

RIV for Robust Authenticated Encryption

Farzaneh Abed¹, Christian Forler²,
Eik List¹, Stefan Lucks¹, and Jakob Wenzel¹

¹ Bauhaus-Universität Weimar, ² Hochschule Schmalkalden
¹ <firstname>.<lastname>@uni-weimar.de, ² cforler@hs-schmalkalden.de

Abstract. Typical AE schemes are supposed to be secure when used as specified. However, they can – and often do – fail miserably when used improperly. As a partial remedy, Rogaway and Shrimpton proposed (nonce-)misuse-resistant AE (MRAE) and the first MRAE scheme SIV (“Synthetic Initialization Vector”). This paper proposes RIV (“Robust Initialization Vector”), which extends the generic SIV construction by an additional call to the internal PRF. RIV inherits the full security assurance from SIV, but unlike SIV and other MRAE schemes, RIV is also provably secure when releasing unverified plaintexts. This follows a recent line of research on “*Robust Authenticated Encryption*”, similar to the CAESAR candidate AEZ.

An AES-based instantiation of RIV runs at less than 1.5 cpb on current x64 processors. Unlike the proposed instantiation of AEZ, which gains speed by relying on reduced-round AES, our instantiation of RIV is provably secure under the single assumption of the AES being secure.

Keywords: robustness, subtle authenticated encryption, provable security

1 Introduction

Authenticated Encryption. A secure authenticated encryption (AE) scheme generates ciphertexts that can not be efficiently distinguished from random bit-strings of the same length as the ciphertext and are infeasible to forge. Typical AE schemes are nonce-based [45], i.e., the user is responsible to supply an additional input that must be unique for every encryption. If a nonce ever repeats, the scheme’s security may fully forfeit. While the concept of unique nonces is simple in theory, it is hard to ensure in practice [19], which led to severe security breaches in the past. Rogaway and Shrimpton [46] defined (nonce-)misuse-resistant AE (MRAE) as notion with the goal of providing full authenticity, and privacy up to the detection of repeated encryptions of the same associated data and message under the same nonce and key. Since then, the topic received significant attention by the community, resulting in a large corpus of MRAE schemes, e.g., [6,10,16,20,22,27,28,29,30,33,43,46].

Robustness aspects of AE are not limited to nonce reuse. “One shortcoming of AE as commonly understood is its idealized, all-or-nothing decryption” [7]. Leaking any information about the message before its authentication has been verified

breaks this assumption. At least five noteworthy recent works strengthened the existing security definitions of robustness.¹ Boldyreva et al. [15] (BDPS) studied the effects when multiple distinct error messages are distinguishable in probabilistic or stateful schemes. Andreeva et al. [4] formalized notions that capture the remaining security under *release of unverified plaintexts* (RUP). Hoang et al. [24] defined *robust AE* (RAE) as a notion for the *best achievable* security of an AE scheme with a user-chosen ciphertext expansion. Badertscher et al. [5] investigated RAE with the frameworks by Maurer and Renner [38,39]. Barwell, Page, and Stam [7] defined *subtle AE* (SAE) as a reference framework for the BDPS, RUP, and RAE notions. The SAE definitions comprise leakage beyond information about the invalid plaintext, which allows to model leakage as a property of the decryption *implementation* rather than as a property of the scheme.

Previous Robust AE Schemes. In spite of so much progress regarding stricter security definitions, the portfolio of dedicated robust AE schemes remains still modest. Among the 57 CAESAR submissions, only four candidates consider robustness against leakage of invalid plaintexts: Julius [6] lacks a security proof; POET [1] and APE [3] concern on-line confidentiality, which cannot provide nonce-misuse resistance in the strong sense of Rogaway and Shrimpton, as has been criticized, e.g., by [25]. Only AEZ [24] provides robust AE. Though, AEZ follows a “proof-then-prune” approach: while the security proof assumes a strong block cipher, the performant instantiation employs four-round AES instead. Since AEZ also defines a key schedule, it appears more as a primitive of its own right than as a block-cipher-based AE scheme.

Beyond CAESAR, Bertoni et al. [12] proposed MR. MONSTER BURRITO, a four-round Feistel network with the round-reduced KECCAK- f permutation in duplex-wrap mode, and the sponge in counter mode for encryption. Shrimpton and Terashima [47] proposed Protected IV (PIV), a framework of strong tweakable ciphers (STPRPs), which generalized the Ψ_3 construction by Coron et al. [17]. PIV is fast (comparable with the construction proposed in this work); though, it requires the block-cipher inverse for decryption. Note that theoretically, more robust AE schemes could be constructed. Hoang et al. [24] showed that the well-known Encode-then-Encipher (EtE) [9] approach achieves RAE security when (a hash of) nonce and associated data are used as tweak. In theory, this implies that a secure STPRP can be transformed into a robust AE scheme, which allows to choose from the schemes that have been developed over the previous decade, e.g., in the domains of full-disk and format-preserving encryption.

Contribution. This work proposes a modular framework, called *Robust IV* (RIV), which provides provable SAE security. RIV is an extension of SIV [26,46] that inherits both the simplicity and the naturally strong security properties of SIV and adds robustness against leakage of invalid plaintexts. We propose an instantiation which runs at less than 1.5 clock cycles per byte (cpb) on current x64 processors.

¹ By robustness, we mean resistance against both nonce misuse and decryption leakage beyond the single error information.

Outline. The remainder of this work is structured as follows: after Section 2 recalls the preliminaries, Section 3 describes the generic RIV framework. Section 4 recalls the relevant notions. Section 5 summarizes our formal security analysis. Section 6 details our instantiation, and Section 7 concludes this work.

2 Preliminaries

We use lowercase letters x, y for indices and integers, uppercase letters X, Y for binary strings and functions, and calligraphic uppercase letters \mathcal{X}, \mathcal{Y} for sets. By ε we denote the empty string. We denote the concatenation of binary strings X and Y by $X \| Y$ and the result of their bitwise XOR by $X \oplus Y$. We indicate the length of X in bits by $|X|$, and write X_i for the i -th block, $X[i]$ for the i -th most significant bit of X , and $X[i..j]$ for the bit sequence $X[i], \dots, X[j]$. $X \leftarrow \mathcal{X}$ denotes that X is chosen uniformly at random from the set \mathcal{X} . We define two sets of particular interest: $\text{Perm}(\mathcal{X})$ be the set of all permutations on \mathcal{X} and $\text{Func}(\mathcal{X}, \mathcal{Y})$ the set of all functions $F : \mathcal{X} \rightarrow \mathcal{Y}$. A uniform random function $\rho : \mathcal{X} \rightarrow \mathcal{Y}$ with domain \mathcal{X} and range \mathcal{Y} is a random variable uniformly distributed over $\text{Func}(\mathcal{X}, \mathcal{Y})$. We define by $X_1, \dots, X_j \stackrel{x}{\leftarrow} X$ the injective splitting of the string X into x -bit blocks such that $X = X_1 \| \dots \| X_j$, with $|X_i| = x$ for $1 \leq i \leq j - 1$, and $|X_j| \leq x$.

For an event E , we denote by $\Pr[E]$ the probability of E . We write $\langle x \rangle_m$ for the binary m -bit-string representation of an integer x and $\langle x \rangle$ for the binary n -bit-string representation of x for an integer n that is clear from the context. If not stated otherwise, we assume representations to be encoded in big-endian manner, i.e., the decimal $\langle 135 \rangle$ is encoded to the n -bit string $000..010000111$.

Universal Hashing. Universal hash functions are well-known components for compressing a message while guaranteeing maximal probabilities about output relations. We briefly recall the definitions that are relevant in this work.

Definition 1 (ϵ -Almost-(XOR-)Universal Hash Functions). Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^*$. Let $\mathcal{H} = \{H \mid H : \mathcal{X} \rightarrow \mathcal{Y}\}$ denote a family of hash functions. \mathcal{H} is called ϵ -almost-universal (ϵ -AU) iff for all distinct elements $X, X' \in \mathcal{X}$, it holds that $\Pr_{H \leftarrow \mathcal{H}}[H(X) = H(X')] \leq \epsilon$. \mathcal{H} is called ϵ -almost-XOR-universal (ϵ -AXU) iff for all distinct elements $X, X' \in \mathcal{X}$ and $Y \in \mathcal{Y}$, it holds that $\Pr_{H \leftarrow \mathcal{H}}[H(X) \oplus H(X') = Y] \leq \epsilon$.

Theorem 1 (Theorem 3 from [14]). Let $\mathcal{X}, \mathcal{Y} \subseteq \{0, 1\}^*$. Further, let $\mathcal{H} = \{H \mid H : \mathcal{X} \rightarrow \mathcal{Y}\}$ be a family of ϵ -AXU hash functions. Then, the family $\mathcal{H}' = \{H' \mid H' : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Y}\}$ with $H'(X, Y) := H(X) \oplus Y$, is ϵ -AU.

Nonce-Based Encryption Schemes. A nonce-based encryption scheme [45] is a tuple $\Pi = (\mathcal{E}, \mathcal{D})$ of deterministic encryption and decryption algorithms $\mathcal{E} : \mathcal{K} \times \mathcal{N} \times \mathcal{M} \rightarrow \mathcal{C}$ and $\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{C} \rightarrow \mathcal{M}$, with associated non-empty key space \mathcal{K} , non-empty nonce space \mathcal{N} , and $\mathcal{M}, \mathcal{C} \subseteq \{0, 1\}^*$ denoting message and ciphertext space, respectively. We often write $\mathcal{E}_K^N(M)$ and $\mathcal{D}_K^N(C)$ as short forms

of $\mathcal{E}(K, N, M)$ and $\mathcal{D}(K, N, C)$. An adversary that never repeats a nonce over its encryption queries is called *nonce-respecting*, and *nonce-ignoring* otherwise. We assume for all $K \in \mathcal{K}$, $N \in \mathcal{N}$, $M \in \mathcal{M}$, and $C \in \mathcal{C}$ *length-preservation*, i.e., $|\mathcal{E}_K^N(M)| = |M|$, *correctness*, i.e., $\mathcal{D}_K^N(\mathcal{E}_K^N(M)) = M$, and *tidiness*, i.e., $\mathcal{E}_K^N(\mathcal{D}_K^N(C)) = C$. We call a nonce-based encryption scheme $\Pi = (\mathcal{E}, \mathcal{D})$ *nonce-keystream-based* iff its encryption algorithm derives a keystream $\kappa_N \subseteq \{0, 1\}^*$, with $|\kappa_N| = |M|$, from the given nonce N and computes the ciphertext as $C \leftarrow \kappa_N \oplus M$. Naturally, the decryption algorithm of such an encryption scheme is identical to its encryption algorithm, i.e., $\mathcal{E}_K^N(M) := \mathcal{D}_K^N(M)$ for all $K \in \mathcal{K}$, $N \in \mathcal{N}$, and $M \in \mathcal{M}$.

Nonce-Based AE Schemes. A nonce-based authenticated encryption scheme (with associated data) [44] is a tuple $\tilde{\Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}})$ of a deterministic encryption algorithm $\tilde{\mathcal{E}} : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{M} \rightarrow \mathcal{C} \times \mathcal{T}$, and a deterministic decryption algorithm $\tilde{\mathcal{D}} : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{C} \times \mathcal{T} \rightarrow \mathcal{M} \cup \{\perp\}$, with associated non-empty key space \mathcal{K} , non-empty nonce space \mathcal{N} , and $\mathcal{H}, \mathcal{M}, \mathcal{C} \subseteq \{0, 1\}^*$ denote the header, message, and ciphertext space, respectively. We define a tag space $\mathcal{T} = \{0, 1\}^\tau$ for a fixed $\tau \geq 0$. We often write $\tilde{\mathcal{E}}_K^{N,H}(M)$ and $\tilde{\mathcal{D}}_K^{N,H}(C, T)$ as short forms of $\tilde{\mathcal{E}}(K, N, H, M)$ and $\tilde{\mathcal{D}}(K, N, H, C, T)$. If a given tuple (N, H, C, T) is valid, $\tilde{\mathcal{D}}_K^{N,H}(C, T)$ returns the corresponding plaintext M , and \perp otherwise. We assume that for all $K \in \mathcal{K}$, $N \in \mathcal{N}$, $H \in \mathcal{H}$, and $M \in \mathcal{M}$ holds *stretch-preservation*: if $\tilde{\mathcal{E}}_K^{N,H}(M) = (C, T)$, then $|C| = |M|$ and $|T| = \tau$, *correctness*: if $\tilde{\mathcal{E}}_K^{N,H}(M) = (C, T)$, then $\tilde{\mathcal{D}}_K^{N,H}(C, T) = M$, and *tidiness*: if $\tilde{\mathcal{D}}_K^{N,H}(C, T) = M \neq \perp$, then $\tilde{\mathcal{E}}_K^{N,H}(M) = (C, T)$, for all $C \in \mathcal{C}$ and $T \in \mathcal{T}$. Note that some notions (e.g., [41]) regard an authenticated ciphertext C with $|C| = |M| + \tau$ instead of an explicitly separated tuple (C, T) .

Subtle AE Schemes. Barwell et al. defined a subtle AE scheme $\tilde{\Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}}, \Lambda)$ as a tuple of deterministic encryption and decryption algorithms $\tilde{\mathcal{E}}$ and $\tilde{\mathcal{D}}$ as above², and an additional deterministic leakage algorithm $\Lambda : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{C} \times \mathcal{T} \rightarrow \{\top\} \cup \mathcal{L}$, with a non-empty leakage space \mathcal{L} and a symbol $\top \notin \mathcal{L}$ to indicate a valid input. This means, for all $K \in \mathcal{K}$, $N \in \mathcal{N}$, $H \in \mathcal{H}$, $C \in \mathcal{C}$, and $T \in \mathcal{T}$ holds: if $\Lambda_K^{N,H}(C, T) = \top$, then $\tilde{\mathcal{D}}_K^{N,H}(C, T) \neq \perp$; moreover, it holds that if $\Lambda_K^{N,H}(C, T) \neq \top$, then $\tilde{\mathcal{D}}_K^{N,H}(C, T) = \perp$.

3 Definition of RIV

Definition 2 (RIV). Let $d, n, \tau \geq 1$. Let $\mathcal{K}_1, \mathcal{K}_2$, and $\mathcal{K} = \mathcal{K}_1 \times \mathcal{K}_2$ be non-empty key sets, \mathcal{N} a non-empty nonce space, $\{0, 1\}^d$ the non-empty domain space, and $\mathcal{H}, \mathcal{M}, \mathcal{C} \subseteq \{0, 1\}^*$ header, message, and ciphertext spaces, respectively, and $\mathcal{T} = \{0, 1\}^\tau$ a tag space. Let further $F : \mathcal{K}_1 \times \{0, 1\}^d \times \mathcal{N} \times \mathcal{H} \times \mathcal{M} \rightarrow \{0, 1\}^n$ be a function and $\Pi = (\mathcal{E}, \mathcal{D})$ a nonce-based encryption scheme with associated key space \mathcal{K}_2 and nonce space $\{0, 1\}^\tau$. Let $F_K^i(\cdot, \cdot, \cdot)$ denote $F_K(\langle i \rangle_d, \cdot, \cdot, \cdot)$. Then, we define the AE scheme $\text{RIV}_{F, \Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}})$ with encryption algorithm

² Though, their definitions denote the authenticated ciphertext (C, T) as C .

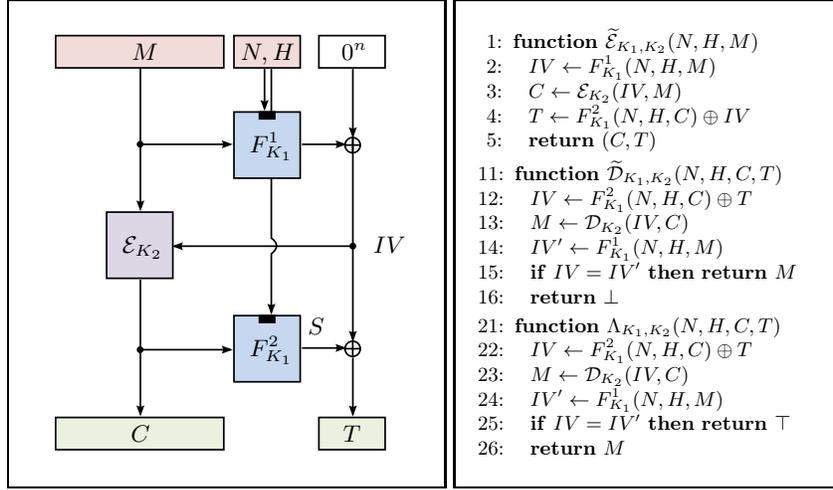


Fig. 1: **Left:** Schematic illustration of the encryption of $\text{RIV}_{F, \Pi}$ with a PRF F and a nonce-based encryption scheme $\Pi = (\mathcal{E}, \mathcal{D})$. **Right:** Definition of encryption and decryption algorithms of $\text{RIV}_{F, \Pi}$, and definition of a plaintext-leaking oracle Λ that will be used in our security analysis.

$\tilde{\mathcal{E}} : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{M} \rightarrow \mathcal{C} \times \mathcal{T}$ and decryption algorithm $\tilde{\mathcal{D}} : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{C} \times \mathcal{T} \rightarrow \mathcal{M} \cup \{\perp\}$, as given in Figure 1.

Definition 3 ($\widehat{\text{RIV}}$). We define the SAE scheme $\widehat{\text{RIV}}_{F, \Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}}, \Lambda)$ with an additional deterministic leakage algorithm $\Lambda : \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{C} \times \mathcal{T} \rightarrow \mathcal{M} \times \{\top\}$, as given in Figure 1.

Feistel Structure and Encode-then-Encipher (EtE). RIV can be seen as an application of the EtE [9] approach by Bellare et al. EtE can generically be used for constructing a robust AE scheme from a tweakable cipher, assuming its enciphering resists chosen-plaintext and chosen-ciphertext attacks [24]. The RIV cipher, however, is essentially an unbalanced three-round Feistel-network.³ It is well-known that such ciphers are secure against chosen-plaintext, but vulnerable to chosen-ciphertext attacks [35] (see also [2,36,42]). RIV is robust *in spite* of its weak enciphering scheme, because its encoding operation has been chosen to specifically cover this weakness.

4 Security Notions

Adversaries and Advantages. An adversary \mathbf{A} is an efficient Turing machine that interacts with a given set of oracles that appear as black boxes to \mathbf{A} . We

³ If the used encryption scheme $\Pi = (\mathcal{E}, \mathcal{D})$ is nonce-*keystream*-based, the RIV cipher is a three-round Feistel network.

use the notation \mathbb{A} for the class of all computationally bounded adversaries and $\mathbf{A}^{\mathcal{O}}$ for the output of \mathbf{A} after interacting with some oracle \mathcal{O} . We write $\Delta_{\mathbf{A}}(\mathcal{O}^L; \mathcal{O}^R) := \sup_{\mathbf{A} \in \mathbb{A}} |\Pr[\mathbf{A}^{\mathcal{O}^L} \Rightarrow 1] - \Pr[\mathbf{A}^{\mathcal{O}^R} \Rightarrow 1]|$ for the advantage of \mathbf{A} to distinguish between oracles \mathcal{O}^L and \mathcal{O}^R . All probabilities are defined over the random coins of the oracles and those of the adversary, if any. We write $\mathbf{Adv}_F^X(q, \ell, t) = \max_{\mathbf{A} \in \mathbb{A}} \{\mathbf{Adv}_F^X(\mathbf{A})\}$ to refer to the maximal advantage over all X -adversaries \mathbf{A} on a given function F that run in time at most t and pose at most q queries consisting of at most ℓ blocks in total to the available oracles. If \mathbf{A} shall distinguish between two sets of oracles $(\mathcal{O}_1^L, \dots, \mathcal{O}_k^L)$ and $(\mathcal{O}_1^R, \dots, \mathcal{O}_k^R)$, we refer to the i -th oracle that \mathbf{A} interacts with by $\mathcal{O}_i \in \{\mathcal{O}_i^L, \mathcal{O}_i^R\}$. By $\mathcal{O}_i \hookrightarrow \mathcal{O}_j$, we denote that \mathbf{A} first queries \mathcal{O}_i and later \mathcal{O}_j with the output of \mathcal{O}_i . Wlog., we assume that \mathbf{A} never asks queries to which it already knows the answer. In the case when \mathbf{A} has access to multiple oracles $\mathcal{O}_1, \dots, \mathcal{O}_k$, we denote by q_i the number of queries and by ℓ_i the maximal number of blocks that \mathbf{A} poses at most to oracle \mathcal{O}_i , $1 \leq i \leq k$.

If \mathcal{O}_i and \mathcal{O}_j represent a family of algorithms indexed by inputs, the indices must match, e.g., when $\tilde{\mathcal{E}}_K^{N,H}(M)$ and $\tilde{\mathcal{D}}_K^{N,H}(C)$ represent encryption and decryption algorithms with a fixed key K and indexed by N and H , then $\tilde{\mathcal{E}}_K \hookrightarrow \tilde{\mathcal{D}}_K$ says that \mathbf{A} first queries $\tilde{\mathcal{E}}_K^{N,H}(M)$ and later $\tilde{\mathcal{D}}_K^{N,H}(C)$.

We define \perp , when in place of an oracle, to always return the invalid symbol \perp . We denote by $\mathcal{S}^{\mathcal{O}}$ an oracle that, given an input X , computes $Y \leftarrow \mathcal{O}(X)$, chooses uniformly at random a value Y' from the space of all possible outputs with $|Y'| = |Y|$, and returns Y' . We assume that $\mathcal{S}^{\mathcal{O}}$ performs lazy sampling, i.e., $\mathcal{S}^{\mathcal{O}}(X)$ returns the same value when queried with the same input X . We often omit the key for brevity, e.g., $\mathcal{S}^{\tilde{\mathcal{E}}}(X)$ will be short for $\mathcal{S}^{\tilde{\mathcal{E}}_K}(X)$.

4.1 Security Definitions for Encryption Schemes

Definition 4 (PRF Advantage). Let $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ be a function with non-empty key space \mathcal{K} , and \mathbf{A} a computationally bounded adversary with access to an oracle, where $K \leftarrow \mathcal{K}$ and $\rho \leftarrow \text{Func}(\mathcal{X}, \mathcal{Y})$. Then, the PRF advantage of \mathbf{A} on F is defined as $\mathbf{Adv}_F^{\text{PRF}}(\mathbf{A}) := \Delta_{\mathbf{A}}(F_K; \rho)$.

Definition 5 (PRP Advantage). Let $n, k \geq 1$ be fixed. Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher and \mathbf{A} a computationally bounded adversary with access to an oracle. Further, let $K \leftarrow \{0, 1\}^k$ and $\pi \leftarrow \text{Perm}(\{0, 1\}^n)$. Then, the PRP advantage of \mathbf{A} on E is defined as $\mathbf{Adv}_E^{\text{PRP}}(\mathbf{A}) := \Delta_{\mathbf{A}}(E_K; \pi)$.

Stinson [48] showed that one can construct an $(\epsilon_1 + \epsilon_2)$ -AU family of hash functions from the consecutive application of an ϵ_1 -AU and an ϵ_2 -AU family of hash functions. From that we can derive the following theorem.

Theorem 2. Let $\mathcal{X}, \mathcal{Y}, \mathcal{Z} \subseteq \{0, 1\}^*$ and let \mathcal{K} be a non-empty set. Further, let $\mathcal{H} = \{H : \mathcal{X} \rightarrow \mathcal{Y}\}$ be a family of ϵ -AU hash functions and let $G : \mathcal{K} \times \mathcal{Y} \rightarrow \mathcal{Z}$ be a function. Then, we can define $F_K(X) := G_K(H(X))$, with independent $K \leftarrow \mathcal{K}$ and $H \leftarrow \mathcal{H}$. Let \mathbf{A} be a PRF adversary on F that asks at most q queries of

at most ℓ blocks in total, and runs in time at most t . Then, there exists a PRF adversary \mathbf{A}_1 on G that asks at most q queries and runs in time $O(t)$ such that

$$\mathbf{Adv}_F^{\text{PRF}}(\mathbf{A}) \leq \mathbf{Adv}_G^{\text{PRF}}(\mathbf{A}_1) + \epsilon \cdot q^2/2.$$

Theorem 2 follows from the fact, that the PRF advantage of F is upper bounded by the maximal PRF advantage on G plus the maximal probability of output collisions of the form $H(X) = H(X')$ over q queries.

Definition 6 (NE Advantage [41]). Let $\Pi = (\mathcal{E}, \mathcal{D})$ be a nonce-based encryption scheme and $K \leftarrow \mathcal{K}$. Let \mathbf{A} be a nonce-respecting adversary with access to an oracle. Then, the NE advantage of \mathbf{A} on Π is defined as $\mathbf{Adv}_{\Pi}^{\text{NE}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\mathcal{E}_K; \mathcal{E}^{\mathcal{E}})$.

We adapt the definition of indistinguishability from random bits from [23] for nonce-based encryption schemes. Note that we strengthen it to adversaries that do not repeat nonces over *all encryption and decryption queries*.

Definition 7 (SRND Advantage). Let $\Pi = (\mathcal{E}, \mathcal{D})$ a nonce-based encryption scheme and $K \leftarrow \mathcal{K}$. Let \mathbf{A} be a nonce-respecting adversary with access to two oracles \mathcal{O}_1 and \mathcal{O}_2 , s.t. \mathbf{A} never asks for $\mathcal{O}_1 \hookrightarrow \mathcal{O}_2$ and never repeats a nonce over all its encryption and decryption queries. Then, we define the SRND advantage of \mathbf{A} on Π as $\mathbf{Adv}_{\Pi}^{\text{SRND}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\mathcal{E}_K, \mathcal{D}_K; \mathcal{E}^{\mathcal{E}}, \mathcal{D}^{\mathcal{D}})$.

4.2 Security Definitions for Nonce-Based AE Schemes

For this subsection, let $\tilde{\Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}})$ be a nonce-based AE scheme, $K \leftarrow \mathcal{K}$, and \mathbf{A} be a computationally bounded adversary on $\tilde{\Pi}$.

Definition 8 (IND-CPA Advantage). Let \mathbf{A} have access to an encryption oracle. Then, the IND-CPA advantage of \mathbf{A} with respect to $\tilde{\Pi}$ is defined as $\mathbf{Adv}_{\tilde{\Pi}}^{\text{IND-CPA}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\tilde{\mathcal{E}}_K; \mathcal{E}^{\tilde{\mathcal{E}}})$.

Definition 9 (INT-CTXT Advantage). Let \mathbf{A} have access to two oracles \mathcal{O}_1 and \mathcal{O}_2 such that \mathbf{A} never queries $\mathcal{O}_1 \hookrightarrow \mathcal{O}_2$. Then, the INT-CTXT advantage of \mathbf{A} on $\tilde{\Pi}$ is defined as $\mathbf{Adv}_{\tilde{\Pi}}^{\text{INT-CTXT}}(\mathbf{A}) := \Pr[\mathbf{A}^{\tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K} \text{ forges}]$, where “forges” means that $\tilde{\mathcal{D}}_K$ returns anything other than \perp for a query of \mathbf{A} .

Definition 10 (NAE Advantage [41]). Let \mathbf{A} have access to two oracles \mathcal{O}_1 and \mathcal{O}_2 such that \mathbf{A} never queries $\mathcal{O}_1 \hookrightarrow \mathcal{O}_2$. Then, the NAE advantage of \mathbf{A} on $\tilde{\Pi}$ is defined as $\mathbf{Adv}_{\tilde{\Pi}}^{\text{NAE}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K; \mathcal{E}^{\tilde{\mathcal{E}}}, \perp)$.

Bellare and Namprempre showed for probabilistic AE that chosen-ciphertext security results from IND-CPA and INT-CTXT security [8]. Fleischmann et al. proved in [19] a generalized theorem for nonce-based AE.

Theorem 3 (Theorem 1 in [19]). Let \mathbf{A} be a computationally bounded NAE adversary on $\tilde{\Pi}$ with access to two oracles \mathcal{O}_1 and \mathcal{O}_2 such that \mathbf{A} never queries

$\mathcal{O}_1 \leftrightarrow \mathcal{O}_2$; \mathbf{A} makes at most q queries of total length of at most ℓ blocks and runs in time at most t . Then, there exist an IND-CPA adversary \mathbf{A}_1 on $\tilde{\Pi}$ and an INT-CTXT adversary \mathbf{A}_2 on $\tilde{\Pi}$, both making at most q queries of at most ℓ blocks and running in time $O(t)$ each, such that

$$\text{Adv}_{\tilde{\Pi}}^{\text{NAE}}(\mathbf{A}) \leq \text{Adv}_{\tilde{\Pi}}^{\text{IND-CPA}}(\mathbf{A}_1) + \text{Adv}_{\tilde{\Pi}}^{\text{INT-CTXT}}(\mathbf{A}_2).$$

4.3 Security Definitions for Subtle AE Schemes

Subtle AE (SAE) defines a compound security notion that provides guarantees for privacy and authenticity under the existence of a leakage oracle. It comprises the notions IND-CPA, INT-CTXT, and an additional notion ERR-CCA.

For this subsection, let $\tilde{\Pi} = (\tilde{\mathcal{E}}, \tilde{\mathcal{D}}, \Lambda)$ be an SAE scheme, $K, K' \leftarrow \mathcal{K} \times \mathcal{K}$ independent keys, and \mathbf{A} a deterministic adversary with access to three oracles $\mathcal{O}_1, \mathcal{O}_2$, and \mathcal{O}_3 such that \mathbf{A} neither queries $\mathcal{O}_1 \leftrightarrow \mathcal{O}_2$ nor $\mathcal{O}_1 \leftrightarrow \mathcal{O}_3$.

Definition 11 (ERR-CCA Advantage). *The ERR-CCA advantage of \mathbf{A} on $\tilde{\Pi}$ is defined as $\text{Adv}_{\tilde{\Pi}}^{\text{ERR-CCA}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K, \Lambda_K; \tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K, \Lambda_{K'})$.*

Definition 12 (SAE Advantage). *The SAE advantage of \mathbf{A} on $\tilde{\Pi}$ is defined as $\text{Adv}_{\tilde{\Pi}}^{\text{SAE}}(\mathbf{A}) := \Delta_{\mathbf{A}}(\tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K, \Lambda_K; \$^{\tilde{\mathcal{E}}}, \perp, \Lambda_{K'})$.*

In the full version of [7], Barwell et al. prove a statement equivalent to Theorem 4. We apply Theorem 3 to decompose their AE security advantage term into the separate advantages for IND-CPA and INT-CTXT.

Theorem 4. *Let \mathbf{A} run in time at most t and ask at most q queries of at most ℓ blocks to its respective oracles. Then, there exist computationally bounded IND-CPA, INT-CTXT, and ERR-CCA adversaries $\mathbf{A}_1, \mathbf{A}_2$, and \mathbf{A}_3 , respectively, on $\tilde{\Pi}$ such that*

$$\text{Adv}_{\tilde{\Pi}}^{\text{SAE}}(\mathbf{A}) \leq \text{Adv}_{\tilde{\Pi}}^{\text{IND-CPA}}(\mathbf{A}_1) + \text{Adv}_{\tilde{\Pi}}^{\text{INT-CTXT}}(\mathbf{A}_2) + \text{Adv}_{\tilde{\Pi}}^{\text{ERR-CCA}}(\mathbf{A}_3),$$

where $\mathbf{A}_1, \mathbf{A}_2$, and \mathbf{A}_3 each make at most q queries of at most ℓ blocks and run in time $O(t)$ each.

Since [4] omitted a compound notion for their security under release of unverified plaintexts, Barwell et al. defined RUPAE as $\Delta_{\mathbf{A}}(\tilde{\mathcal{E}}_K, \tilde{\mathcal{D}}_K, \mathcal{V}_K; \$^{\tilde{\mathcal{E}}}, \tilde{\mathcal{D}}_{K'}, \perp)$ [7, Theorem 3, Corollary 2]. They showed that the maximal SAE advantage on an AE scheme $\tilde{\Pi}$ is, with a reduction term, also equivalent to the maximal RUPAE advantage. Moreover, they showed that – again with a reduction term – it is also equivalent to the maximal robust-AE advantage on $\tilde{\Pi}$ with fixed stretch τ .

5 Security Results for Generic RIV

This section summarizes our security results. For the remainder of this section, let $d, n, \tau \geq 1$ be integers, $\mathcal{K}_1, \mathcal{K}_2$ be non-empty key spaces, and $K_1, K_2 \leftarrow \mathcal{K}_1 \times \mathcal{K}_2$ be independent keys, $F : \mathcal{K}_1 \times \{0, 1\}^d \times \mathcal{N} \times \mathcal{H} \times \mathcal{M} \rightarrow \{0, 1\}^n$, and $\Pi = (\mathcal{E}, \mathcal{D})$ be a nonce-based encryption scheme with associated key space \mathcal{K}_2 .

Theorem 5. *Let \mathbf{A} be a computationally bounded SAE adversary on $\widehat{\text{RIV}}_{F,\Pi}$ which asks at most q queries of at most ℓ blocks in total and runs in time at most t . Then, there exists a computationally bounded PRF adversary \mathbf{A}_1 on F that asks at most $2q$ queries of at most $2(d + n\ell)$ bits and runs in time $O(t)$, and a computationally bounded SRND adversary \mathbf{A}_2 on Π that asks at most q queries of at most ℓ blocks in total and runs in time $O(t)$ such that*

$$\text{Adv}_{\widehat{\text{RIV}}_{F,\Pi}}^{\text{SAE}}(\mathbf{A}) \leq \frac{8q^2 + 3q}{2^n} + 4 \cdot \left(\text{Adv}_F^{\text{PRF}}(\mathbf{A}_1) + \text{Adv}_{\Pi}^{\text{SRND}}(\mathbf{A}_2) \right).$$

Due to space limitations, the proof can be found in the full version of this paper⁴. We can derive the following corollary for the NAE advantage on $\text{RIV}_{F,\Pi}$ in the absence of a plaintext-leaking oracle.

Corollary 1. *Let \mathbf{A} be a computationally bounded NAE adversary on $\text{RIV}_{F,\Pi}$ which asks at most q queries of at most ℓ blocks in total and runs in time at most t . Then, there exist a computationally bounded PRF adversary \mathbf{A}_1 on F that asks at most $2q$ queries of at most $2(d + n\ell)$ bits and runs in time $O(t)$, and a computationally bounded SRND adversary \mathbf{A}_2 on Π that asks at most q queries of at most ℓ blocks in total and runs in time $O(t)$, such that*

$$\text{Adv}_{\text{RIV}_{F,\Pi}}^{\text{NAE}}(\mathbf{A}) \leq \frac{2q^2 + q}{2^n} + 2 \cdot \left(\text{Adv}_F^{\text{PRF}}(\mathbf{A}_1) + \text{Adv}_{\Pi}^{\text{SRND}}(\mathbf{A}_2) \right).$$

The proof can be found can be found in the full version of this paper.

Proof Ideas. The intuition of our proofs is the following: in encryption direction, for every fresh tuple of nonce, header, and message, F will produce a fresh $IV \leftarrow F_K^1(N, H, M)$ that has not occurred before with overwhelming probability. Since Π is SRND-secure, \mathcal{E} will produce a randomly chosen ciphertext. The second invocation of F with a fresh ciphertext then produces a random tag. To determine the privacy advantage of the scheme, we have to bound only the PRF-advantage on F , the SRND-security of \mathcal{E} , and the probabilities of random collisions of IV s from the birthday paradox.

In decryption direction, whenever the nonce, header, or ciphertext changes, $IV \leftarrow F_K^2(N, H, C)$ will be a random value up to the birthday bound. Since Π is SRND-secure, a fresh IV (regarded over all encryption *and* decryption queries) will produce a fresh pseudorandom plaintext. Thus, even when the adversary learns the decrypted (invalid) message, M will provide it with no information about other plaintexts as long as the IV does not repeat. When an adversary changes N , H , or C and manages to cancel the difference by a fresh tag, the second call to $F_K^1(N, H, M)$ will yield a random IV' that differs from IV with probability close to $1/2^n$. Thus, a similar argumentation as for the encryption also applies to the inverse direction. Finally, the domain separation from the first parameter to F protects against choices of $(N, H, M) = (N, H, C)$.

⁴ The full version of this paper will soon appear on ePrint.

6 Instantiation

Pseudo-Dot-Product Hashing. Let $n, m \geq 1$ with even m and let $\mathcal{X} = \bigcup_{i=1}^{m/2} \{0, 1\}^{2in}$. Given a set of m pair-wise independent key words $K = (K_1, \dots, K_m)$ and an m -word input $M = (M_1, \dots, M_m)$, with $M_i, K_i \in \{0, 1\}^n$, $1 \leq i \leq m$, a *pseudo-dot-product* (PDP) family of hash functions $\mathcal{H} = \{H : \mathcal{X} \times \mathcal{X} \rightarrow \{0, 1\}^{2n}\}$ is defined as

$$H_K(M) := \sum_{i=1}^{m/2} (M_{2i-1} + K_{2i-1}) \cdot (M_{2i} + K_{2i}).$$

Bernstein [11] credits it to Winograd [51] and classifies it as $(m, \lceil m/2 \rceil)$ -design, i.e., it requires m independent key words and $\lceil m/2 \rceil$ multiplications to process m message words. If modular additions and multiplications are performed within the rings \mathbb{Z}_{2^n} and $\mathbb{Z}_{2^{2n}}$, the construction is known as NH, to be $1/2^n$ -AU, and is used in variants in UMAC [13], VMAC [18,32], and HS1 [33]. All these constructions employ a multi-stage hashing process: the input is first compressed with NH, before the results are used as inputs in a usual polynomial hash (and optionally further processed by an inner-product hash). To obtain a slightly higher security margin and efficiency, we consider a recently proposed variant, called CLHASH.

6.1 CLHASH

CLHASH [34] is a family of multi-stage hash functions that produces 64-bit hashes and employs a PDP family of hash functions CLNH, which resembles NH, but replaces modular additions and multiplications with XORs and carry-less multiplications in $\mathbb{GF}(2^{64})/p(x)$ with the irreducible polynomial $p(x) = x^{64} + x^4 + x^3 + x + 1$. Therefore, CLNH can exploit the `vpclmulqdq` instruction for 64-bit carry-less multiplication which was originally introduced for boosting the performance of GCM [21].

CLHASH[m] splits a given message M into $(64m)$ -bit blocks (M_1, \dots, M_s) , and pads the final block with zeroes such that its length becomes a multiple of 128 bits. Each block M_j is compressed with CLNH to a 128-bit value A_j . If the message consists of only a single block, the message length $|M|$ is multiplied with an independent key $K_L \in \{0, 1\}^{64}$ and XORed to the result; the result is reduced to a 64-bit value modulo $p(x) = x^{64} + x^4 + x^3 + x + 1$ and returned.

For longer messages, the values A_j are processed by a polynomial hash with an independent key $K_P \in \{0, 1\}^{128}$ and reduced modulo $q(x) = x^{127} + x + 1$. For efficiency, the two most significant bits of K_P are fixed to zero, and a lazy reduction modulo $x^{128} + x^2 + x$ is used instead without affecting security.

The 128-bit result of the polynomial hash is then reduced to a 64-bit value by another application of CLNH with two further independent key words $K_{A_1}, K_{A_2} \in \{0, 1\}^{64}$. The result H is finally XORed with the hashed length to account for inputs of variable lengths, and is reduced to a 64-bit value.

Algorithm 1 Definition of $\text{CLHASH}^T[m, t]$ with a hash length of $64t$ bits, a block length of $m/8$ bytes, and t Toeplitz iterations.

<pre> 101: function CLHASH^T[m, t]_K(M) 102: (K_N, K_P, K_A, K_L) ← KEYGEN(K) 103: s ← max(⌈64 · M /m⌉, 1) 104: (M₁, ..., M_s) ←^{64m} M 105: M_s ← PAD₁₂₈(M_s) 106: for i ← 1 to t do 107: for j ← 1 to s do 108: K_j ← K_{N(2i-1)..m+2(i-1)} 109: A_j ← CLNH[m]_{K_j}(M_j) 110: if s = 1 then 111: H₁ ← A₁ 112: else 113: K_{P_i} ← K_{P_i} mod 2¹²⁶ 114: O_i ← POLY_{K_{P_i}}(A₁, ..., A_s) 115: H_i ← CLNH[2]_{K_{A_i}}(O_i) 116: H_i ← HASHLEN_{K_{L_i}}(H_i, M) 117: return (H₁ ... H_t) 201: function CLNH[m]_{K_j}(M_j) 202: return ⊕_{i=1}^m (M_{j_{2i-1}} ⊕ K_{j_{2i-1}}) 203: · (M_{j_{2i}} ⊕ K_{j_{2i}}) </pre>	<pre> 301: function KEYGEN(K) 302: κ ← 64(m + 2t - 2) 303: K_N ← K[1..κ] 304: K_P ← K[(κ + 1)..(κ + 128t)] 305: κ ← κ + 128t 306: K_A ← K[(κ + 1)..(κ + 128t)] 307: κ ← κ + 128t 308: K_L ← K[(κ + 1)..(κ + 64t)] 309: return (K_N, K_P, K_A, K_L) 401: function POLY_{K_P}(A₁, ..., A_s) 402: return ⊕_{i=1}^s A_i · K_P^{s-i} 403: mod (2¹²⁸ + 4 + 2) 501: function HASHLEN_{K_L}(H_i, M) 502: return (H_i ⊕ (K_L · M)) 503: mod (2⁶⁴ + 27) 601: function PAD_n(X) 602: if (X mod n = 0) then 603: return X 604: return X 0^{n- X mod n} </pre>
--	---

In [34], the authors show that CLHASH is XOR-universal for messages of up to $b = 8m$ bytes, and ϵ -AXU for messages of up to N bytes.

Theorem 6 (Lemma 9 in [34]). *Let $N \geq 1$ denote the maximal message length in bytes, $m \geq 2$ be even, and $b = 8m$ the key size of CLNH. Then, CLHASH as defined above is ϵ -AXU with*

$$\epsilon \leq \epsilon_{\text{CLNH}[m]} + \epsilon_{\text{POLY}} + \epsilon_{\text{CLNH}[2]} \leq \frac{1}{2^{64}} + \frac{N/b - 1}{2^{126}} + \frac{1}{2^{64}},$$

where the terms stem from the facts that CLNH[m] is an $\epsilon_{\text{CLNH}[m]}$ -AU, and the polynomial hash an ϵ_{POLY} -AXU family of hash functions.

The recommended values $N \leq 2^{64}$ and $b = 1024$ yield $\epsilon \leq 2.004/2^{-64}$. The construction requires $b + 40$ bytes of key material: b bytes for CLNH, a 16-byte value K_P for the polynomial hash, two eight-byte values $K_A[1], K_A[2]$ for the final call to CLNH, and an eight-byte value K_L for hashing the input length.

Toeplitz Extension. To obtain a hash function with 128-bit security, one can process the same message twice under independent keys and concatenate the results. Doubling the key lengths of K_P , K_A , and K_L increases their keys to 80 bytes. Since doubling the key length for CLNH would absurdly increase the key material, we use the Toeplitz extension [31,37] instead. Let $K_{i..j}$ be short for

K_i, \dots, K_j , $1 \leq i \leq j$. Given an ϵ -AU family of hash functions $H : \{0, 1\}^{mn} \times \{0, 1\}^{mn} \rightarrow \{0, 1\}^n$ which compresses an m -word input with an m -word key, one can derive a hash function $H^t : \{0, 1\}^{(m+2t-2)n} \times \{0, 1\}^{mn} \rightarrow \{0, 1\}^{tm}$ by

$$H_{K_{1..(m+2t-2)}}^t(M) := H_{K_{1..m}}(M) \parallel H_{K_{3..(m+2)}}(M) \parallel \dots \parallel H_{K_{(2t-1)..(m+2t-2)}}(M).$$

So, the i -th call to H employs the key shifted by $2i-2$ words. In total, the key size increases slightly from m to $m+2(t-1)$ words. We refer to the Toeplitz version of CLNH by $\text{CLNH}^T[m, t]$, and to that of CLHASH[m] by $\text{CLHASH}^T[m, t]$. Algorithm 1 provides a specification. In total, $\text{CLHASH}^T[m, t]$ requires $(8m + 56t - 16)$ bytes of key material, which corresponds to $(8m + 96)$ bytes for $t = 2$.

Definition 13 (Toeplitz CLHASH). *Let $n = 64$, $t \geq 1$, $m \geq 2$ be even. Let $\mathcal{X} = \bigcup_{i=1}^{m/2} \{0, 1\}^{2in}$. Let further $\mathcal{K}_N = \{0, 1\}^{64m+128(t-1)}$, $\mathcal{K}_P = \{0, 1\}^{128t}$, $\mathcal{K}_A = \{0, 1\}^{128t}$, $\mathcal{K}_L = \{0, 1\}^{64t}$, and $\mathcal{K} = \mathcal{K}_N \times \mathcal{K}_P \times \mathcal{K}_A \times \mathcal{K}_L$. The family of keyed hash functions $\text{CLHASH}^T[m, t] : \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^{64t}$ is defined in Algorithm 1.*

Theorem 7. *For any fixed $n, t \geq 1$, and even $m \geq 2$, $\text{CLNH}^T[m, t]$ is 2^{-nt} -AU on equal-length strings.*

The proof of Theorem 7 can be found in the full version of this paper.

Theorem 8. *Let $N \leq 2^{64}$ be the maximal message length in bytes, $t \geq 1$, $m \geq 2$ be even, and $b = 8m$ the key length for CLNH in bytes. Then, $\text{CLHASH}^T[m, t]$ is an ϵ^t -AXU family of hash functions with*

$$\epsilon \leq \epsilon_{\text{CLNH}[m]} + \epsilon_{\text{POLY}} + \epsilon_{\text{CLNH}[2]} \leq \frac{1}{2^{64}} + \frac{N/b - 1}{2^{126}} + \frac{1}{2^{64}} \leq \frac{3}{2^{64}}.$$

The proof of Theorem 8 follows from Theorem 7 and the fact that the keys for the individual iterations of polynomial, inner-product, and length hashing steps are chosen uniformly from their respective spaces and pairwise independently for each iteration. We can derive that $\text{CLHASH}^T[m, 2]$ is ϵ -AXU for $\epsilon \leq 9/2^{128}$ when $m \geq 2$.

6.2 Constructing a PRF

Let $n, d \geq 1$, and $\mathcal{N}, \mathcal{H}, \mathcal{M}$ be as in Section 3. For brevity, we define $\mathcal{Y} := \{0, 1\}^d \times \mathcal{N} \times \mathcal{H} \times \mathcal{M}$. Let $\text{ENCODE} : \mathcal{Y} \rightarrow \{0, 1\}^*$ define an injective encoding function. Then, we can construct a PRF from the composition of ENCODE , a family of ϵ -AU hash functions $\mathcal{H}' = \{H' | H' : \{0, 1\}^* \rightarrow \{0, 1\}^n\}$, and a block cipher $E : \mathcal{K}_2 \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, with independent keys $K_1 \in \mathcal{K}_1$ determining the hash function, and $K_2 \in \mathcal{K}_2$ for the cipher. We call the construction $\text{EHE}[\text{ENCODE}, \mathcal{H}', E] : \mathcal{Y} \rightarrow \{0, 1\}^n$ (for *Encode-Hash-Encrypt*) and define it as

$$\text{EHE}[\text{ENCODE}, \mathcal{H}', E]_{K_1, K_2}(D, N, H, M) := E_{K_2}(\mathcal{H}'_{K_1}(\text{ENCODE}(D, N, H, M))).$$

We write $\text{EHE}[\mathcal{H}', E]$ or even EHE as short forms of $\text{EHE}[\text{ENCODE}, \mathcal{H}', E]$ when the components are clear from the context. The injective encoding excludes collisions between distinct inputs. From Theorem 2, and applying the PRF/PRP switching lemma, we can derive the following theorem.

Algorithm 2 Encryption of nonce-based XOR-CTR, instantiated with a block cipher $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$, with $n, k \geq 1$.

```

1: function XOR-CTR[ $E$ ]. $\mathcal{E}_K^N(M)$ 
2:    $IV \leftarrow E_K(N)$ 
3:    $m \leftarrow \lceil |M|/n \rceil$ 
4:    $\kappa \leftarrow E_K(IV \oplus \langle 0 \rangle) \parallel \dots \parallel E_K(IV \oplus \langle m-1 \rangle)$ 
5:   return  $C \leftarrow M \oplus \kappa[\text{first } |M| \text{ bits}]$ 

```

Theorem 9. Let $\pi \leftarrow \text{Perm}(\{0, 1\}^n)$. Further, let $\text{EHE}[\text{ENCODE}, \mathcal{H}', \pi]$, \mathcal{H}' , and ENCODE be defined as above. Let \mathbf{A} be a computationally bounded adversary that asks at most q queries of at most ℓ blocks and runs in time at most t . Then

$$\text{Adv}_{\text{EHE}[\text{ENCODE}, \mathcal{H}', \pi]}^{\text{PRF}}(\mathbf{A}) \leq \binom{q}{2} \cdot \left(\frac{1}{2^n} + \epsilon \right).$$

6.3 Encryption

When starting counter-mode encryption from a random value and incrementing by modular addition, one has to either consider potential carry bits or to reduce the security by fixing a maximal message length. Wang et al. [50] proposed to replace modular addition by XOR, which avoids the need for concerning carry bits. Let $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher. We define $\text{XOR-CTR}[E] = (\mathcal{E}, \mathcal{D})$ as the nonce-based encryption scheme with encryption algorithm $\text{XOR-CTR}[E].\mathcal{E} : \{0, 1\}^k \times \mathcal{N} \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ and associated non-empty nonce-space \mathcal{N} , as defined in Algorithm 2.

We denote by $\text{XOR-CTR}[\pi, \pi']$ a version of XOR-CTR with two independent n -bit permutations π and π' , where π is used for encrypting the nonce and π' for producing the keystream. Then, $\text{XOR-CTR}[\pi, \pi']$ is almost identical to the $\text{CTR2}[\pi, \pi']$ construction in [45], with the difference that the former replaces the addition of IV and counter modulo 2^n by XOR. Since this change does not affect the probability of block-cipher inputs to repeat, the NE advantage of XOR-CTR is given by Theorem 10, which adapts Theorem 3 in [45].

Theorem 10. Let $\pi, \pi' \leftarrow \text{Perm}(\{0, 1\}^n) \times \text{Perm}(\{0, 1\}^n)$ be independent permutations and \mathbf{A} be a nonce-respecting NE adversary, which runs in time at most t and poses at most q queries to its oracles with at most ℓ blocks. Then

$$\text{Adv}_{\text{XOR-CTR}[\pi, \pi']}^{\text{NE}}(\mathbf{A}) \leq \frac{\ell^2}{2^n}.$$

From the fact that encryption and decryption of $\text{XOR-CTR}[\pi, \pi']$ are identical operations, we can derive the following theorem.

Theorem 11. There exists a reduction of a nonce-respecting SRND adversary \mathbf{A} with access to two oracles on $\text{XOR-CTR}[\pi, \pi']$ to a nonce-respecting NE adversary \mathbf{A}' on $\text{XOR-CTR}[\pi, \pi']$ such that

$$\text{Adv}_{\text{XOR-CTR}[\pi, \pi']}^{\text{SRND}}(\mathbf{A}) \leq \text{Adv}_{\text{XOR-CTR}[\pi, \pi']}^{\text{NE}}(\mathbf{A}'),$$

Algorithm 3 Definition of our instantiation $\text{RIV}_{F,\Pi}$. Message and header lengths are restricted to multiple of eight bits, and nonces/IVs/tags are 128 bits: $n = \tau = 128$, and $d = 4$. Here, we leave the key size of $\text{CLHASH}^\top[m, 2]$, m , as a parameter to study its impact on performance later.

101: function $\tilde{\mathcal{E}}_{SK}(N, H, M)$ 102: $(K_1, K_2) \leftarrow \text{KEYGEN}(SK)$ 103: $IV \leftarrow \text{EHE}_{K_1, K_2}^1(N, H, M)$ 104: $C \leftarrow \text{XOR-CTR}[E].\mathcal{E}_{K_2}(IV, M)$ 105: $T \leftarrow \text{EHE}_{K_1, K_2}^2(N, H, C) \oplus IV$ 106: return (C, T) 201: function $\text{KEYGEN}(SK)$ 202: $K_2 \leftarrow E_{SK}(\langle 0 \rangle)$ 203: $\kappa \leftarrow (8m + 96)/16$ 204: $K_1 \leftarrow E_{SK}(\langle 1 \rangle) \parallel \dots \parallel E_{SK}(\langle \kappa \rangle)$ 205: return (K_1, K_2) 301: function $\text{EHE}_{K_1, K_2}^D(N, H, X)$ 302: $Y \leftarrow \text{ENCODE}(D, N, H, X)$ 303: return $E_{K_2}(\mathcal{H}'_{K_1}(Y))$ 401: function $\text{PAD}_n(X)$ 402: if $(X \bmod n = 0)$ then 403: return X 404: return $X \parallel 0^{n- X \bmod n}$	501: function $\tilde{\mathcal{D}}_{SK}(N, H, C, T)$ 502: $(K_1, K_2) \leftarrow \text{KEYGEN}(SK)$ 503: $IV \leftarrow \text{EHE}_{K_1, K_2}^2(N, H, C) \oplus T$ 504: $M \leftarrow \text{XOR-CTR}[E].\mathcal{D}_{K_2}(IV, C)$ 505: $IV' \leftarrow \text{EHE}_{K_1, K_2}^1(N, H, M)$ 506: if $(IV = IV')$ then 507: return M 508: return \perp 601: function $\text{ENCODE}(D, N, H, X)$ 602: $\bar{H} \leftarrow \text{PAD}_{128}(H)$ 603: $\bar{X} \leftarrow \text{PAD}_{128}(X)$ 604: $\bar{L} \leftarrow \langle D \rangle_d \parallel \langle H /8 \rangle_{60} \parallel \langle X /8 \rangle_{64}$ 605: return $(\bar{H} \parallel N \parallel \bar{X} \parallel \bar{L})$ 701: function $\mathcal{H}'_{K_1}(X)$ 702: return $\text{CLHASH}^\top[m, 2]_{K_1}(X)$ 801: function $E_{K_2}(X)$ 802: return $\text{AES-128}_{K_2}(X)$
--	--

where both \mathbf{A} and \mathbf{A}' ask at most q queries of at most ℓ blocks to their available oracle(s) and run in time $O(t)$.

6.4 Instantiation of RIV

We instantiate $\text{RIV}_{F,\Pi}$ with $\text{EHE}[\text{ENCODE}, \mathcal{H}', E]$ for F , with $\text{CLHASH}^\top[m, 2]$ as family of universal hash functions \mathcal{H}' , and $\text{XOR-CTR}[E]$ for Π , with the AES-128 as E . Algorithm 3 provides a specification. Our instantiation $\text{RIV}_{F,\Pi}$ expects a 128-bit user-supplied secret key SK , from which the remaining key material is derived by calling $E_{SK}(\cdot)$ iteratively in counter mode. The secret key is not used further. RIV uses $n = \tau = 128$, i.e., n -bit tags, and n -bit IVs for the counter mode. Moreover, the nonce space is fixed to 128 bits: $\mathcal{N} = \{0, 1\}^n$. For F , it employs a four-bit domain separation, i.e., $d = 4$, and an injective encoding function $\text{ENCODE} : \{0, 1\}^d \times \mathcal{N} \times \mathcal{H} \times \mathcal{M} \rightarrow \{0, 1\}^*$, as defined in Algorithm 3. Header and message lengths are restricted to multiple of eight bits. The maximal number of header and message bytes to be encrypted under the same key are 2^{60} bytes each. So, the maximal number of bytes for RIV is less than 2^{62} bytes. We recommend that at most 2^{50} bytes be encrypted under the same key.

Using a Single Key for the Block Cipher. There are four uses of the block cipher E in RIV: in the first invocation of EHE, for encrypting the IV , for

generating the keystream in XOR-CTR[E], and in the second invocation of EHE. If four more calls to the AES key schedule would be tolerable, one could use four independent keys. Alternatively, we use a single key for the uses of E , and have to consider the security impact in the following theorem. Its proof can be found in the full version of this paper.

Theorem 12. *Let $\text{RIV}_{F,\Pi}$ be defined as in Algorithm 3. Let $K_1, K_2 \leftarrow \mathcal{K}$ be independent keys. We replace the calls to E by independent random permutations $\pi_1, \pi_2, \pi_3, \pi_4 \leftarrow \text{Perm}(\{0, 1\}^n)^4$. Let \mathbf{A} be a computationally bounded adversary that has access to three oracles $\mathcal{O}_1, \mathcal{O}_2$, and \mathcal{O}_3 for encryption, decryption, and leakage, respectively. \mathbf{A} shall distinguish between a real setting of $\text{RIV}_{F,\Pi}$ as above with a single-keyed block cipher E , and $\text{RIV}_{F,\Pi}$ which uses four independent uniformly chosen permutations $\pi^1, \pi^2, \pi^3, \pi^4 \leftarrow \text{Perm}(\{0, 1\}^n)$ with π^1 used in EHE^1 , π^2 used in EHE^2 , and π^3, π^4 used for XOR-CTR[π^3, π^4]. \mathbf{A} asks at most q queries of at most ℓ blocks and runs in time at most t . Then, we can upper bound the distinguishing advantage of \mathbf{A} by*

$$16.5\ell^2 \cdot \max\{\epsilon, 1/2^n\} + \text{Adv}_E^{\text{PRP}}(\ell + 3q, O(t)).$$

Theorem 13. *Let $d = 4$, $n = \tau = 128$, and $m \geq 2$ be even. Let $\text{RIV}_{F,\Pi}$ be as given in Algorithm 3 and let $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$ be computationally bounded IND-CPA, INT-CTXT, and ERR-CCA adversaries on $\text{RIV}_{F,\Pi}$, respectively, which run each in time at most t and ask at most q queries of at most ℓ blocks in total. Then, it holds that*

$$\begin{aligned} \text{Adv}_{\text{RIV}_{F,\Pi}}^{\text{IND-CPA}}(\mathbf{A}) &\leq \frac{2q^2 + \ell^2}{2^n} + q^2\epsilon + \delta_E, \\ \text{Adv}_{\text{RIV}_{F,\Pi}}^{\text{INT-CTXT}}(\mathbf{A}) &\leq \frac{2q^2 + q + \ell^2}{2^n} + q^2\epsilon + \delta_E, \\ \text{Adv}_{\text{RIV}_{F,\Pi}}^{\text{ERR-CCA}}(\mathbf{A}) &\leq \frac{8q^2 + 2q + 2\ell^2}{2^n} + 2q^2\epsilon + \delta_E, \end{aligned}$$

where $\delta_E = 16.5\ell^2 \cdot \epsilon + \text{Adv}_E^{\text{PRP}}(\ell + 3q, O(t))$ and $\epsilon \leq 9/2^{128}$.

The proof follows from Theorems 5, 8, 9, and 11, and those of the lemmata from Section 5 that can be found in the full version of this paper.

6.5 Performance of RIV

We implemented reference and optimized versions of RIV in C.⁵ Since the default key length for one iteration CLNH of $b = 1024$ bytes (which corresponds to $\text{CLHASH}^T[128, 2]$) appeared high, we tested also a variant with a smaller key size of $b = 256$ bytes for CLNH ($\text{CLHASH}^T[32, 2]$). Table 1 summarizes the results of our benchmarks. Our code was compiled using gcc v4.9.3 with options `-O3 -maes -mavx2 -mpclmul -march=native`, and run on (1) an Intel

⁵ Our code is open to the public domain: <https://github.com/medsec/riv>.

Platform	Instance	b	Message length (bytes)							
			128	256	512	1024	2048	4096	8192	16384
Haswell	RIV	256	3.81	2.78	2.14	1.81	1.62	1.48	1.40	1.37
	RIV	1024	3.53	2.13	1.81	1.49	1.37	1.29	1.25	1.22
	RIV (2-pass)	256	1.71	1.40	1.26	1.14	1.08	1.04	1.01	0.99
	RIV (2-pass)	1024	2.20	1.60	1.17	1.08	1.01	0.97	0.94	0.92
Broadwell	RIV	256	3.16	2.41	1.84	1.49	1.38	1.26	1.20	1.15
	RIV	1024	3.13	2.11	1.56	1.34	1.16	1.09	1.04	1.02
	RIV (2-pass)	256	2.16	1.67	1.30	1.09	1.03	0.95	0.92	0.90
	RIV (2-pass)	1024	2.19	1.50	1.14	1.01	0.92	0.86	0.84	0.82

Table 1: Performance results on Intel Haswell and Broadwell, respectively, in cycles per byte for the encryption with optimized implementations of RIV and a reduced version, which omits the second call to F . b denotes the key length for CLNH in bytes. Details regarding our setup are provided in the text.

Core i5-4200M (Haswell) at 2.50 GHz, and (2) on an Intel i5-5200 (Broadwell) at 2.20 GHz, both with the TurboBoost, SpeedStep, and HyperThreading technologies *disabled*. For measuring, we used the median of 10000 encryptions, omitting the cost for key setup, using the `rdtsc` instruction.

Our results show that RIV can run at less than 1.5 cpb on Haswell. Interestingly, a *SIV-like* reduced version of RIV, which is an easily obtained byproduct that simply omits the second call to F , represents a performant MRAE scheme with ≤ 1.04 cpb. This is slightly faster than the $4867/4096 \approx 1.17$ cpb reported for the manually assembly-optimized AES-GCM-SIV [22] and 1.06 cpb for the version of MRO with four-round BLAKE2b in [20], concerning messages of at least four KiB length on Haswell. Clearly, the reported performance of AEZv4 of about 0.7 cpb is unrivaled. Though, our construction provides a slightly higher security margin. Moreover, the security of AEZv4 bases on heuristic assumptions on four-round AES.

7 Conclusion

This work described a modular framework RIV for the construction of provably secure subtle AE schemes by extending the SIV framework from two to three passes. The obvious strength of RIV resides in the simplicity of its structure: it allows a straight-forward transformation of existing SIV-based constructions into subtle AE schemes. We proved the security in the standard model under notions that strive for ideal security goals; a further step could be to prove *achievable* security in the RAE setting with fixed stretch. Moreover, since the generic RIV construction bases only on PRF assumptions, this leaves open the possibility for proofs in the indistinguishability setting [40]. RIV is slightly less efficient than

earlier STPRP constructions, i.e., it employs three additional calls to an n -bit PRP, compared to e.g., a single call in HCTR-based [50] constructions. Since the use of a nonce-based encryption scheme $(\mathcal{E}, \mathcal{D})$ poses only the requirement on the IV to be a nonce, it might look to be sufficient to have two calls to universal hash functions instead of to calls to a PRF F . Yet, at least the outputs from the first invocation of F , $F_{K_1}^1(\cdot, \cdot, \cdot)$ must be unpredictable in order to prevent leaking information about the message in the tag. A potential future work can be to further study reductions of the design to target even higher efficiency. Nevertheless, we proposed an instantiation that is highly efficient on current x64 platforms and avoids the weak-key issues that were reported for GHASH-based polynomials in HCTR instantiations [49].

Acknowledgments. We thank all reviewers of the FSE 2016 for their helpful comments, Daniel Lemire and Owen Kaser for valuable notes on CLHASH, and Guy Barwell, Daniel Page, and Martijn Stam for insights into their work on subtle authenticated encryption. Christian Forler received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement no. 307952 and from the Silicon Valley Community Foundation under the Cisco Systems project *Misuse-Resistant Authenticated Encryption for Complex and Low-End Systems* (MIR-ACLE).

References

1. Farzaneh Abed, Christian Forler, David McGrew, Eik List, Scott Fluhrer, Stefan Lucks, and Jakob Wenzel. Pipelineable On-line Encryption. In Carlos Cid and Christian Rechberger, editors, *FSE*, volume 8540 of *LNCS*, pages 205–223. Springer, 2014.
2. Ross J. Anderson and Eli Biham. Two Practical and Provably Secure Block Ciphers: BEAR and LION. In Dieter Gollmann, editor, *FSE*, volume 1039 of *LNCS*, pages 113–120. Springer, 1996.
3. Elena Andreeva, Begül Bilgin, Andrey Bogdanov, Atul Luykx, Florian Mendel, Bart Mennink, Nicky Mouha, Qingju Wang, and Kan Yasuda. PRIMATES. <http://competitions.cr.yt.to/caesar-submissions.html>, 2014.
4. Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. How to Securely Release Unverified Plaintext in Authenticated Encryption. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT (1)*, volume 8873 of *LNCS*, pages 105–125. Springer, 2014.
5. Christian Badertscher, Christian Matt, Ueli Maurer, Phillip Rogaway, and Björn Tackmann. Robust Authenticated Encryption and the Limits of Symmetric Cryptography. In *IMA Int. Conf.*, volume 9496 of *LNCS*, pages 112–129. Springer, 2015.
6. Lear Bahack. Julius. <http://competitions.cr.yt.to/caesar-submissions.html>, 2014.
7. Guy Barwell, Daniel Page, and Martijn Stam. Rogue Decryption Failures: Reconciling AE Robustness Notions. In Jens Groth, editor, *IMA Int. Conf.*, volume 9496 of *LNCS*, pages 94–111. Springer, 2015.

8. Mihir Bellare and Chanathip Namprempre. Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. In Tatsuaki Okamoto, editor, *ASIACRYPT*, volume 1976 of *LNCS*, pages 531–545. Springer, 2000.
9. Mihir Bellare and Phillip Rogaway. Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In Tatsuaki Okamoto, editor, *ASIACRYPT*, pages 317–330, 2000.
10. Mihir Bellare, Phillip Rogaway, and David Wagner. The EAX Mode of Operation. In Bimal K. Roy and Willi Meier, editors, *FSE*, volume 3017 of *LNCS*, pages 389–407. Springer, 2004.
11. Daniel J. Bernstein. Polynomial evaluation and message authentication. <http://cr.yp.to/papers>, permanent ID: *b1ef3f2d385a926123e1517392e20f8c*, 2, 2007.
12. Guido Bertoni, Joan Daemen, Michael Peeters, Gilles Van Assche, and Ronny Van Keer. Using Keccak technology for AE: Ketje, Keyak and more, August 22 2014. SHA-3 2014 Workshop, UC Santa Barbara.
13. John Black, Shai Halevi, Hugo Krawczyk, Ted Krovetz, and Phillip Rogaway. UMAC: Fast and Secure Message Authentication. In Michael J. Wiener, editor, *CRYPTO*, pages 216–233, 1999.
14. Martin Boesgaard, Thomas Christensen, and Erik Zenner. Badger - A Fast and Provably Secure MAC. In John Ioannidis, Angelos D. Keromytis, and Moti Yung, editors, *ACNS*, volume 3531 of *LNCS*, pages 176–191, 2005.
15. Alexandra Boldyreva, Jean Paul Degabriele, Kenneth G. Paterson, and Martijn Stam. On Symmetric Encryption with Distinguishable Decryption Failures. In Shiho Moriai, editor, *FSE*, volume 8424 of *LNCS*, pages 367–390. Springer, 2013.
16. Debrup Chakraborty and Palash Sarkar. On Modes of Operations of a Block Cipher for Authentication and Authenticated Encryption. *IACR Cryptology ePrint Archive*, 2014:627, 2014.
17. Jean-Sébastien Coron, Yevgeniy Dodis, Avradip Mandal, and Yannick Seurin. A Domain Extender for the Ideal Cipher. In Daniele Micciancio, editor, *TCC*, volume 5978 of *LNCS*, pages 273–289. Springer, 2010.
18. Wei Dai and Ted Krovetz. VHASH security. *IACR Cryptology ePrint Archive*, 2007:338, 2007.
19. Ewan Fleischmann, Christian Forler, and Stefan Lucks. McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes. In Anne Canteaut, editor, *FSE*, pages 196–215, 2012.
20. Robert Granger, Philipp Jovanovic, Bart Mennink, and Samuel Neves. Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 (to appear)*, LNCS, 2016.
21. Shay Gueron and Michael E. Kounavis. Efficient Implementation of the Galois Counter Mode using a Carry-Less Multiplier and a Fast Reduction Algorithm. *Inf. Process. Lett.*, 110(14-15):549–553, 2010.
22. Shay Gueron and Yehuda Lindell. GCM-SIV: Full Nonce Misuse-Resistant Authenticated Encryption at Under One Cycle per Byte. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM Conference on Computer and Communications Security*, pages 109–119. ACM, 2015.
23. Shai Halevi and Phillip Rogaway. A Parallelizable Enciphering Mode. In Tatsuaki Okamoto, editor, *CT-RSA*, volume 2964 of *LNCS*, pages 292–304. Springer, 2004.
24. Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. Robust Authenticated-Encryption AEZ and the Problem That It Solves. In Elisabeth Oswald and Marc

- Fischlin, editors, *EUROCRYPT (1)*, volume 9056 of *LNCS*, pages 15–44. Springer, 2015.
25. Viet Tung Hoang, Reza Reyhanitabar, Phillip Rogaway, and Damian Vizár. Online Authenticated-Encryption and its Nonce-Reuse Misuse-Resistance. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO (1)*, volume 9215 of *LNCS*, pages 493–517. Springer, 2015.
 26. ISO/IEC. *19772:2009, Information technology – Security techniques – Authenticated Encryption*, 2009.
 27. Tetsu Iwata and Kan Yasuda. BTM: A Single-Key, Inverse-Cipher-Free Mode for Deterministic Authenticated Encryption. In Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini, editors, *Selected Areas in Cryptography*, pages 313–330, 2009.
 28. Tetsu Iwata and Kan Yasuda. HBS: A Single-Key Mode of Operation for Deterministic Authenticated Encryption. In Orr Dunkelman, editor, *FSE*, pages 394–415, 2009.
 29. Jérémy Jean, Ivica Nikolić, and Thomas Peyrin. Deoxys. <http://competitions.cr.yt.to/caesar-submissions.html>, 2014.
 30. Jérémy Jean, Ivica Nikolić, and Thomas Peyrin. Joltik. <http://competitions.cr.yt.to/caesar-submissions.html>, 2014.
 31. Hugo Krawczyk. LFSR-based Hashing and Authentication. In Yvo Desmedt, editor, *CRYPTO*, pages 129–139, 1994.
 32. Ted Krovetz. Message Authentication on 64-Bit Architectures. In Eli Biham and Amr M. Youssef, editors, *Selected Areas in Cryptography*, volume 4356 of *LNCS*, pages 327–341. Springer, 2006.
 33. Ted Krovetz. HS1-SIV. <http://competitions.cr.yt.to/caesar-submissions.html>, 2014.
 34. Daniel Lemire and Owen Kaser. Faster 64-bit universal hashing using carry-less multiplications. *Journal of Cryptographic Engineering*, pages 1–15, 2015.
 35. Michael Luby and Charles Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.
 36. Stefan Lucks. Faster Luby-Rackoff Ciphers. In Dieter Gollmann, editor, *FSE*, volume 1039 of *LNCS*, pages 189–203. Springer, 1996.
 37. Yishay Mansour, Noam Nisan, and Prasoona Tiwari. The Computational Complexity of Universal Hashing. In Harriet Ortiz, editor, *STOC*, pages 235–243. ACM, 1990.
 38. Ueli Maurer. Constructive Cryptography - A New Paradigm for Security Definitions and Proofs. In Sebastian Mödersheim and Catuscia Palamidessi, editors, *TOSCA*, volume 6993 of *LNCS*, pages 33–56. Springer, 2011.
 39. Ueli Maurer and Renato Renner. Abstract Cryptography. In Bernard Chazelle, editor, *ICS*, pages 1–21. Tsinghua University Press, 2011.
 40. Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology. In Moni Naor, editor, *Theory of Cryptography*, pages 21–39. Springer, 2004.
 41. Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering Generic Composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT*, volume 8441 of *LNCS*, pages 257–274. Springer, 2014.
 42. Moni Naor and Omer Reingold. On the Construction of Pseudorandom Permutations: Luby-Rackoff Revisited. *J. Cryptology*, 12(1):29–66, 1999.
 43. Reza Reyhanitabar, Serge Vaudenay, and Damian Vizár. Misuse-Resistant Variants of the OMD Authenticated Encryption Mode. In Sherman S. M. Chow, Joseph K.

- Liu, Lucas Chi Kwong Hui, and Siu-Ming Yiu, editors, *ProvSec*, volume 8782 of *LNCS*, pages 55–70. Springer, 2014.
44. Phillip Rogaway. Authenticated-Encryption with Associated-Data. In Vijayalakshmi Atluri, editor, *ACM Conference on Computer and Communications Security*, pages 98–107, 2002.
 45. Phillip Rogaway. Nonce-Based Symmetric Encryption. In Bimal K. Roy and Willi Meier, editors, *FSE*, volume 3017 of *LNCS*, pages 348–359. Springer, 2004.
 46. Phillip Rogaway and Thomas Shrimpton. A Provable-Security Treatment of the Key-Wrap Problem. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *LNCS*, pages 373–390. Springer, 2006.
 47. Thomas Shrimpton and R. Seth Terashima. A Modular Framework for Building Variable-Input-Length Tweakable Ciphers. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT (1)*, volume 8269 of *LNCS*, pages 405–423. Springer, 2013.
 48. Douglas R. Stinson. Universal Hashing and Authentication Codes. *Des. Codes Cryptography*, 4(4):369–380, 1994.
 49. Zhelei Sun, Peng Wang, and Liting Zhang. Weak-Key and Related-Key Analysis of Hash-Counter-Hash Tweakable Enciphering Schemes. In Ernest Foo and Douglas Stebila, editors, *ACISP*, volume 9144 of *LNCS*, pages 3–19. Springer, 2015.
 50. Peng Wang, Dengguo Feng, and Wenling Wu. HCTR: A Variable-Input-Length Enciphering Mode. In Dengguo Feng, Dongdai Lin, and Moti Yung, editors, *CISC*, volume 3822 of *LNCS*, pages 175–188. Springer, 2005.
 51. Shmuel Winograd. A new algorithm for inner product. *Computers, IEEE Transactions on*, 100(7):693–694, 1968.