1 is adapted from [FOR17, Theorem 3] and [GLR17, Theorem 3].

Table 1: Examples for functions $\mathsf{KeyGen} : \{0,1\}^k \times \mathbb{N}^2 \to \{0,1\}^{k_e} \times \{0,1\}^{k_m}$ that compute $(K_e, K_m) \leftarrow \mathsf{Keygen}(K, k_e, k_m)$ and their (ϵ_k, t)-right collision resistance. Let $G : \mathcal{K} \times \mathbb{N} \to \{0,1\}^{k_e} \times \{0,1\}^{k_m}$ be a (ϵ_G, t)-secure PRNG.

KeyGen	ϵ_k
(K_e, K_e) for $k_m = k_e$	0
$(K_e, K_e \oplus \theta)$ for $\theta \in \{0,1\}^{k_m} \setminus \{0\}$ and $k_m = k_e$	0
$G(K, k_m, k_e)$	$\epsilon_G + 2^{-k_m}$
$(K[k_e + k_m - 1..k_m], K[k_m - 1..0])$ for $k = k_e + k_m$	1
$(K, 1 \| K[k-2..0])$ for $k = k_e = k_m$	1

Theorem 1. Let $\Pi[\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{MAC}]$ be an EtM scheme such that KeyGen is (ϵ_k, t_1)-right collision-resistant and MAC is $(\epsilon_{\mathsf{maccr}}, t_2)$-collision-resistant for some $t_1 = O(t)$ and $t_2 = O(t)$. Then, for any t-bounded CMT-4 adversary $\mathcal{A}$ against Π, it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \epsilon_{\mathsf{maccr}} + \epsilon_k \,.$$

Proof. Suppose an adversary $\mathcal{A}$ outputs challenge values (K_1, N_1, A_1, M_1) and (K_2, N_2, A_2, M_2) with corresponding ciphertexts (C_1, T_1) and (C_2, T_2). We bound the probability that $(C_1, T_1) = (C_2, T_2) = (C, T)$. We define a sequence of hybrid games G_0 through G_2 as follows, where we introduce Boolean variables E_i, for $i \in \{0,1,2\}$ such that E_i is true if and only if the adversary wins in Game G_i.

Game G_0. This is the original `cmt4` game in the real world.

Game G_1. Game G_1 is almost identical to Game G_0 but adds the aspect that G_1 terminates if $K_{m_1} = K_{m_2} \wedge K_1 \neq K_2$. The probability of this event is at most

$$|\Pr[\mathsf{E}_0] - \Pr[\mathsf{E}_1]| \leq \epsilon_k \,.$$

Game G_2. Game G_2 is almost identical to G_1 except that G_2 also terminates if $(N_1, A_1, K_1) = (N_2, A_2, K_2)$ and $M_1 \neq M_2$. Since this is impossible from the assumption that the encryption scheme is correct and tidy, it follows that

$$|\Pr[\mathsf{E}_1] - \Pr[\mathsf{E}_2]| = 0 \,.$$

Finally, the adversary wins Game G_2 if it is successful with $(K_{m_1}, N_1, A_1) \neq (K_{m_1}, N_2, A_2)$, $K_{m_1} = K_{m_2} \longleftrightarrow K_1 = K_2$[1]. and $\mathsf{Mac}(K_{m_1}, N_1, A_1, C) = \mathsf{Mac}(K_{m_2}, N_2, A_2, C)$. This can happen only if there is a collision against the MAC given none of the previous conditions occurs. As a result, we can upper bound the probability by

$$\Pr[\mathsf{E}_2] \leq \epsilon_{\mathsf{maccr}} \,.$$

To sum up,

$$\mathrm{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) = \Pr[\mathsf{E}_0] \leq \left(\sum_{i=1}^{2} |\Pr[\mathsf{E}_{i-1}] - \Pr[\mathsf{E}_i]| \right) + \Pr[\mathsf{E}_2] \leq \epsilon_{\mathsf{maccr}} + \epsilon_k \,,$$

which yields our claim in Theorem 1. □

[1]Two-way implication: $K_1 = K_2 \implies K_{m_1} = K_{m_2}$ and $K_{m_1} = K_{m_2} \implies K_1 = K_2$

Secure and Insecure Examples of KeyGen. Theorem 1 shows that CMT-4 security of EtM schemes relies on the right collision resistance of the KeyGen function as well as the collision resistance of the MAC. KeyGen functions with low ϵ_k can be found easily. For concreteness, we listed a few intuitive secure examples in Table 1, alongside two negative examples with $\epsilon_k = 1$. The next section can therefore concentrate on the collision resistance of MACs in and for EtM-based schemes.

5 Collision Resistance of Leveled Leakage-resilient MACs

In this section, we show the collision resistance of HBC [BGP+19], HTBC [BGP+19] (the MAC used in TEDT and Romulus-T), and LRMAC1 [BGPS21]. While we are unaware of concrete AEAD schemes that employ LRMAC1, establishing its suitability for CMT-4-secure AEAD is relevant as it offers useful leakage resilience.

5.1 Collision Resistance of Hash-then-BC (HBC)

The MAC Hash-then-BC [BGP+19] is defined as follows. Given a hash function $H : \mathcal{K}_h \times \mathcal{M} \to \{0,1\}^n$ and a block cipher $E : \mathcal{K}_m \times \{0,1\}^n \to \{0,1\}^n$, the authentication of a message $M \in \mathcal{M}$ to a tag $T \in \{0,1\}^n$ with $\mathsf{HBC}[H,E]_{K_m,K_h}$ under keys $K_m, K_h \in \mathcal{K}_m \times \mathcal{K}_h$ is computed as

$$T \leftarrow E(K_m, H_{K_h}(M)) .$$

First, we define the collision-resistance game as follows. The adversary $\mathcal{A}$ gets the hash-function key K_h at the beginning of the game. From here on, we drop all further occurrences of hash-function keys and will proceed similarly in all following games. $\mathcal{A}$ asks q_e chosen-key queries to the ideal-cipher oracle E and obtains the corresponding outputs. If a query (K_i, X_i) is in forward direction, $\mathcal{A}$ obtains $Y_i \leftarrow E(K_i, X_i)$; if a query (K_i, Y_i) is in backward direction, it obtains $X_i \leftarrow E^{-1}(K_i, Y_i)$. At the end of its interactions, $\mathcal{A}$ outputs two pairs (K_{m_1}, M_1) and (K_{m_2}, M_2) and wins if and only if $\mathsf{HBC}[H,E](K_{m_1}, M_1) = \mathsf{HBC}[H,E](K_{m_2}, M_2)$.

Theorem 2. Let $E : \mathcal{K}_m \times \{0,1\}^n \to \{0,1\}^n$ be an ideal cipher and $H : \mathcal{M} \to \{0,1\}^n$ be a $(\epsilon_{\mathsf{cr}}, t_1)$-collision-resistant and $(\epsilon_{\mathsf{epre}}, t_2)$-everywhere-preimage-resistant hash function. Then for any adversary $\mathcal{A}$ that runs in time t and makes $q_e \le 2^{n-1}$ queries to the ideal cipher, such that $t_1 = O(t + q_e)$ and $t_2 = O(t + q_e)$, $\mathsf{HBC}[H,E]$ is (ϵ, t)-collision-resistant for

$$\epsilon \le q_e \epsilon_{\mathsf{epre}} + \epsilon_{\mathsf{cr}} + \frac{2q_e^2 + 1}{2^n} .$$

Proof. Suppose $\mathcal{A}$ outputs (K_{m_1}, M_1) and (K_{m_2}, M_2) such that $\mathsf{HBC}[H,E](K_{m_1}, M_1) = \mathsf{HBC}[H,E](K_{m_2}, M_2)$. We define a sequence of hybrid games as follows: Let E_i be the event that the adversary wins in Game G_i for $i \in \{0, \ldots, 3\}$.

Game G_0. The real-world game.

Game G_1. Game G_1 is almost identical to Game G_0 but terminates if one of the following events happens during the ideal-cipher queries of $\mathcal{A}$.

- Two forward queries with different keys produce the same output. The probability of this event is upper bounded by $\binom{q_e}{2}/(2^n - q_e) \le q_e^2/2^n$.
- A backward query with input T is followed by a forward query with output T with a different key. The probability of this event is at most $q_e^2/2^n$.

It follows that

$$|\Pr[\mathsf{E}_0] - \Pr[\mathsf{E}_1]| \le \frac{2q_e^2}{2^n} .$$

Game G_2. Game G_2 is almost identical to G_1 except that G_2 terminates if $M_1 \neq M_2$ and $H(M_1) = H(M_2)$. The probability of this event is at most

$$|\Pr[\mathsf{E}_1] - \Pr[\mathsf{E}_2]| \leq \epsilon_{\mathsf{cr}}.$$

Game G_3. We define that Game G_3 keeps a set $\mathcal{X} = \{X_i : i \in [q_e]\}$, where the values $X_i \leftarrow E^{-1}(K_i, Y_i)$ represent the responses of backward ideal-cipher queries with key K_i and input Y_i. Game G_3 is almost identical to G_2 except that G_3 terminates also if $H(M_1) \in \mathcal{X}$ or $H(M_2) \in \mathcal{X}$. Then, from the definition of everywhere-pre-image resistance, we have

$$|\Pr[\mathsf{E}_2] - \Pr[\mathsf{E}_3]| \leq q_e \epsilon_{\mathsf{epre}}.$$

Finally, we study the probability that $\mathcal{A}$ wins in Game G_3. Then, one of the following cases must have occurred.

- Case 1: $K_{m_1} = K_{m_2}$. This case implies that $M_1 \neq M_2$. If $H(M_1) = H(M_2)$, there exists an adversary against the collision resistance of the hash function H. Otherwise, if $H(M_1) \neq H(M_2)$, a collision of the tags $T_1 = T_2$ is impossible since E is a permutation for the same key K_m. This case cannot happen: if a hash collision existed, the game would terminate.

- Case 2: $K_{m_1} \neq K_{m_2}$. In this case, a collision can happen only if $E(K_{m_1}, H(M_1)) = E(K_{m_2}, H(M_2))$. Note that $H(M_1) \notin \mathcal{X}$ and $H(M_2) \notin \mathcal{X}$ by assumption. Thus, a collision can happen only randomly with a probability of

$$\Pr[\mathsf{E}_3] = \Pr[E(K_{m_1}, H(M_1)) = E(K_{m_2}, H(M_2))] = 1/2^n\,.$$

Our result follows from the sum of the individual bounds. □

5.2 Collision Resistance of Hash-then-TBC (HTBC)

The MAC HTBC [BGP+19] is defined as follows. Let $H : \mathcal{M} \to \{0,1\}^n \times \{0,1\}^\tau$ be a hash function and $\widetilde{E} : \mathcal{K}_m \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be a TBC. Then, for a given message M and keys $K_m \in \mathcal{K}_m$, $\mathsf{HTBC}[H, \widetilde{E}]_{K_m}$ computes a tag T as

$$T \leftarrow \widetilde{E}(K_m, W, V) \quad \text{where} \quad (V, W) \leftarrow H(M)\,.$$

We define the collision-resistance game similarly to that for HBC.

Theorem 3. Let $\widetilde{E} : \mathcal{K}_m \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be an ideal cipher and $H : \mathcal{M} \to \{0,1\}^\tau \times \{0,1\}^n$ be a $(\epsilon_{\mathsf{cr}}, t_1)$-collision-resistant and $(\epsilon_{\mathsf{epre}}, t_2)$-everywhere-preimage-resistant hash function. Then, for any adversary $\mathcal{A}$ that runs in time t and makes $q_e \leq 2^{n-1}$ queries to the ideal cipher, such that $t_1 = O(t + q_e)$ and $t_2 = O(t + q_e)$, $\mathsf{HTBC}[H, \widetilde{E}]$ is (ϵ, t)-collision-resistant for

$$\epsilon \leq q_e \epsilon_{\mathsf{epre}} + \epsilon_{\mathsf{cr}} + \frac{2q_e^2 + 1}{2^n}\,.$$

Proof. Suppose $\mathcal{A}$ outputs (K_{m_1}, M_1) and (K_{m_2}, M_2) such that $\mathsf{HTBC}[H, \widetilde{E}](K_{m_1}, M_1) = \mathsf{HTBC}[H, \widetilde{E}](K_{m_2}, M_2)$. Again, we define a sequence of hybrid games G_0 through G_3 and define E_i as the event that the adversary wins in game G_i.

Game G_0. This is the real-world game.

Game G_1. Game G_1 differs from G_0 in the fact that it terminates if one of the following events happens during the ideal-cipher queries.

- Two forward queries with different keys produce the same output. This probability is bounded by $q_e^2/2^n$.
- A backward query with input T is followed by a forward query with output T with a different key. This probability is also bounded by $q_e^2/2^n$.

We obtain

$$|\Pr[\mathsf{E}_0] - \Pr[\mathsf{E}_1]| \leq \frac{2q_e^2}{2^n}.$$

Game G_2. Game G_2 is almost identical to G_1 but adds the fact that it terminates if $M_1 \neq M_2$ and $H(M_1) = H(M_2)$. The probability of this event can be bounded by

$$|\Pr[\mathsf{E}_1] - \Pr[\mathsf{E}_2]| \leq \epsilon_{\mathsf{cr}}\,.$$

Game G_3. We adopt the definition of the set of backward-query responses $\mathcal{X}$ from Game G_3 of the proof of HBC. Besides it, Game G_3 adds to G_2 only the fact that G_3 also terminates if $H(M_1) \in \mathcal{X}$ or $H(M_2) \in \mathcal{X}$. From the definition of everywhere-pre-image resistance, we obtain

$$|\Pr[\mathsf{E}_2] - \Pr[\mathsf{E}_3]| \leq q_e \epsilon_{\mathsf{epre}}\,.$$

Finally, we study the probability that $\mathcal{A}$ wins in Game G_3. Similar to Theorem 2, the adversary cannot win if $(K_{m_1}, W_1) = (K_{m_2}, W_2)$. However, if $K_{m_1} \neq K_{m_2}$, a collision can happen if $\widetilde{E}(K_{m_1}, W_1, V_1) = \widetilde{E}(K_{m_2}, W_2, V_2)$. For this event, we have to consider two mutually exclusive cases that cover all possibilities as follows.

- Case 1: $M_1 = M_2$. In this case, we have $H(M_1) = H(M_2) = (V, W)$. Thus, the adversary will be successful if it can find two keys K_{m_1}, K_{m_2} such that $\widetilde{E}(K_{m_1}, W, V) = \widetilde{E}(K_{m_2}, W, V)$. This is impossible since the game would terminate as defined in either G_1 or G_3.
- Case 2: $M_1 \neq M_2$. If $H(M_1) = H(M_2)$, the game would terminate as defined in G_2. Otherwise, a collision can happen only if $\widetilde{E}(K_{m_1}, W_1, V_1) = \widetilde{E}(K_{m_2}, W_2, V_2)$. If these two queries had appeared in any ideal-cipher queries, the conditions that allowed this collision to occur would have led the game to terminate.

If $K_{m_1} = K_{m_2}$ and $W_1 \neq W_2$, then $H(M_1) = H(M_2)$ and $M_1 = M_2$ are impossible, but the analysis of the case when $M_1 \neq M_2$ is the same. This means that $\widetilde{E}(K_{m_1}, W_1, V_1) = \widetilde{E}(K_{m_2}, W_2, V_2)$ holds but at least one of these queries must have not appeared in any ideal-cipher query, as the game would have terminated otherwise. If none of the above happens, then a collision can happen only randomly with probability at most

$$\Pr[\mathsf{E}_3] = \Pr[\widetilde{E}(K_{m_1}, W_1, V_1) = \widetilde{E}(K_{m_2}, W_2, V_2)] \leq \frac{1}{2^n}\,.$$

Our claim in Theorem 3 follows from adding E_3 to the transition differences. □

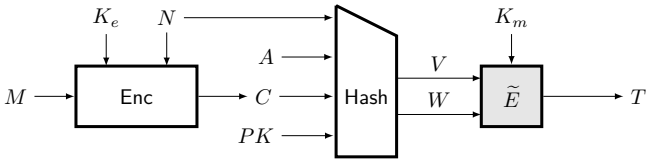

Figure 3: High-level structure of TEDT.

5.3 Collision Resistance of TEDT Variants

TEDT. The authentication part of the original TEDT [BGP+20] (depicted in Figure 3) and Romulus-T [IKMP20] follow exactly our HTBC format. Therein, the hash function H is based on Hirose's compression function and Merkle-Damgård strengthening, which outputs a $2n$-bit value $V\|W$. Note that for any tuple of key and nonce, the adversary can choose a suitable message to achieve any desirable ciphertext. Thus, finding (K_1, U_1) and (K_2, U_2), where $U_i \leftarrow \mathsf{pad}(N_i, A_i, C_i, PK)$ and PK denotes the public key for multi-user security, that lead to the same tag T is equivalent to breaking CMT-4 security.

TEDT2. The authentication function of TEDT2 [Lis21] (depicted in Figure 4) also follows our HTBC format, except for the fact that $8\|N\|W$ is used as a tweak in the final TBC call. Let N_1 and N_2 be two nonces corresponding to the same (C, T) output. Then, we will have two cases:

Case 1: $N_1 = N_2$. Then, the analysis is exactly the same as in Theorem 3.

Case 2: $N_1 \neq N_2$. In this case, N can be seen as part of the hash output. Thus, in this case, the commitment can be broken only by finding a collision in the ideal-cipher queries. Then, we consider two subcases depending on ideal-cipher queries concerning the T-producing TBC call:

- Both queries are forward ideal-cipher queries: Here, the adversary will be successful only if it can find two different tweaks producing the same tag. The success probability is upper bounded by $\frac{q_e^2}{2^n}$.
- At least one of the queries is a backward ideal-cipher query: In this case, the success probability can be upper bounded by ϵ_{epre}.

Thus, the analysis is the same as for HTBC with a hash function $H^{'}(N, A, C) = N\|H(A, C)$.

5.4 Collision Resistance of LRMAC1

The MAC LRMAC1 [BGPS21] is defined as follows. Let $H : \mathcal{M} \to \{0,1\}^\tau \times \{0,1\}^n$ be a hash function and $\widetilde{E} : \mathcal{K}_m \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be a TBC. Given $K_m \in \mathcal{K}_m$ and a message $M \in \mathcal{M}$, $\mathsf{LRMAC1}[H, \widetilde{E}]$ computes the authentication tag as

$$T \leftarrow \widetilde{E}(K_m, V, 0^n) \quad \text{where} \quad V \leftarrow H(M)\,.$$

We define the collision-resistance game similarly as in the case of HBC.

Theorem 4. Let $\widetilde{E} : \mathcal{K}_m \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be an ideal cipher and $H : \mathcal{M} \to \{0,1\}^\tau \times \{0,1\}^n$ be a $(\epsilon_{\mathsf{cr}}, t_1)$-collision-resistant hash function. Then, for any adversary $\mathcal{A}$ that runs in time t and makes $q_e \leq 2^{n-1}$ queries to the ideal cipher, such that $t_1 = O(t+q_e)$, $\mathsf{LRMAC1}[H, \widetilde{E}]$ is (ϵ, t)-collision-resistant for

$$\epsilon \leq \epsilon_{\mathsf{cr}} + \frac{q_e^2 + 2q_e + 5}{2^n}\,.$$

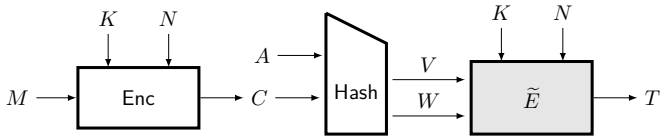

Figure 4: High-level structure of TEDT2.

Proof. Suppose $\mathcal{A}$ outputs (K_{m_1}, M_1) and (K_{m_2}, M_2) such that $\mathsf{LRMAC1}[H, \widetilde{E}](K_{m_1}, M_1) = \mathsf{LRMAC1}[H, \widetilde{E}](K_{m_2}, M_2)$. Again, we will define a sequence of hybrid games and use E_i as the event that the adversary wins in game G_i.

Game G_0. The real-world game.

Game G_1. The game terminates if one of the following events happens during the ideal-cipher queries.

- Two forward queries with different keys produce the same output. This probability is bounded by $q_e^2/2^n$.
- A backward query with input T outputs 0^n. This probability is bounded by $2q_e/2^n$.

It follows that

$$|\Pr[\mathsf{E}_0] - \Pr[\mathsf{E}_1]| \leq \frac{q_e^2 + 2q_e}{2^n}\,.$$

Game G_2. Game G_2 is almost identical to G_1 but adds the fact that G_2 will terminate if $M_1 \neq M_2$ and $H(M_1) = V_1 = V_2 = H(M_2)$. The probability of this event is bounded by

$$|\Pr[\mathsf{E}_1] - \Pr[\mathsf{E}_2]| \leq \epsilon_{\mathsf{cr}}\,.$$

Finally, we study the probability that $\mathcal{A}$ wins in G_2. We have to consider only the case that $V_1 \neq V_2$ as $V_1 = V_2$ would lead the game to terminate.

- Case 1: $(K_1, V_1, 0^n, T)$ appeared in an ideal-cipher query. Then,

$$\Pr\left[\widetilde{E}(K_2, V_2, 0^n) = T\right] \leq \frac{1}{2^n - q_2} \leq \frac{2}{2^n}\,.$$

- Case 2: $(K_2, V_2, 0^n, T)$ appeared in an ideal-cipher query. Then,

$$\Pr\left[\widetilde{E}(K_1, V_1, 0^n) = T\right] \leq \frac{1}{2^n - q_1} \leq \frac{2}{2^n}\,.$$

- Case 3: neither TBC call appeared in any ideal-cipher query. Then,

$$\Pr\left[\widetilde{E}(K_1, V_1, 0^n) = \widetilde{E}(K_2, V_2, 0^n)\right] \leq \frac{1}{2^n}\,.$$

Thus, $\Pr[\mathsf{E}_2] \leq 5/2^n$. The bound in Theorem 4 follows from adding it to all transition probabilities. □

Increasing the tag size and bit-security level. In all considered MACs, the security is bounded by half of the tag size. One way around this limitation is to use a Double-Block-Length (DBL) construction for the TGF. For instance, the TGF of LRMAC1 can be replaced by $\widetilde{E}(K, H(M), 0^n) \| \widetilde{E}(K, H(M), 1^n)$, that is Hirose's DBL compression function with an initial value of 0^n [Hir06]. The construction is still invertible and compatible with CIML2 security. Similar standard constructions can be found for HBC and HTBC.

Table 2: Different variants of KET and the requirements on their components for CMT-4 security. (R)CR = (right) collision resistance.

	Component				
Scheme	KeyGen	KDF	Enc	TGF	Result
KET-1	-	CR	RCR	CR	Theorem 5
KET-1a	LCR	CR	RCR	partial CR	Theorem 6
KET-2	RCR	-	RCR	CR	Theorem 7
KET-2a	RCR	CR	partial RCR	CR	Theorem 8

6 CMT-4 Security of Single-pass Leveled Schemes

In this section, we study the second blueprint KET from Figure 2. We call this blueprint KET as a short-hand for its three components: a KDF, an Enc function, and a TGF. KET can be seen as the paradigm underlying single-pass leveled leakage-resilient schemes such as Triplex [SPS+22] or Multiplex [PSS24].

In this section, we establish three goals. First, we show that the KET composition is CMT-4-secure when each component satisfies a specific set of collision-resistance properties. Second, we show that it can fulfill the compact commitment, wherein verifying the tag suffices to verify the commitment. Finally, we show that if the keys used in the first and last components are identical (or generated by KeyGen having specific CR properties), we can relax the collision-resistance requirements for certain components.

6.1 CMT-4 Security of the Generic KET scheme

We begin with the generic KET scheme wherein the keys in the KDF and the TGF are independent, i.e. no constraints are imposed on their keys. For such schemes, we demonstrate that achieving CMT-4 security requires collision resistance in all three components, KDF, Enc, and TGF, with the minor relaxation that we require only right collision resistance for Enc, i.e. collision resistance for the part of its outputs that are used in the TGF.

Definition 2. Let $\Pi[\mathsf{KeyGen}, \mathsf{KDF}, \mathsf{Enc}, \mathsf{TGF}]$ be a nonce-based AEAD scheme. If, for a given key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \mathcal{A}$, and a message $M \in \mathcal{M}$, it encrypts M to a ciphertext (C, T) as

$$(K_e, K_m) \leftarrow \mathsf{KeyGen}(K) \quad IV \leftarrow \mathsf{KDF}(K_e, N),$$
$$(C, FV) \leftarrow \mathsf{Enc}(IV, A, M), \quad T \leftarrow \mathsf{TGF}(K_m, FV)\,,$$

then, we call Π a KET-1 scheme.

We will study four relevant variants of this scheme which differ in their assumptions posed on their individual components. Table 2 summarizes their properties.

Theorem 5. Let $\Pi[\mathsf{KeyGen}, \mathsf{KDF}, \mathsf{Enc}, \mathsf{TGF}]$ be a KET-1 scheme such that

- the KDF is $(\epsilon_{\mathsf{kdf}}, t_1)$-collision-resistant,
- Enc is $(\epsilon_{\mathsf{enc}}, t_2)$-right collision-resistant, and
- the TGF is $(\epsilon_{\mathsf{tgf}}, t_3)$-collision-resistant

for some $t_1 = O(t)$, $t_2 = O(t)$, and $t_3 = O(t)$. Then, for any t-bounded CMT-4 adversary $\mathcal{A}$ against Π it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \epsilon_{\mathsf{kdf}} + \epsilon_{\mathsf{enc}} + \epsilon_{\mathsf{tgf}}\,. \tag{1}$$

Proof. Since there is no restriction on KeyGen inputs and outputs, the CMT-4 security for the context (K, N, A, M) follows from the CMT-4 security of the context (K_m, K_e, N, A, M). Suppose an adversary $\mathcal{A}$ outputs challenge values (N_1, A_1, K_1, M_1) and (N_2, A_2, K_2, M_2) with corresponding ciphertexts (C_1, T_1) and (C_2, T_2). We upper bound the probability that $(C_1, T_1) = (C_2, T_2) = (C, T)$. We consider the following disjoint cases that cover all possibilities.

Case 1: $K_{m_1} \neq K_{m_2}$. In this case, there must be a collision against the TGF.

Case 2: $K_{m_1} = K_{m_2}$, $(A_1.M_1) \neq (A_2, M_2)$. Then, we consider the following subcases.

- Case 2a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 2b: $FV_1 = FV_2$. In this case, there is a right collision against the Enc function.

Case 3: $K_{m_1} = K_{m_2}$, $(A_1.M_1) = (A_2, M_2)$, $(K_{e_1}, N_1) \neq (K_{e_2}, N_2)$. Again, we study two subcases.

- Case 3a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 3b: $FV_1 = FV_2$, $IV_1 \neq IV_2$. Then, there is a right-output collision against the Enc function.
- Case 3c: $FV_1 = FV_2$, $IV_1 = IV_2$. Then, there is a collision against the KDF.

For each of the cases above, the advantage of the adversary is bounded by the collision resistance property of the individual components, as given in Equation 1. □

6.2 CMT-4 Security of the KET-1a scheme

Theorem 5 does not require any collision resistance property for KeyGen and holds even when the keys K_e and K_m are independent. However, if KeyGen is left-collision-resistant, we can lift the requirement of full i.e. (FV, K_m)-collision resistance from the TGF. Instead, left collision resistance on the values FV will suffice, as captured by the following theorem.

Theorem 6. Let Π[KeyGen, KDF, Enc, TGF] be a KET-1a scheme such that

- KeyGen is (ϵ_k, t_1)-left collision-resistant,
- the KDF is $(\epsilon_{\mathsf{kdf}}, t_2)$-collision-resistant,
- Enc is $(\epsilon_{\mathsf{enc}}, t_3)$-right collision-resistant, and
- the TGF is $(\epsilon_{\mathsf{tgf}}, t_4)$-$FV$-collision-resistant (partial CR only on input FV but not on input K_m).

for some $t_1 = O(t)$, $t_2 = O(t)$, $t_3 = O(t)$, and $t_4 = O(t)$. Then, for any t-bounded CMT-4 adversary $\mathcal{A}$ against Π, it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \epsilon_{\mathsf{kdf}} + \epsilon_{\mathsf{enc}} + \epsilon_{\mathsf{tgf}} + \epsilon_k \,. \tag{2}$$

Proof. Suppose an adversary $\mathcal{A}$ outputs challenge values (N_1, A_1, K_1, M_1) and (N_2, A_2, K_2, M_2) with corresponding ciphertexts (C_1, T_1) and (C_2, T_2). We upper bound the probability that (C_1, T_1) = (C_2, T_2) = (C, T), where we have to consider the following disjoint cases.

Case 1: $(A_1, M_1) \neq (A_2, M_2)$. We have to study two subcases.

- Case 1a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 1b: $FV_1 = FV_2$. Then, there is a right collision against Enc.

Case 2: $(A_1, M_1) = (A_2, M_2)$, $(K_1, N_1) \neq (K_2, N_2)$. Since KeyGen is left collision-resistant, the probability of $K_1 \neq K_2 \ \wedge \ K_{e_1} = K_{e_2}$ is bounded by ϵ_k. Accounting for ϵ_k by a standard hybrid argument, we can safely assume $K_1 \neq K_2$ and therefore $K_{e_1} \neq K_{e_2}$ in the remainder. Thus, this case reduces to $(K_{e_1}, N_1) \neq (K_{e_2}, N_2)$.

- Case 2a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 2b: $FV_1 = FV_2$, $IV_1 \neq IV_2$. Then, there is a right collision against the Enc.
- Case 2c: $FV_1 = FV_2$, $IV_1 = IV_2$. Then, there is a collision against the KDF, as $(K_{e_1}, N_1) \neq (K_{e_2}, N_2)$.

For each of the cases above, the advantage of the adversary is bounded by the collision resistance property of its respective three components, as given in Equation 2. □

Next, we consider variants of KET that use the nonce as an additional input of the encryption function. For those variants, collision resistance of the KDF is not necessary. This is intuitive since we can view the next scheme as KET-1a where N is appended to the output of the KDF.

Definition 3. Let Π[KeyGen, KDF, Enc, TGF] be a nonce-based AEAD scheme. If, for a given key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \mathcal{A}$, Π encrypts a message $M \in \mathcal{M}$ to a ciphertext (C, T) as

$$\begin{aligned} (K_e, K_m) &\leftarrow \mathsf{KeyGen}(K) & IV &\leftarrow \mathsf{KDF}(K_e, N), \\ (C, FV) &\leftarrow \mathsf{Enc}(IV, N, A, M), & T &\leftarrow \mathsf{TGF}(K_m, FV)\,, \end{aligned}$$

then, we call Π an KET-2 scheme.

Theorem 7. Let Π[KeyGen, KDF, Enc, TGF] be an KET-2 scheme such that

- the KeyGen function is (ϵ_k, t_1)-right collision-resistant,
- and Enc is $(\epsilon'_{\mathsf{enc}}, t_2)$-right collision-resistant, and
- the TGF is $(\epsilon_{\mathsf{tgf}}, t_3)$-collision-resistant

for $t_1 = O(t)$, $t_2 = O(t)$ and $t_3 = O(t)$. Then, for any adversary $\mathcal{A}$ running in time at most t against the CMT-4 security of Π, it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \epsilon'_{\mathsf{enc}} + \epsilon_{\mathsf{tgf}} + \epsilon_k\,. \tag{3}$$

Proof. Suppose $\mathcal{A}$ outputs challenge values (N_1, A_1, K_1, M_1) and (N_2, A_2, K_2, M_2) with corresponding ciphertexts (C_1, T_1) and (C_2, T_2). We bound the probability that $(C_1, T_1) = (C_2, T_2) = (C, T)$. We consider the following disjoint cases.

Case 1: $K_1 \neq K_2$. Here, we have the following subcases.

- Case 1a: $K_{m_1} = K_{m_2}$. Then, there is a right collision on K_m against the KeyGen function.
- Case 1b: $K_{m_1} \neq K_{m_2}$. Then, there is a collision against the TGF.

Case 2: $K_1 = K_2$, $(N_1, A_1, M_1) \neq (N_2, A_2, M_2)$. Again, we have two subcases.

- Case 2a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 2b: $FV_1 = FV_2$. Then, there is a right-output collision against Enc.

For each of the cases above, the advantage of the adversary is bounded by the collision resistance property of the three components, as given in Equation 3. □

Finally, we consider a special case of KET-2 that we call KET-2a, where Enc is only collision-resistant when IV, A, or C change, i.e. it may be easy to find (IV, A, C, N_1) and (IV, A, C, N_2) such that $FV_1 = FV_2$. However, if $(IV_1, A_1) \neq (IV_2, A_2)$, then collisions are hard to find. The following theorem demonstrates that, despite this restrictive assumption on the collision resistance of Enc, we can still attain CMT-4 security by imposing a milder condition. In this case, it is essential to also assume that KDF is collision-resistant.

Theorem 8. Let $\Pi[\mathsf{KeyGen}, \mathsf{KDF}, \mathsf{Enc}, \mathsf{TGF}]$ be a KET-2a scheme such that

- the KeyGen function is (ϵ_k, t_1)-right collision-resistant,
- the KDF is $(\epsilon_{\mathsf{kdf}}, t_2)$-collision-resistant,
- Enc is $(\epsilon'_{\mathsf{enc}}, t_3)$-$(IV, A)$-right collision-resistant i.e Enc is RCR only on input (IV, A), and
- the TGF is $(\epsilon_{\mathsf{tgf}}, t_4)$-collision-resistant

for $t_1 = O(t)$, $t_2 = O(t)$, $t_3 = O(t)$, and $t_4 = O(t)$. Then, for any adversary $\mathcal{A}$ running in time at most t against the CMT-4 security of Π, it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \epsilon_{\mathsf{kdf}} + \epsilon'_{\mathsf{enc}} + \epsilon_{\mathsf{tgf}} + \epsilon_k \,. \tag{4}$$

Proof. Suppose an adversary $\mathcal{A}$ outputs challenge values (N_1, A_1, K_1, M_1) and (N_2, A_2, K_2, M_2) with corresponding ciphertexts (C_1, T_1) and (C_2, T_2). We bound the probability that $(C_1, T_1) = (C_2, T_2) = (C, T)$ with the following cases.

Case 1: $K_1 \neq K_2$. Here, we have the following two subcases.

- Case 1a: $K_{m_1} = K_{m_2}$. Then, there is a right-output collision against the KeyGen function, i.e. on K_m.
- Case 1b: $K_{m_1} \neq K_{m_2}$. Then, there is a collision against the TGF.

Case 2: $K_1 = K_2$, $(N_1, A_1, M_1) \neq (N_2, A_2, M_2)$.

- Case 2a: $FV_1 \neq FV_2$. Then, there is a collision against the TGF.
- Case 2b: $FV_1 = FV_2$.
  - If $IV_1 \neq IV_2$, there is a (IV, A)-right-output collision against the Enc function.
  - Otherwise, if $IV_1 = IV_2$ and $N_1 \neq N_2$ holds, there is collision against the KDF.
  - Finally, if $IV_1 = IV_2$ and $N_1 = N_2$ hold, we must have $A_1 \neq A_2$ by injectivity of Enc over the message space when the other parameters remain unchanged. $A_1 \neq A_2$ implies that there will be a collision against (IV, A)-right-output collision of Enc.

For each of the cases above, the advantage of the adversary is bounded by the collision resistance property of the four components, as given in Equation 4. □

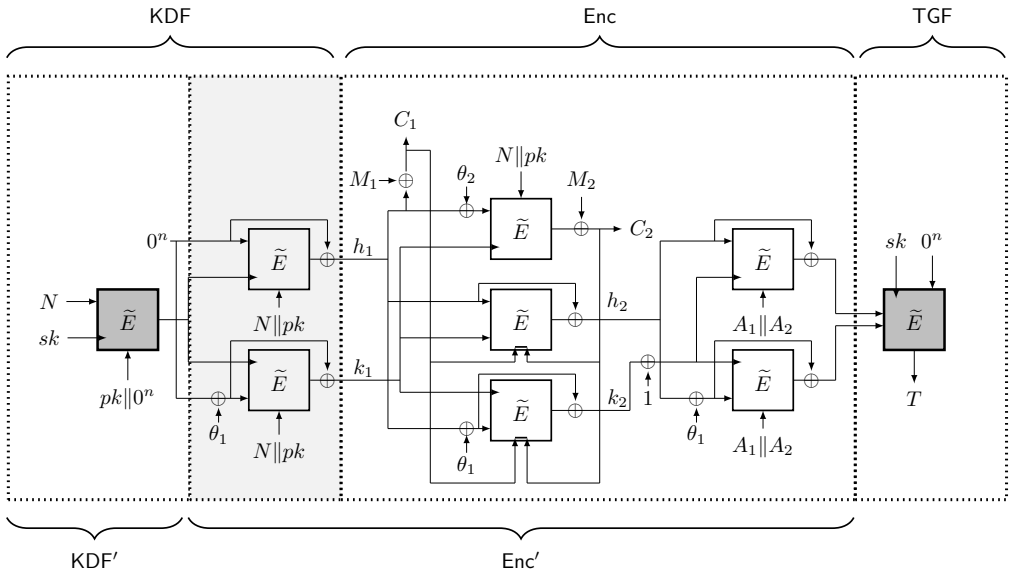

Figure 5: Encryption with Triplex.

7 Triplex as an Instantiation of KET-2

In this section, we demonstrate the usefulness of the KET blueprint by showing that the recent single-pass scheme Triplex [SPS+22] can be viewed as an instance of KET-2.

Triplex. Triplex operates with a KDF that consists of three TBCs, a protected call followed by two parallel calls to an unprotected TBC. The KDF takes a key $K = sk\|pk$ – that combines a secret part sk with a public part pk for higher multi-user security – and a nonce N and produces a $2n$-bit output $IV = h_1\|k_1$. The encryption function of Triplex takes various inputs including pk, N, A, M, and $IV = h_1\|k_1$, and outputs a ciphertext C along with $FV = V\|W$. Its TGF is essentially a single TBC call. It takes sk as the key, $V\|W$ as the tweak, and a fixed input 0^n to generate a tag T. Note that both Enc and KDF take N as input and both the KDF and the TGF use the same key sk.

There are multiple ways to view Triplex, and each one leads to the application of a different theorem. We will view the CMT-4 security of Triplex as an application of Theorem 7. We can consider pk as part of the nonce instead of the key since it is not utilized as a key anywhere. This simplification allows us to view Triplex as a specific instance of the generic KET-2 construction. According to Theorem 7, for achieving CMT-4 security, we need to demonstrate collision resistance of the TGF and right-output collision resistance of Enc.

Corollary 1. Let $\Pi[\mathsf{KDF}, \mathsf{Enc}, \mathsf{TGF}]$ denote Triplex. Then, there exists an $(\epsilon_{\mathsf{cr}}, t_1)$-collision-resistant hash function H, such that for any adversary $\mathcal{A}$ running in time at most t against the CMT-4 security of Triplex, it holds that

$$\mathbf{Adv}^{\mathtt{cmt4}}_{\Pi}(\mathcal{A}) \leq \frac{q_e^2 + 2q_e + 5}{2^n} + \epsilon_{\mathsf{cr}}\,,$$

where $t_1 = O(t)$.

Proof. First, we will redefine the KDF function of Triplex. This is done by moving the two parallel TBC calls from out of the KDF into the Enc function. We denote the modified KDF and Enc functions as KDF′ and Enc′, respectively, as visiaulized in Figure 5. This

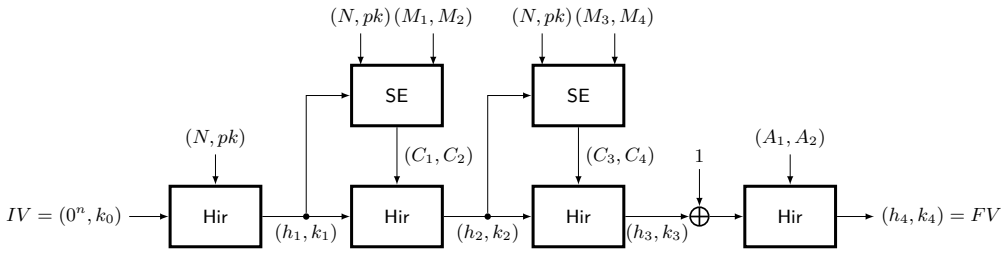

Figure 6: Alternative visualization of the modified encryption function Enc' of Triplex.

change does not affect the scheme's security but only moves the boundary of where the KDF ends and Enc begins. In this representation, KDF is not collision-resistant, and we apply Theorem 7. We know that $\epsilon_k = 0$ since sk is used as both K_e and K_m. Moreover, the collision resistance of the TGF is similar to that of the TGF used in LRMAC1. In other words, we obtain from the analysis of Theorem 4 that

$$\epsilon_{\mathsf{tgf}} \leq \frac{q_e^2 + 2q_e + 5}{2^n} .$$

What remains is to bound ϵ_{enc}. Note that the function Enc' can be visualized as shown in Figure 6, wherein the bottom part is the Triplex hash function of the input $\mathsf{pad}(IV, N, pk, A, C)$ for some injective padding function. The top symmetric encryption (SE) component computes ciphertext C being input to the hash function. Thus, if $N_1 = N_2$, the top part (SE) is bijective and FV-collision-resistant. If $N_1 \neq N_2$, it is still FV-collision-resistant since N is part of the input to the hash function. Then,

$$\epsilon_{\mathsf{enc}} \leq \epsilon_{\mathsf{cr}} .$$

Finally, the hash function H used in Triplex is the Merkle-Damgård with Permutation (MDP) hash function [HPY07] instantiated with Hir, Hirose's double-block-length function [Hir06]. From the indifferentiability of this MDPH hash function, we have that ϵ_{cr} is negligible, which implies the commitment security of Triplex. □

8 Conclusion

In this paper, we studied the CMT-4 security of two families of leveled leakage-resilient schemes: Grade-3 schemes based on EtM and single-pass Grade-2 schemes. In both cases, we give positive results. We show that EtM is committing as long as the keys satisfy a particular definition of dependence and the MAC is collision-resistant. We give positive results on the collision resistance of different leakage-resilient MACs: HBC, HTBC and LRMAC1, and apply this to show the CMT-4 security of TEDT. We also discuss how to increase the security by increasing the tag size. For single-pass schemes, we give different variants with different assumptions on their components and show that the recently proposed scheme, Triplex, achieves CMT-4 security up to half the tag size.

We believe our work shows an interesting connection between context commitment and leakage-resilient schemes. Even though the two security goals are different and not implied by each other, the underlying design principles allow for efficient schemes that achieve both goals.

An interesting future direction is to study how to design leakage-resilient schemes that are also committing beyond half the tag size. Another direction is to study if there is a connection that can be derived between CIML2 security and CMT-4 security.

Acknowledgments

We would like to thank Ling Song and Yaobin Shen for organizing the Asian Symmetric Workshop (ASK) 2023 in Guangzhou, China, where this work was initiated. Mustafa Khairallah is funded by the Wallenberg-NTU Postdoctoral Presidential Fellowship. Eik List has been supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – LI 4223/1-1.

References

[BBC+20] Davide Bellizia, Olivier Bronchain, Gaëtan Cassiers, Vincent Grosso, Chun Guo, Charles Momin, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Mode-Level vs. Implementation-Level Physical Security in Symmetric Cryptography - A Practical Guide Through the Leakage-Resistance Jungle. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO I*, volume 12170 of *Lecture Notes in Computer Science*, pages 369–400. Springer, 2020.

[BGP+19] Francesco Berti, Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Strong Authenticity with Leakage Under Weak and Falsifiable Physical Assumptions. In Zhe Liu and Moti Yung, editors, *Inscrypt*, volume 12020 of *Lecture Notes in Computer Science*, pages 517–532. Springer, 2019.

[BGP+20] Francesco Berti, Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. TEDT, a Leakage-Resist AEAD Mode for High Physical Security Applications. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(1):256–320, 2020.

[BGPS21] Francesco Berti, Chun Guo, Thomas Peters, and François-Xavier Standaert. Efficient Leakage-Resilient MACs Without Idealized Assumptions. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT II*, volume 13091 of *Lecture Notes in Computer Science*, pages 95–123. Springer, 2021.

[BH22] Mihir Bellare and Viet Tung Hoang. Efficient Schemes for Committing Authenticated Encryption. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT II*, volume 13276 of *Lecture Notes in Computer Science*, pages 845–875. Springer, 2022.

[DEM+17] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, and Thomas Unterluggauer. ISAP – Towards Side-Channel Secure Authenticated Encryption. *IACR Trans. Symmetric Cryptol.*, 2017(1):80–105, 2017.

[DEMS21] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schläffer. Ascon v1.2: Lightweight authenticated encryption and hashing. *J. Cryptol.*, 34(3):33, 2021.

[DGRW18] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast Message Franking: From Invisible Salamanders to Encryptment. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO I*, volume 10991 of *Lecture Notes in Computer Science*, pages 155–186. Springer, 2018.

[DJS19] Jean Paul Degabriele, Christian Janson, and Patrick Struck. Sponges Resist Leakage: The Case of Authenticated Encryption. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT II*, volume 11922 of *Lecture Notes in Computer Science*, pages 209–240. Springer, 2019.

[FOR17] Pooya Farshim, Claudio Orlandi, and Razvan Rosie. Security of Symmetric Primitives under Incorrect Usage of Keys. *IACR Trans. Symmetric Cryptol.*, 2017(1):449–473, 2017.

[GLL19] Shay Gueron, Adam Langley, and Yehuda Lindell. AES-GCM-SIV: Nonce Misuse-Resistant Authenticated Encryption. *RFC*, 8452:1–42, 2019.

[GLR17] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. Message Franking via Committing Authenticated Encryption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO III*, volume 10403 of *Lecture Notes in Computer Science*, pages 66–97. Springer, 2017.

[Hir06] Shoichi Hirose. Some Plausible Constructions of Double-Block-Length Hash Functions. In Matthew J. B. Robshaw, editor, *FSE*, volume 4047 of *Lecture Notes in Computer Science*, pages 210–225. Springer, 2006.

[HPY07] Shoichi Hirose, Je Hong Park, and Aaram Yun. A Simple Variant of the Merkle-Damgård Scheme with a Permutation. In Kaoru Kurosawa, editor, *ASIACRYPT*, volume 4833 of *Lecture Notes in Computer Science*, pages 113–129. Springer, 2007.

[IKMP20] Tetsu Iwata, Mustafa Khairallah, Kazuhiko Minematsu, and Thomas Peyrin. Duel of the titans: The romulus and remus families of lightweight AEAD algorithms. *IACR Trans. Symmetric Cryptol.*, 2020(1):43–120, 2020.

[JNPS21] Jérémy Jean, Ivica Nikolic, Thomas Peyrin, and Yannick Seurin. The deoxys AEAD family. *J. Cryptol.*, 34(3):31, 2021.

[KSW23] Juliane Krämer, Patrick Struck, and Maximiliane Weishäupl. Committing AE from Sponges: Security Analysis of the NIST LWC Finalists. Cryptology ePrint Archive, Paper 2023/1525, 2023. https://eprint.iacr.org/2023/1525.

[LGR21] Julia Len, Paul Grubbs, and Thomas Ristenpart. Partitioning oracle attacks. In Michael D. Bailey and Rachel Greenstadt, editors, *30th USENIX Security Symposium, USENIX Security 2021, August 11-13, 2021*, pages 195–212. USENIX Association, 2021.

[Lis21] Eik List. TEDT2 - Highly Secure Leakage-Resilient TBC-Based Authenticated Encryption. In Patrick Longa and Carla Ràfols, editors, *LATINCRYPT Colombia, October 6-8, 2021, Proceedings*, volume 12912 of *Lecture Notes in Computer Science*, pages 275–295. Springer, 2021.

[NRS14] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering Generic Composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 257–274. Springer, 2014.

[PSS24] Thomas Peters, Yaobin Shen, and François-Xavier Standaert. Multiplex: TBC-based Authenticated Encryption with Sponge-Like Rate. *IACR Cryptol. ePrint Arch.*, page 294, 2024.

[Rog04] Phillip Rogaway. Nonce-Based Symmetric Encryption. In Bimal K. Roy and Willi Meier, editors, *FSE*, volume 3017 of *Lecture Notes in Computer Science*, pages 348–359. Springer, 2004.

[RS06] Phillip Rogaway and Thomas Shrimpton. A Provable-Security Treatment of the Key-Wrap Problem. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 373–390. Springer, 2006.

[SCM08] Joseph Salowey, Abhijit Choudhury, and David A. McGrew. AES Galois Counter Mode (GCM) Cipher Suites for TLS. *RFC*, 5288:1–8, 2008.

[SPS+22] Yaobin Shen, Thomas Peters, François-Xavier Standaert, Gaëtan Cassiers, and Corentin Verhamme. Triplex: an Efficient and One-Pass Leakage-Resistant Mode of Operation. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2022(4):135–162, 2022.

[SW24] Patrick Struck and Maximiliane Weishäupl. Constructing Committing and Leakage-Resilient Authenticated Encryption. *IACR Trans. Symmetric Cryptol.*, 2024(1):497–528, 2024.