

Parallelizable and Authenticated Online Ciphers

Elena Andreeva^{1,2}, Andrey Bogdanov³, Atul Luykx^{1,2}, Bart Mennink^{1,2},
Elmar Tischhauser^{1,2}, and Kan Yasuda^{1,4}

¹ Department of Electrical Engineering, ESAT/COSIC, KU Leuven, Belgium.

² iMinds, Belgium.

³ Department of Mathematics, Technical University of Denmark, Denmark.

⁴ NTT Secure Platform Laboratories, Japan.

Abstract. Online ciphers encrypt an arbitrary number of plaintext blocks and output ciphertext blocks which only depend on the preceding plaintext blocks. All online ciphers proposed so far are essentially serial, which significantly limits their performance on parallel architectures such as modern general-purpose CPUs or dedicated hardware. We propose the first parallelizable online cipher, COPE. It performs two calls to the underlying block cipher per plaintext block and is fully parallelizable in both encryption and decryption. COPE is proven secure against chosen-plaintext attacks assuming the underlying block cipher is a strong PRP. We then extend COPE to create COPA, the first parallelizable, online authenticated cipher with nonce-misuse resistance. COPA only requires two extra block cipher calls to provide integrity. The privacy and integrity of the scheme is proven secure assuming the underlying block cipher is a strong PRP. Our implementation with Intel AES-NI on a Sandy Bridge CPU architecture shows that both COPE and COPA are about *5 times faster* than their closest competition: TC1, TC3, and McOE-G. This high factor of advantage emphasizes the paramount role of parallelizability on up-to-date computing platforms.

Keywords: Block cipher, tweakable cipher, online cipher, authenticated encryption, nonce-misuse resistance, parallelizability, AES

1 Introduction

Online ciphers. A cipher which takes input of arbitrary length is said to be an *online cipher* if it can output ciphertext blocks as it is receiving the plaintext blocks. Specifically, the i th ciphertext block should only depend on the key and the first i plaintext blocks. This desirable functionality known more generally as online data processing is characteristic for other cryptographic primitives such as standard encryption schemes like CTR, CBC, OFB, and CFB.

The first theoretical treatment of online ciphers was put forward by Bellare, Boldyreva, Knudsen, and Namprempre [4]. They introduce the online ciphers HCBC1 and HCBC2, both of which require the use of two keys, one for the underlying block cipher and the other for the almost-xor-universal hash family [24]. Subsequently, Nandi [21] proposed two more efficient online ciphers MHCBC and

MCBC. MHCBC improves upon HCBC2 by using a smaller hashing key with a finite field multiplication as universal hash function, whereas MCBC does not even require a universal hash function, thus needing only one block cipher key and calling the block cipher twice per plaintext block. Rogaway and Zhang in [27] recast the formalism of Bellare et al. [4] in terms of tweakable block ciphers [17] and provide three generalizations of the previous online ciphers: TC1, TC2, and TC3.

Authenticated encryption from online ciphers. While online authenticated encryption (AE) schemes are not a novelty⁵, presently we are aware of only one family of online *and* misuse-resistant AE schemes, McOE [11]. McOE makes use of the online cipher TC3 [27] to build its general structure and adds two calls to the tweakable cipher to achieve authenticity. To process messages of arbitrary lengths, McOE applies a tag splitting method, similar to the ciphertext stealing method [9].

Bellare et al. [4] give a few generic transformations to turn an online cipher into a secure authenticated encryption scheme.

Problem statement. All existing online ciphers are highly sequential and none of them offer any possibility for parallelizing the computation between distinct block cipher calls. The only exception can be seen in TC1, which allows parallelization only in decryption but not in encryption. As a consequence, the McOE AE schemes are not parallelizable either, due to the fact that they are based on existing online ciphers.

At the same time, in the overwhelming majority of cases in practice, the underlying cipher is AES which is very well parallelizable on many platforms. Parallelization is a rather inherent feature of hardware implementations, both in ASIC and FPGA. Also in general-purpose software, parallelizable encryption algorithms have profited in terms of performance due to the bitslice approach for a long time already [18, 6, 14]. However, with the introduction of the hardware supported AES by Intel in general-purpose x86 CPUs as an instruction set AES-NI in Intel Westmere, Sandy Bridge, and Ivy Bridge — followed by the AMD adoption of AES-NI in AMD Bulldozer and Piledriver — the parallelizability of the AES modes of operation has become of truly paramount importance. With AES-NI, using a parallelizable mode of operation enables performance advantages of a factor 3 and higher — see, for instance, the case of the (serial) CBC encryption vs (parallel) CBC decryption [1].

Our contributions. We present the first parallelizable online cipher, COPE, and the first parallelizable online authenticated encryption scheme with nonce-misuse resistance, COPA.

COPE: Our novel design is illustrated in Fig. 1. To process a single plaintext block two block cipher calls are required. A secret mask (*tweak*) is applied

⁵ Examples of online AE schemes include EAX [5], GCM [19], and OCB1-3 [26, 25, 16].

to the plaintext block and used as input to the first block cipher call. Then the output of the second block cipher call is masked again to produce the ciphertext block.

By introducing dummy masks, each block cipher call can be viewed as an instance of the XEX construction [25], which uses the so-called “doubling” mask generation. Our basic design only deals with message lengths that are a multiple of the block length. In order to handle messages of arbitrary lengths we use the technique prescribed in the XLS domain extender by Ristenpart et al. in [23]. In contrast with previous designs, our scheme only uses a *single key* and a *single cryptographic primitive*, namely a block cipher.

COPE is proven IND-CPA up to the birthday bound of $n/2$ -bit security, where n denotes the block size of the underlying block cipher.

COPA: We transform COPE to support authentication, while maintaining parallelizability. The modifications are limited to computing an XOR sum of the plaintext data and using two extra block cipher calls; these can be seen in Fig. 2. The scheme also supports associated data in a way similar to how PMAC1 [25] operates. The privacy and integrity of COPA are proven up to the birthday bound.

To illustrate the impact of the parallelizability of our online schemes, we implement them with AES-NI on an Intel Sandy Bridge processor. We systematically compare the performance we attain with the online ciphers TC1, TC3, and MCBC as well as the online AE scheme McOE-G when instantiated with the AES. When compared to these closest online competitors, which are all explicitly not parallelizable, our modes provide performance improvements between a factor of 4.5 and 5, being below 2 cycles per byte on a single core. We expect almost a linear speed-up when several cores are available.

Organization of the paper. The remainder of the paper is organized as follows. We recall the necessary background on block ciphers in Section 2. Section 3 provides the specification of our new parallel modes. Sections 4 and 5 deal with the security proofs. Section 6 gives AES-NI implementations of our modes along with a systematic comparison to the state-of-the-art schemes. We conclude in Section 7.

2 Preliminaries

2.1 Block Ciphers

A block cipher $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a function that takes as input a key $k \in \mathcal{K}$ and a plaintext $M \in \{0, 1\}^n$, and produces a ciphertext $C = E(k, M)$. We sometimes write $E_k(\cdot) = E(k, \cdot)$. For a fixed key k , a block cipher is a permutation on n bits, and we denote the inverse permutation (decryption function) by E_k^{-1} .

Let $\text{Perm}(n)$ be the set of all permutations on n bits. When writing $x \stackrel{\$}{\leftarrow} X$ for some finite set X we mean that x is sampled uniformly from X . We write $\Pr[\mathbf{A} \mid \mathbf{B}]$ to denote the probability of event \mathbf{A} given \mathbf{B} .

Definition 1. Let E be a block cipher. The $\text{prp}\pm 1$ advantage of a distinguisher \mathcal{D} is defined as

$$\mathbf{Adv}_E^{\text{prp}\pm 1}(\mathcal{D}) = \left| \Pr_k[\mathcal{D}^{E_k, E_k^{-1}} = 1] - \Pr_{\pi}[\mathcal{D}^{\pi, \pi^{-1}} = 1] \right|.$$

Here, \mathcal{D} is a distinguisher with oracle access to either (E_k, E_k^{-1}) or (π, π^{-1}) . The probabilities are taken over $k \stackrel{\$}{\leftarrow} \mathcal{K}$, $\pi \stackrel{\$}{\leftarrow} \text{Perm}(n)$ and random coins of \mathcal{D} , if any. By $\mathbf{Adv}_E^{\text{prp}\pm 1}(t, q)$ we denote the maximum advantage taken over all distinguishers that run in time t and make q queries.

We shall also write $E_k^{\pm 1}$ for (E_k, E_k^{-1}) . Similarly, $\pi^{\pm 1}$ means (π, π^{-1}) , and so on.

2.2 Binary Fields

The set $\{0, 1\}^n$ of bit strings can be considered as the finite field $\text{GF}(2^n)$ consisting of 2^n elements. To do this, we represent an element of $\text{GF}(2^n)$ as a polynomial over the field $\text{GF}(2)$ of degree less than n . A string $a_{n-1}a_{n-2} \cdots a_1a_0 \in \{0, 1\}^n$ corresponds to the polynomial $a_{n-1}\mathbf{x}^{n-1} + a_{n-2}\mathbf{x}^{n-2} + \cdots + a_1\mathbf{x} + a_0 \in \text{GF}(2^n)$. The addition in the field is just the addition of polynomials over $\text{GF}(2)$ (that is, bitwise XOR, denoted by \oplus). To define multiplication in the field, we fix an irreducible polynomial $f(\mathbf{x})$ of degree n over the field $\text{GF}(2)$. Given two elements $a(\mathbf{x}), b(\mathbf{x}) \in \text{GF}(2^n)$, their product is defined as $a(\mathbf{x})b(\mathbf{x}) \bmod f(\mathbf{x})$ —polynomial multiplication over the field $\text{GF}(2)$ reduced modulo $f(\mathbf{x})$. We simply write $a(\mathbf{x})b(\mathbf{x})$ and $a(\mathbf{x}) \cdot b(\mathbf{x})$ to mean the product in the field $\text{GF}(2^n)$.

The set $\{0, 1\}^n$ can be also regarded as a set of integers ranging from 0 through $2^n - 1$. A string $a_{n-1}a_{n-2} \cdots a_1a_0 \in \{0, 1\}^n$ corresponds to the integer $a_{n-1}2^{n-1} + a_{n-2}2^{n-2} + \cdots + a_12 + a_0 \in [0, 2^n - 1]$. We often write elements of $\text{GF}(2^n)$ as integers, based on these conversions. So, for example, “2” means \mathbf{x} , “3” means $\mathbf{x} + 1$, and “7” means $\mathbf{x}^2 + \mathbf{x} + 1$. When we write multiplications such as $2 \cdot 3$ and 7^2 , we mean those in the field $\text{GF}(2^n)$.

2.3 XE and XEX Constructions of Tweakable Ciphers

Given a block cipher $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a secret mask $\Delta \in \{0, 1\}^n$, the ciphers

$$E'_{k, \Delta}(x) \stackrel{\text{def}}{=} E_k(x \oplus \Delta) \text{ or } E'_{k, \Delta}(x) \stackrel{\text{def}}{=} E_k(x \oplus \Delta) \oplus \Delta$$

behave like another block cipher independent of E_k (up to some bound). In the case of $E'_{k, \Delta}$, adversaries are allowed to make only forward queries, whereas

$E'_{k,\Delta}$ accepts both encryption and decryption queries. Now consider a set of secret masks $\{\Delta_i\}_{i \in \mathcal{T}}$, where Δ_i and Δ_j may not be necessarily independent. An index $i \in \mathcal{T}$ is called a tweak, which is not secret. We obtain a tweakable cipher $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \rightarrow \{0,1\}^n$ by defining $\tilde{E}_{k,i} \stackrel{\text{def}}{=} E'_{k,\Delta_i}$, and similarly $\tilde{E}_{k,j}$. We consider $\tilde{E}_{k,i}$ and $\tilde{E}_{k,j}$ together, where $i \in \mathcal{T}_0$, $j \in \mathcal{T}_1$ and $\mathcal{T}_0 \cap \mathcal{T}_1 = \emptyset$.

Definition 2. Let \tilde{E}, \tilde{E} be tweakable ciphers. The twk advantage of a distinguisher \mathcal{D} is defined as

$$\text{Adv}_{\tilde{E}, \tilde{E}}^{\text{twk}}(\mathcal{D}) = \left| \Pr_k[\mathcal{D}^{\tilde{E}_{k,i}, \tilde{E}_{k,j}^{\pm 1}} = 1] - \Pr_{\pi_i, \pi_j}[\mathcal{D}^{\pi_i, \pi_j^{\pm 1}} = 1] \right|.$$

Here, \mathcal{D} is a distinguisher with oracle access to a series of permutations. The tweaks run over $i \in \mathcal{T}_0$ and $j \in \mathcal{T}_1$ where $\mathcal{T}_0 \cap \mathcal{T}_1 = \emptyset$. By $\text{Adv}_{\tilde{E}, \tilde{E}}^{\text{twk}}(t, q)$ we denote the maximum advantage taken over all distinguishers that run in time t and make q queries in total.

The doubling method [25] enables us to produce many different values of the mask Δ from just one secret value $L \stackrel{\text{def}}{=} E_k(0)$. Namely, the masks are produced as $\Delta = 2^\alpha 3^\beta 7^\gamma L$ for varying indices of α, β and γ . To do this, we need to choose our irreducible polynomial $f(x)$ carefully. First, $f(x)$ needs to be primitive, meaning that 2 generates the whole multiplicative group. Second, we make sure that $\log_2 3$ and $\log_2 7$ are both “huge.” Third, we check if $\log_2 3$ and $\log_2 7$ are “apart enough” (modulo $2^n - 1$). We impose these conditions to ensure that values $2^\alpha 3^\beta 7^\gamma$ do not collide or become equal to 1. For example, when $n = 128$, the irreducible polynomial $f(x) = x^{128} + x^7 + x^2 + x + 1$ satisfies these requirements, making values $2^\alpha 3^\beta 7^\gamma$ all distinct and not equal to 1 for $\alpha \in [-2^{108}, 2^{108}]$ and $\beta, \gamma \in [-2^7, 2^7]$ [25], except for $(\alpha, \beta, \gamma) = (0, 0, 0)$. So we obtain tweakable ciphers $\tilde{E}_{k,\alpha\beta\gamma}$ and $\tilde{E}_{k,\alpha\beta\gamma}$.

Lemma 1 (XE and XEX [25]). Let $\mathcal{T}_0, \mathcal{T}_1 = \{(\alpha, \beta, \gamma)\}$ be two sets of integer triples such that $2^\alpha 3^\beta 7^\gamma$ are all distinct and not equal to 1, in particular $\mathcal{T}_0 \cap \mathcal{T}_1 = \emptyset$. Then the permutations $\{\tilde{E}_{k,\alpha\beta\gamma}\}_{\mathcal{T}_0} \cup \{\tilde{E}_{k,\alpha\beta\gamma}^{\pm 1}\}_{\mathcal{T}_1}$ are indistinguishable from independently random permutations $\{\pi_{\alpha\beta\gamma}\}_{\mathcal{T}_0} \cup \{\pi_{\alpha\beta\gamma}^{\pm 1}\}_{\mathcal{T}_1}$. Specifically, for given t, q , there exists a $t' \approx t$ such that

$$\text{Adv}_{\tilde{E}, \tilde{E}}^{\text{twk}}(t, q) \leq \frac{9.5q^2}{2^n} + \text{Adv}_E^{\text{prp}\pm 1}(t', 2q).$$

3 COPE and COPA: Design and Specification

We define COPE and COPA. COPE is an online cipher secure against chosen plaintext attacks. COPE makes two calls to the underlying block cipher per message block. COPA is an authenticated online cipher that builds on COPE. The additional cost of producing a tag is kept minimal—a message checksum and two extra block cipher calls. COPA accepts associated data input.

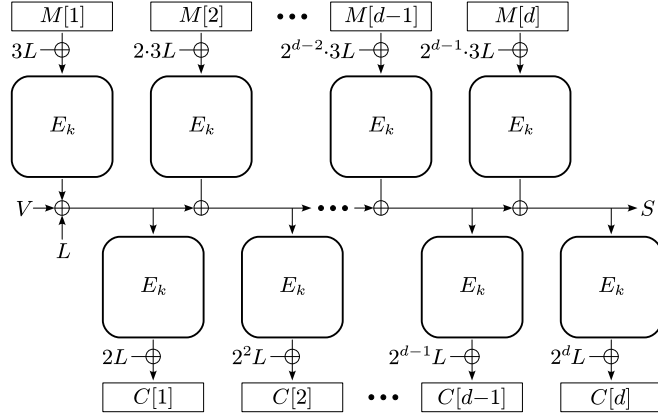


Fig. 1: Online cipher COPE. Set $V \stackrel{\text{def}}{=} 0$ for COPE. Variable S will be used later by COPA.

In this section we assume that the message length is a positive multiple of n . The length of associated data can be fractional. In App. A we show how to handle fractional messages with COPE and COPA. At the end of this section we give the design rationale for our constructions, explaining our choice of operations.

3.1 COPE Definition

Let $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be an n -bit block cipher, and denote $L \stackrel{\text{def}}{=} E_k(0)$. The encryption and decryption procedures of the COPE online cipher on a message $M[1]M[2] \cdots M[d]$ of d n -bit blocks and on a ciphertext $C[1]C[2] \cdots C[d]$ are then defined as:

<p>COPE-ENCRYPT $\mathcal{E}[E]$:</p> $\frac{V[0] \leftarrow L, \Delta_0 \leftarrow 3L, \Delta_1 \leftarrow 2L}{\text{for } i = 1, \dots, d \text{ do}}$ $V[i] \leftarrow E_k(M[i] \oplus \Delta_0) \oplus V[i-1]$ $C[i] \leftarrow E_k(V[i]) \oplus \Delta_1$ $\Delta_0 \leftarrow 2\Delta_0, \Delta_1 \leftarrow 2\Delta_1$ <p>end for</p>	<p>COPE-DECRYPT $\mathcal{E}^{-1}[E]$:</p> $\frac{V[0] \leftarrow L, \Delta_0 \leftarrow 3L, \Delta_1 \leftarrow 2L}{\text{for } i = 1, \dots, d \text{ do}}$ $V[i] \leftarrow E_k^{-1}(C[i] \oplus \Delta_1)$ $M[i] \leftarrow E_k^{-1}(V[i] \oplus V[i-1]) \oplus \Delta_0$ $\Delta_0 \leftarrow 2\Delta_0, \Delta_1 \leftarrow 2\Delta_1$ <p>end for</p>
---	--

The encryption operation is illustrated in Fig. 1.

3.2 COPA Definition

The core of the authenticated online cipher COPA is identical to COPE. The only differences are that first, an authentication tag T is generated after the COPE cipher invocation, and second, that associated data (if any) is processed before the cipher iteration to produce a value V that is XOR-ed into the first

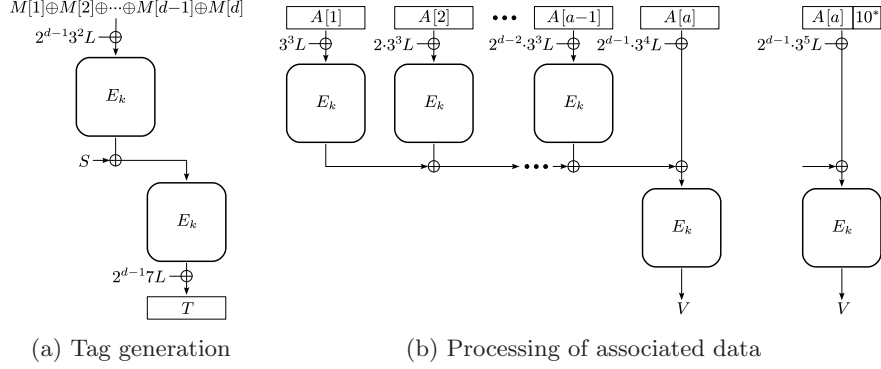


Fig. 2: Authenticated online cipher COPA: tag generation and processing of associated data.

intermediate block chaining (see Fig. 1): $V[0] \leftarrow V \oplus L$. If there is no associated data, then we set $V \stackrel{\text{def}}{=} 0$.

The tag T is computed by keeping a XOR checksum of the message blocks $\Sigma \stackrel{\text{def}}{=} M[1] \oplus \dots \oplus M[d]$ and computing

$$T \leftarrow E_k(E_k(\Sigma \oplus 2^{d-1}3^2L) \oplus S) \oplus 2^{d-1}7L,$$

with $S \stackrel{\text{def}}{=} V[d]$ denoting the last intermediate value in COPE’s block chaining, as in Fig. 1. The tag computation is illustrated in Fig. 2a. The value V is generated as follows: any associated data $A[1], \dots, A[a]$ is padded (if not a multiple of n bits) by a one and as many zeroes as necessary to obtain a multiple of the block size n . These blocks are then processed by a PMAC1-like [25] iteration as illustrated in Fig. 2b. Here, the block “ $A[a]10^*$ ” replaces the block “ $A[a]$ ” if $A[a]$ itself is not n bits. Tag verification occurs by checking if

$$S \oplus E_k(\Sigma \oplus 2^{d-1}3^2L) = E_k^{-1}(T \oplus 2^{d-1}7L),$$

where the tag is rejected if the equality is not true.

3.3 COPE and COPA for Arbitrary-Length Messages

We explain how to extend our schemes to accept “fractional” messages M in App. A. Here the length $|M|$ is not necessarily a positive multiple of the block size n . Note that simply using 10^* padding to M would result in ciphertext expansion. The methods described in App. A avoid such expansion with minimal loss of efficiency.

3.4 Design Rationale

One could combine universal hashing with a block cipher to design an AE scheme. Indeed, McOE-G [11] follows this approach. However, we decided to avoid the use

of universal hashing, for three reasons. First, the use of universal hashing would result in additional implementation cost, in particular with hardware. Second, recent study shows that there is an issue of weak keys with polynomial-based hashing [22]. Third, on the latest Intel CPUs, one call of AES is faster than one multiplication over the finite field $\text{GF}(2^{128})$, which is an operation used for polynomial-based hashing.

There has been discussion of whether one should use the doubling method or Gray code to produce tweak masks. We decided to use doubling, for three reasons. First, doubling provides us with the framework of XE and XEX constructions, which makes our constructions and proofs simple and easy to analyze. Neither our constructions nor our proofs can be directly translated into a Gray code version, as it is not immediately clear which masks we should use for the construction to make the proof work. Second, although it was reported that Gray code performs better than doubling on Intel CPUs [16], recent study shows that the doubling method can be implemented equally efficiently [3]. Third, the speedup of Gray code mask generation requires precomputation and memory, whereas doubling does not.

4 Privacy of COPE

4.1 Security Definition of Online Ciphers

We use the security definition of online ciphers from Rogaway and Zhang [27]. Let $(\{0, 1\}^n)^+$ denote the set of strings whose length is a positive multiple of n bits and is at most $2^n \cdot n$ bits. An online cipher $\mathcal{E} : \mathcal{K} \times (\{0, 1\}^n)^+ \rightarrow (\{0, 1\}^n)^+$ is a function such that it is a permutation on every block of n bits, having the additional feature that the outputs are the same for a common prefix. In other words, the first $|M|$ bits of $\mathcal{E}_k(M\|N)$ and $\mathcal{E}_k(M\|N')$ are the same for any $M, N, N' \in (\{0, 1\}^n)^+$. So an online cipher \mathcal{E}_k yields a permutation of i -th blocks, where the permutation is determined by the prefix (i.e. the first $i - 1$ blocks). Let $\text{OPerm}(n)$ be the set of all such permutations $\pi : (\{0, 1\}^n)^+ \rightarrow (\{0, 1\}^n)^+$.

Definition 3. Let \mathcal{E} be an online cipher. The IND-CPA advantage of a distinguisher \mathcal{D} is defined as

$$\text{Adv}_{\mathcal{E}}^{\text{cpa}}(\mathcal{D}) = \left| \Pr_k[\mathcal{D}^{\mathcal{E}_k} = 1] - \Pr_{\pi}[\mathcal{D}^{\pi} = 1] \right|.$$

Here, \mathcal{D} is a distinguisher with oracle access to either \mathcal{E}_k or π . The probabilities are taken over $k \xleftarrow{\$} \mathcal{K}$, $\pi \xleftarrow{\$} \text{OPerm}(n)$ and random coins of \mathcal{D} , if any. By $\text{Adv}_{\mathcal{E}}^{\text{cpa}}(t, q, \sigma, \ell)$ we denote the maximum advantage taken over all distinguishers that run in time t and make q queries, each of length at most ℓ blocks, and of total length at most σ blocks.




Fig. 3: IND-CPA proofs of COPE: introducing dummy masks rewriting the scheme in terms of XEX.

4.2 IND-CPA Proof Sketch

This section gives a sketch of the proof showing that COPE is secure against chosen-plaintext attacks with respect to privacy (IND-CPA). The details of the proof can be found in the full paper [2].

Theorem 1. *Let $\mathcal{E}[E]$ denote COPE, where E is the underlying block cipher. We have*

$$\mathbf{Adv}_{\mathcal{E}[E]}^{\text{cpa}}(t, q, \sigma, \ell) \leq \frac{38\sigma^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 4\sigma) + \frac{(\ell + 1)(q - 1)^2}{2^n},$$

where $t' \approx t$.

The proof consists of two steps. First, we rewrite COPE in terms of XEX constructions.⁶ We introduce dummy masks to the state values, as shown in Fig. 3. The block cipher calls in the upper layer are now $\tilde{E}_{k,\alpha-1,1,0}$ and those in the lower layer $\tilde{E}_{k,\alpha,0,0}$. Note that the “ L ” initially XORed to the state now disappears. We use Lem. 1 to replace the block cipher calls in the upper layer with random permutations $\pi_{\alpha-1,1,0}$ and those in the lower layer with $\pi_{\alpha,0,0}$ (for $\alpha = 1, 2, \dots$). Such a replacement costs us

$$\frac{9.5 \cdot (2\sigma)^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 2 \cdot 2\sigma) = \frac{38\sigma^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 4\sigma).$$

We write $\mathcal{E}[\pi]$ to denote the COPE scheme making calls to independent random permutations $\pi_{\alpha\beta\gamma}$ rather than to a block cipher.

Second, we show that $\mathcal{E}[\pi]$ behaves exactly the same as the ideal functionality, as long as collisions of state values do not occur. Define variables $V[\alpha]$ of state values as $V[\alpha] \stackrel{\text{def}}{=} \bigoplus_{i=1}^{\alpha} \pi_{i-1,1,0}(M[i])$ which is also equal to $\pi_{\alpha,0,0}^{-1}(C[\alpha])$.

⁶ The reason why our IND-CPA COPE is based on XEX constructions, and not on XEs, is because our COPA, which gives decryption oracle access to adversaries, builds upon COPE.

We look for *collisions* of these variables. Here by a “collision” roughly we mean the same value of $V[\alpha]$ coming from different prefixes $M[1]M[2] \cdots M[\alpha]$ and $M'[1]M'[2] \cdots M'[\alpha]$, for some α . More precisely, we have a collision of $V[\alpha] = V'[\alpha]$ if we have $V[\alpha - 1] \neq V'[\alpha - 1]$ and $V[\alpha] = V'[\alpha]$, which implies we must have $M[\alpha] \neq M'[\alpha]$ and also $M[i] \neq M'[i]$ for some $i < \alpha$. Let \mathbf{C} denote the event that a collision of $V[\alpha]$ occurs for some α .

Claim. Unless \mathbf{C} occurs, $\mathcal{E}[\pi]$ is indistinguishable from the ideal functionality. Furthermore, we have $\Pr[\mathcal{D}^{\mathcal{E}[\pi]} \text{ sets } \mathbf{C}] \leq (\ell + 1)/2^n$.

5 Privacy and Integrity of COPA

5.1 Security Definition of Authenticated Online Ciphers

Also for authenticated online ciphers, we use the IND-CPA security advantage of Def. 3, except that the ideal encryption oracle now has an additional random function that maps $\{0, 1\}^* \times (\{0, 1\}^n)^+$ to $\{0, 1\}^n$, corresponding to $(A, M) \mapsto T$.

We use the notion of integrity of authenticated encryption schemes from Fleischmann et al. [11]. By \perp , we denote a function that returns \perp on every input.

Definition 4. Let \mathcal{E} be an online cipher. The integrity advantage of a distinguisher \mathcal{D} is defined as

$$\text{Adv}_{\mathcal{E}}^{\text{int}}(\mathcal{D}) = \left| \Pr_k[\mathcal{D}^{\mathcal{E}_k^{\pm 1}} = 1] - \Pr_k[\mathcal{D}^{\mathcal{E}_k, \perp} = 1] \right|.$$

Here, \mathcal{D} is a distinguisher with oracle access to either $(\mathcal{E}_k, \mathcal{E}_k^{-1})$ or (\mathcal{E}_k, \perp) . To avoid a trivial win, we assume that the distinguisher does not make a query (A, C, T) if it has made a query (A, M) to the encryption oracle and obtained (C, T) from the oracle. By $\text{Adv}_{\mathcal{E}}^{\text{int}}(t, q, \sigma, \ell)$ we denote the maximum advantage taken over all distinguishers that run in time t and make q queries, each of length at most ℓ blocks, and of total length at most σ blocks.

5.2 Privacy of COPA

We now give a proof sketch of the IND-CPA security of COPA. The details can be found in the full paper [2].

Theorem 2. Let $\mathcal{E}[E]$ denote COPA, where E is the underlying block cipher. We have

$$\text{Adv}_{\mathcal{E}[E]}^{\text{cpa}}(t, q, \sigma, \ell) \leq \frac{39(\sigma + q)^2}{2^n} + \text{Adv}_E^{\text{prp}\pm 1}(t', 4(\sigma + q)) + \frac{(\ell + 2)(q - 1)^2}{2^n},$$

where $t' \approx t$.

The IND-CPA security analysis of COPE carries over, with only minor modifications. First, we introduce dummy masks in a similar way (to the encryption part), and replace all XE (in the associated-data part) and XEX constructions by random permutations. This replacement costs us

$$\frac{9.5 \cdot (2\sigma + 2q)^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 2 \cdot 2(\sigma + q)) = \frac{38(\sigma + q)^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 4(\sigma + q)).$$

Write $\mathcal{E}[\pi, \boldsymbol{\pi}]$ to denote the COPA scheme calling random permutations instead of a block cipher.

Next, we again use the collision event \mathbf{C} , but introduce two more events. One is \mathbf{A} , which is the event that we have a collision of V for two different associated data. Recall that for $A = \emptyset$, we have $V = 0$. The other is \mathbf{T} , which is the event that we have a collision of tag values for messages of the same length (or more precisely, a collision of input values to a random permutation that produces tags).

Claim. Unless $\mathbf{A} \vee \mathbf{C} \vee \mathbf{T}$ occurs, $\mathcal{E}[\pi, \boldsymbol{\pi}]$ is indistinguishable from the ideal functionality.

Lemma 2 (PMAC1, [25]). *The function $H[\pi] : \{0, 1\}^* \rightarrow \{0, 1\}^n$ ($A \mapsto V$) is indistinguishable from a random function $\Phi : \{0, 1\}^* \rightarrow \{0, 1\}^n$. Specifically, the distinguishing advantage (defined accordingly, only forward queries) is at most $\sigma^2/2^n$. Here, $\{0, 1\}^*$ denotes the set of strings whose length is at most $2^n \cdot n$ bits.*

So now we replace XE and XEX constructions with random permutations and $H[\pi]$ with a random function Φ . Denote the scheme by $\mathcal{E}[\Phi, \boldsymbol{\pi}]$. Then we have the following.

Claim. We have $\Pr[\mathcal{D}^{\mathcal{E}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{A}] \leq q^2/2^n$ and $\Pr[\mathcal{D}^{\mathcal{E}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{C} \vee \mathbf{T} \mid \neg \mathbf{A}] \leq (\ell + 2)(q - 1)^2/2^n$.

5.3 Integrity of COPA

The integrity proof of COPA is more involved than the privacy proofs and we include the full proof in this paper. We prove the following theorem:

Theorem 3. *Let $\mathcal{E}[E]$ denote COPA, where E is the underlying block cipher. We have*

$$\mathbf{Adv}_{\mathcal{E}[E]}^{\text{int}}(t, q, \sigma, \ell) \leq \frac{39(\sigma + q)^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 4(\sigma + q)) + \frac{(\ell + 2)(q - 1)^2}{2^n} + \frac{2q}{2^n},$$

where $t' \approx t$.

Let \mathbf{F} denote the event that the decryption oracle \mathcal{E}_k^{-1} returns something other than \perp . Clearly the two games are the same as long as the event \mathbf{F} does not occur, so we have

$$\Pr[\mathcal{D}^{\mathcal{E}_k^{\pm 1}} = 1] - \Pr[\mathcal{D}^{\mathcal{E}_k, \perp} = 1] \leq \Pr[\mathcal{D}^{\mathcal{E}_k^{\pm 1}} \text{ sets } \mathbf{F}].$$

In the rest of this section we shall bound this probability. First, as usual, we replace block cipher calls with random permutations $\pi, \boldsymbol{\pi}$. Then we replace the PMAC1 part of processing associated data with a random function Φ . These all together cost us (cf. the proof of Thm. 2)

$$\frac{38(\sigma + q)^2}{2^n} + \frac{\sigma^2}{2^n} + \mathbf{Adv}_E^{\text{prp}\pm 1}(t', 4(\sigma + q)).$$

Removing “Privacy Part.” Define events \mathbf{A} , \mathbf{C} and \mathbf{T} as we have done in the privacy proof of Thm. 2. Note that these events are defined in terms of variables $V[\alpha]$, where we also define $V[0] \stackrel{\text{def}}{=} V$ and $V'[\alpha + 1]$ the input value to the block cipher that produces a tag. We define these values as being set only by the queries to the encryption oracle \mathcal{E} . We do not let queries to the decryption oracle \mathcal{E}^{-1} affect variables $V[\cdot], V'[\cdot]$, whether or not it returns a message or \perp . Set $\mathbf{E} \stackrel{\text{def}}{=} \mathbf{A} \vee \mathbf{C} \vee \mathbf{T}$.

Next we define similar events \mathbf{A}' , \mathbf{C}' and \mathbf{T}' . These are exactly the same as the previous ones, except that now we consider only those events (i.e. collisions of $V[\cdot]$ or $V'[\cdot]$) that occur prior to a forgery (that is, under the condition $\neg \mathbf{F}$). Again, set $\mathbf{E}' \stackrel{\text{def}}{=} \mathbf{A}' \vee \mathbf{C}' \vee \mathbf{T}'$.

When we consider event \mathbf{F} , we would like to assume that we are under the condition $\neg \mathbf{E}'$, meaning that the encryption oracle \mathcal{E} has behaved ideally so far (till forgery). To do this, we use the inequality

$$\Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{F}] \leq \Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}'] + \Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{E}'].$$

We shall construct a distinguisher \mathcal{D}' that breaks the privacy of the encryption oracle \mathcal{E} . The distinguisher \mathcal{D}' uses \mathcal{D} , and the query complexity of \mathcal{D}' is at most that of \mathcal{D} . Specifically, \mathcal{D}' starts running \mathcal{D} , answering \mathcal{E} -queries using its \mathcal{E} oracle, and whenever \mathcal{D} makes a query to the decryption oracle \mathcal{E}^{-1} , \mathcal{D}' replies with a \perp .

Claim. We have $\Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}[\Phi, \boldsymbol{\pi}]} \text{ sets } \mathbf{E}'] \leq q^2/2^n + (\ell + 2)(q - 1)^2/2^n$.

Proof. Note that if $\mathcal{D}^{\mathcal{E}^{\pm 1}}$ sets \mathbf{E}' , then till this event \mathcal{D}' simulates the environment of \mathcal{D} correctly. Hence we get $\Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}} \text{ sets } \mathbf{E}'] \leq \Pr[\mathcal{D}'^{\mathcal{E}} \text{ sets } \mathbf{E}']$, which is less than $q^2/2^n + (\ell + 2)(q - 1)^2/2^n$ as shown in the privacy proof. \square

Passing to a Single-Query Adversary. So it remains to evaluate the probability that \mathcal{D} sets \mathbf{F} under the condition $\neg \mathbf{E}'$. We shall construct a forger \mathcal{D}_1 from \mathcal{D} . The forger \mathcal{D}_1 makes multiple queries to the encryption oracle \mathcal{E} but makes only one query to the decryption oracle \mathcal{E}^{-1} at the end of its run. We define \mathcal{D}_1 as follows: it chooses a random index $i \in [1, q]$. It then runs \mathcal{D} , answering its \mathcal{E} -queries using the \mathcal{E} oracle of \mathcal{D}_1 and answering the queries to the decryption oracle \mathcal{E}^{-1} with \perp . When \mathcal{D} makes the i -th query (A^*, C^*, T^*) to the decryption oracle, \mathcal{D}_1 outputs the query (A^*, C^*, T^*) and stops (or more precisely, makes that query to the decryption oracle \mathcal{E}^{-1} and stops.)

Claim. We have $\Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}[\Phi, \pi]} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}'] \leq q \Pr[\mathcal{D}_1^{\mathcal{E}^{\pm 1}[\Phi, \pi]} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}']$.

Proof. Let \mathbf{F}_h denote the event that at the h -th query the decryption oracle \mathcal{E}^{-1} returns something other than \perp for the first time; that is, the oracle has returned only \perp so far. Clearly these are disjoint events, and we have $\mathbf{F} = \bigvee_{h=1}^q \mathbf{F}_h$. Then, under the events $\neg \mathbf{E}'$ and $i = h$, the forger \mathcal{D}_1 correctly simulates the game of \mathcal{D} . Therefore, we get $\Pr[\mathcal{D}_1^{\mathcal{E}^{\pm 1}} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}'] \geq \Pr[(i = h) \wedge \mathcal{D}^{\mathcal{E}^{\pm 1}} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}'] \geq (1/q) \Pr[\mathcal{D}^{\mathcal{E}^{\pm 1}} \text{ sets } \mathbf{F} \mid \neg \mathbf{E}']$. \square

Evaluating Forgery Probabilities. Let (A^*, C^*, T^*) denote the (non-trivial) query made by \mathcal{D}_1 to the decryption oracle $\mathcal{E}^{-1}[\Phi, \pi]$. We shall evaluate the probability that this would make \mathcal{E}^{-1} return something other than \perp . To evaluate the probability, we shall consider the following cases.

Lemma 3 (Case 1). *If A^* or T^* is new, or C^* contains a new block, then the probability of a forgery is at most $2/2^n$.*

Proof. If A^* is new, then it means that it triggers the random function Φ and yields a fresh random value $V \leftarrow \Phi(A^*)$. This value is XORed to the value that is input to the block cipher to produce the tag, which must be equal to T^* . All other values XORed to the value are independent of V . Hence, regardless of the values C^*, T^* , the probability of such an event is at most $1/2^n$.

Say that A^* is not new, but that C^* contains a new block. Let $C^*[\alpha]$ be one of the new blocks. The decryption invokes $\pi_{\alpha, 0, 0}^{-1}(C^*[\alpha])$, which is sampled from the set of at least $2^n - q$ points. Therefore, the probability of a forgery is at most $1/(2^n - q) \leq 2/2^n$, assuming $q \leq 2^{n-1}$.

Say that A^* is not new, C^* does not contain a new block, but T^* is new. This is similar to the previous case. This would trigger a fresh point of $\pi_{d^*-1, 0, 1}^{-1}(T^*)$, where d^* denotes the number of blocks in the message M^* . The point is sampled from the set of at least $2^n - q$ points. Therefore, the probability of a forgery is at most $1/(2^n - q) \leq 2/2^n$. \square

Lemma 4 (Case 2). *If A^* and T^* are old, and C^* consists of old blocks, then the probability of a forgery is at most $2/2^n$.*

Proof. To handle this case, we introduce some notation. For the query (A^*, C^*, T^*) in question, divide C^* into blocks as $C^*[1]C^*[2] \cdots C^*[d^*] \leftarrow C^*$ and define $C^*[0] \stackrel{\text{def}}{=} A^*$ and $C^*[d^* + 1] \stackrel{\text{def}}{=} T$. We then focus on a pair of adjacent “blocks” $(C^*[\alpha - 1], C^*[\alpha])$ for $\alpha = 1, 2, \dots, d^* + 1$. We call a pair *old* if it (as a pair) has already appeared in some previous query made to the encryption oracle \mathcal{E} and in the corresponding value returned by the oracle. That is, if \mathcal{D} has made a query (A', M') to the oracle and got (C', T') back, then we check if the pair in question $(C^*[\alpha - 1], C^*[\alpha])$ is contained in (A', C', T') —that is, we check if $(C^*[\alpha - 1], C^*[\alpha]) = (C'[\alpha - 1], C'[\alpha])$ holds, where $C'[0]$ and $C'[d' + 1]$ are defined similarly. We do this for all previous queries. We call the pair $(C^*[\alpha - 1], C^*[\alpha])$ *new* otherwise.

Note that the query (A^*, C^*, T^*) always contains a new pair. If (A^*, C^*, T^*) contains no new pairs, then, given the non-triviality of the query, we must have observed a collision, contradicting the assumption $\neg \mathbf{E}'$.

We now make a distinction among new pairs $(C^*[\alpha - 1], C^*[\alpha])$ based on the decrypted message block $M^*[\alpha]$ from the two adjacent ciphertext blocks. We say that a pair is *collapsing* if there exists a previous query (A', M') made by \mathcal{D} to the encryption oracle \mathcal{E} such that $M'[\alpha] = M^*[\alpha]$.

There exists a new pair $(C^[\alpha - 1], C^*[\alpha])$ that is not collapsing.* This case means that we trigger a random sampling of $\pi_{\alpha,1,0}^{-1}$ to compute $M^*[\alpha]$. Then, note that the value $\Sigma^* = M^*[1] \oplus M^*[2] \oplus \dots \oplus M^*[d^*]$ is already uniquely determined by the values $C^*[d^*]$ and T^* (via Fig. 2a). There are at least $2^n - q$ possible values for $M^*[\alpha]$, and the message blocks must sum up to this particular value Σ^* , which happens with a probability at most $1/(2^n - q) \leq 2/2^n$.

All new pairs in (A^, C^*, T^*) are collapsing.* This final case is quite different from the previous ones above, as we do not have any fresh sampling of permutations $\pi_{\alpha,\beta,\gamma}^{\pm 1}$ or the random function Φ in evaluating $\mathcal{E}^{-1}[\Phi, \pi](A^*, C^*, T^*)$. To tackle this case, we shall convert this forgery game into one where the adversary \mathcal{D}° tries to find multiple collisions by outputting the following set of values, without making any query to the oracles:

1. $r \in [1, \ell]$,
2. $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_r \leq \ell + 1$,
3. $(A_1, M_1), (A_2, M_2), \dots, (A_r, M_r)$, and
4. $(A'_1, M'_1), (A'_2, M'_2), \dots, (A'_r, M'_r)$.

The adversary \mathcal{D}° “wins” if the submitted values form a multi-collision in the following sense: (A_i, M_i) and (A'_i, M'_i) collides at the α_i -th block, for all $i \in [1, r]$. The adversary \mathcal{D}° runs \mathcal{D}_1 , simulating the \mathcal{E} oracle with an ideal functionality. Note that this simulation is correct under the condition $\neg \mathbf{E}'$. When \mathcal{D}_1 outputs (A^*, C^*, T^*) , \mathcal{D}° first checks for new pairs contained in it. Let $1 \leq \alpha_1 < \alpha_2 < \dots < \alpha_r \leq \ell + 1$ be the positions of new pairs. Then \mathcal{D}° checks the history of values (C, T) that it returned. Note that under $\neg \mathbf{E}'$, a block $C^*[\alpha]$ determines a unique prefix AM . We choose (A_i, M_i) to be the prefix determined by $C^*[\alpha_i]$. To choose (A'_i, M'_i) , let $A'_i M''$ be the prefix determined by $C^*[\alpha_i - 1]$. Then \mathcal{D}° chooses randomly, from the previously queried values, a message block $M[\alpha] \neq M_i[\alpha]$. Set $M'_i \stackrel{\text{def}}{=} M'' M[\alpha]$. The adversary \mathcal{D}° does this for $i = 1, 2, \dots$ except for the last block.

- If $\alpha_r < d^* + 1$, then we know the message checksum $\Sigma^* = M^*[1] \oplus \dots \oplus M^*[d^*]$, so \mathcal{D}° does not have to guess the value of $M'_{\alpha_r}[\alpha_r]$.
- If $\alpha_r = d^* + 1$, then we simply set the last input value to be the checksum of all previous (guessed) message blocks.

Now we observe that as long as all the guesses of the message blocks are correct, \mathcal{D}° wins if \mathcal{D}_1 succeeds in forgery of this type. It should be noted that the values returned by \mathcal{E} are independent of the success probabilities in question, under the

event $\neg \mathbf{E}'$. Therefore, for a fixed r ,

$$\Pr[\mathcal{D}^\circ \text{ wins} \mid r] \geq \frac{1}{q-1} \cdots \frac{1}{q-1} \Pr[\mathcal{D}_1 \text{ forges} \mid r] = \frac{1}{(q-1)^{r-1}} \Pr[\mathcal{D}_1 \text{ forges} \mid r].$$

We then calculate $\Pr[\mathcal{D}^\circ \text{ wins} \mid r]$. We do this by lazy sampling of the permutations, and we see that, for a fixed r ,

$$\Pr[\mathcal{D}^\circ \text{ wins} \mid r] \leq \frac{1}{2^n - 1} \cdot \frac{1}{2^n - 1} \cdots \frac{1}{2^n - 1} = \frac{1}{(2^n - 1)^r}.$$

Hence by varying r we get in total

$$\Pr[\mathcal{D}_1 \text{ forges}] \leq \sum_{i=1}^{\ell} \frac{(q-1)^{i-1}}{(2^n - 1)^i} \Pr[i = r] \leq \frac{1}{(2^n - 1)} \sum_{i=1}^{\ell} \Pr[i = r] = \frac{1}{(2^n - 1)}.$$

6 Efficient Parallel Implementation

6.1 The Setting

We compare our schemes to some prominent existing online ciphers: TC1 and TC3 [27] being the most efficient previous schemes; and MCBC [21] as a representative for a scheme relying only on block cipher invocations (as opposed to tweakable block ciphers or universal hashing). The modes HCBC1 and MHCBC are implicitly covered by TC1 and TC3, and HCBC2 has a performance inferior to TC3.

For the case of authenticated online ciphers, we exclude modes of operation and dedicated designs that are based on a nonce and rely on its non-reuse (e.g., GCM [19], OCB [16], ALE [7], and AEGIS [28]). Therefore, we compare our COPA design to the McOE family of authenticated encryption algorithms [11], which, to the best of our knowledge, is the only other online scheme not relying on the non-reuse of a nonce. We focus on the McOE-G instance, since McOE-X itself is not secure [20], featuring a key recovery with birthday complexity.

For the concrete instantiation of all schemes, we use the AES-128 [10] as the underlying block cipher, and multiplication in $\text{GF}(2^{128})$ as an almost XOR-universal hash function [15]. As target platform for the implementations, we chose the recent generation of Intel microprocessors (Westmere or later) which support the AES-NI instruction set [12] and carryless multiplication [13].

6.2 Implementation characteristics of COPE and COPA

The online modes proposed in this paper can utilize parallelized execution of block cipher calls in two ways: for messages longer than one block, the encryptions of subsequent message blocks can be carried out independently of each other once the respective masks have been XORed. The same holds for the second series of block cipher calls, once the chaining XORs have been executed.

Table 1: Software performance of (authenticated) online ciphers based on the AES on the Intel Sandy Bridge platform (AES-NI). All numbers are given in cycles per byte (cpb).

Algorithm	message length (bytes)						
	128	256	512	1024	2048	4096	8192
CTR	1.74	1.27	1.05	0.93	0.86	0.83	0.82
TC1	9.00	8.75	8.65	8.60	8.56	8.56	8.56
TC3	9.08	8.82	8.72	8.67	8.63	8.63	8.62
MCBC	11.66	11.00	10.68	10.52	10.44	10.40	10.38
COPE	2.56	2.08	1.89	1.78	1.72	1.70	1.69
McOE-G	10.85	9.73	9.14	8.90	8.74	8.69	8.66
COPA	3.78	2.85	2.31	2.06	1.94	1.88	1.85

This parallelism can be exploited in a single-core scenario by pipelining the block cipher rounds for several consecutive block cipher invocations. Similarly, these invocations can be processed independently by multiple threads, with the recombination being the computation of the chaining. Note that both scenarios can be combined when multiple cores with pipelined block cipher calls are available, which is typically the case for Intel’s AES-NI architecture.

On the recent Sandy and Ivy Bridge platforms, the AES round function can be computed at a latency of 8 cycles with a throughput of 1 cycle. Consequently, to fully utilize the pipeline, our implementation issues 8 AES round function evaluations on the next 8 consecutive blocks (independent data and same key). The tweak masks are computed using dedicated multiplication routines for 2^α , 3^β and $7^\gamma \in \text{GF}(2^{128})$. By contrast, the general $\text{GF}(2^{128})$ multiplication needed for TC1, TC3, and McOE-G is implemented using the PCLMULQDQ carryless multiplication instruction followed by modular reduction.

6.3 Performance measurements

We provide performance data for the (authenticated) encryption of messages of length $16 \cdot 2^b$ bytes, with $3 \leq b \leq 10$. The performance of AES-CTR is provided as a reference point. All measurements were taken on a single core of an Intel Core i5-2520M CPU at 2500 MHz, averaged over $5 \cdot 10^5$ repetitions, processing one message at a time. Our findings are summarized in Table 1. All numbers are given in cycles per byte (cpb).

One can observe that for all message lengths, the parallelizability of the proposed schemes results in speed-ups of factor 4.5 – 5 in comparison to the existing modes, at least for somewhat longer messages. By fully utilizing the pipeline, our schemes are only marginally slower than two times AES-CTR, which implies that the overhead imposed by the computation of the masks and the chaining is kept at a minimum. The authenticated mode COPA carries the additional overhead of two more AES calls plus field arithmetic for finalization, but this quickly becomes insignificant as the message length increases. Note,

however, that some constant overhead in comparison to the unauthenticated mode remains even for very long messages: this can be attributed to the fact that the computation of the checksum does not allow overwriting the message blocks, leading to increased register pressure. We also note that with the availability of carryless multiplication, TC1 and TC3 can be implemented more efficiently than the purely block cipher-based MCBC which was created with the goal to improve performance by avoiding field arithmetic.

The performance of our parallelizable schemes COPE and COPA can be further improved by utilizing multiple cores. Our implementation of multithreaded encryption confirms the intuition that one can expect a nearly linear speedup when using multiple cores for computing our schemes (i.e., the cost is < 1 cpb for two cores and so on).

7 Conclusion

By presenting COPE, our work provides the first solution for a parallelizable online cipher. Building on COPE, we go on to construct COPA, the first parallelizable and nonce-misuse resistant online authenticated encryption scheme. Our implementations of COPE and COPA with Intel AES-NI on a Sandy Bridge processor architecture benefit strongly from the parallelism, which gives us speedups of about factor 5 in comparison to existing (serial) online ciphers TC1, TC3, MCBC and the online AE scheme McOE-G.

Our designs additionally employ only a single key and use only a block cipher as a building block—as opposed to tweakable block ciphers or universal hash functions. We prove that our cipher COPE is an IND-CPA secure online permutation. The privacy result is also carried over to COPA. The integrity proof of COPA uses a technique of converting a forgery to a set of multiple collisions. It seems that the technique has not been used before by security proofs of parallelizable authenticated encryption mode or message authentication code. The technique may be applicable to other new types of parallelizable modes of operation. We leave it as an interesting open problem to construct a scheme with less primitive calls but with comparable security guarantees.

ACKNOWLEDGMENTS. This work has been funded in part by the IAP Program P6/26 BCRYPT of the Belgian State (Belgian Science Policy), in part by the European Commission through the ICT program under contract ICT-2007-216676 ECRYPT II, in part by the Research Council KU Leuven: GOA TENSE, and in part by the Research Fund KU Leuven, OT/08/027. Elena Andreeva is supported by a Postdoctoral Fellowship from the Flemish Research Foundation (FWO-Vlaanderen). Bart Mennink is supported by a Ph.D. Fellowship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

References

1. Akdemir, K., Dixon, M., Feghali, W., Fay, P., Gopal, V., Guilford, J., Erdinc Ozturk, G.W., Zohar, R.: Breakthrough AES Performance with Intel AES New Instructions. Intel white paper (January 2010)
2. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Tischhauser, E., Yasuda, K.: Parallelizable and authenticated online ciphers. *Cryptology ePrint Archive* (2013), full version of this paper
3. Aoki, K., Iwata, T., Yasuda, K.: How Fast Can a Two-Pass Mode Go? A Parallel Deterministic Authenticated Encryption Mode for AES-NI (Extended Abstract of Work in Progress). *Directions in Authenticated Ciphers (DIAC)* (July 2012)
4. Bellare, M., Boldyreva, A., Knudsen, L.R., Namprempre, C.: Online Ciphers and the Hash-CBC Construction. In: Kilian, J. (ed.) *CRYPTO*. *Lecture Notes in Computer Science*, vol. 2139, pp. 292–309. Springer (2001)
5. Bellare, M., Rogaway, P., Wagner, D.: The EAX Mode of Operation. In: Roy, B.K., Meier, W. (eds.) *FSE*. *Lecture Notes in Computer Science*, vol. 3017, pp. 389–407. Springer (2004)
6. Bernstein, D.J., Schwabe, P.: New AES Software Speed Records. In: Chowdhury et al. [8], pp. 322–336
7. Bogdanov, A., Mendel, F., Regazzoni, F., Rijmen, V., Tischhauser, E.: ALE: AES-Based Lightweight Authenticated Encryption. In: *FSE'13*. *Lecture Notes in Computer Science*, Springer (2013), to appear
8. Chowdhury, D.R., Rijmen, V., Das, A. (eds.): *Progress in Cryptology - INDOCRYPT 2008*, 9th International Conference on Cryptology in India, Kharagpur, India, December 14-17, 2008. *Proceedings, Lecture Notes in Computer Science*, vol. 5365. Springer (2008)
9. Daemen, J.: *Hash Function and Cipher Design: Strategies Based on Linear and Differential Cryptanalysis*. Ph.D. thesis, Katholieke Universiteit Leuven, Leuven, Belgium (1995)
10. Daemen, J., Rijmen, V.: *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer (2002)
11. Fleischmann, E., Forler, C., Lucks, S.: McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes. In: Canteaut, A. (ed.) *FSE*. *Lecture Notes in Computer Science*, vol. 7549, pp. 196–215. Springer (2012)
12. Gueron, S.: *Intel Advanced Encryption Standard (AES) Instructions Set*. Intel white paper (September 2012)
13. Gueron, S., Kounavis, M.: Intel Carry-Less Multiplication Instruction and its Usage for Computing the GCM mode. Intel white paper (September 2012)
14. Käsper, E., Schwabe, P.: Faster and Timing-Attack Resistant AES-GCM. In: Clavier, C., Gaj, K. (eds.) *CHES*. *Lecture Notes in Computer Science*, vol. 5747, pp. 1–17. Springer (2009)
15. Krawczyk, H.: LFSR-based Hashing and Authentication. In: Desmedt, Y. (ed.) *CRYPTO*. *Lecture Notes in Computer Science*, vol. 839, pp. 129–139. Springer (1994)
16. Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) *FSE*. *Lecture Notes in Computer Science*, vol. 6733, pp. 306–327. Springer (2011)
17. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. In: Yung, M. (ed.) *CRYPTO*. *Lecture Notes in Computer Science*, vol. 2442, pp. 31–46. Springer (2002)

18. Matsui, M., Nakajima, J.: On the Power of Bitslice Implementation on Intel Core2 Processor. In: Paillier, P., Verbauwhede, I. (eds.) CHES. Lecture Notes in Computer Science, vol. 4727, pp. 121–134. Springer (2007)
19. McGrew, D.A., Viega, J.: The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT. Lecture Notes in Computer Science, vol. 3348, pp. 343–355. Springer (2004)
20. Mendel, F., Mennink, B., Rijmen, V., Tischhauser, E.: A Simple Key-Recovery Attack on McOE-X. In: Pieprzyk, J., Sadeghi, A.R., Manulis, M. (eds.) Cryptology and Network Security. LNCS, vol. 7712, pp. 23 – 31. Springer (2012)
21. Nandi, M.: Two New Efficient CCA-Secure Online Ciphers: MHCBC and MCBC. In: Chowdhury et al. [8], pp. 350–362
22. Procter, G., Cid, C.: On Weak Keys and Forgery Attacks against Polynomial-based MAC Schemes. In: FSE 2013 (2013)
23. Ristenpart, T., Rogaway, P.: How to Enrich the Message Space of a Cipher. In: Biryukov, A. (ed.) FSE. Lecture Notes in Computer Science, vol. 4593, pp. 101–118. Springer (2007)
24. Rogaway, P.: Bucket Hashing and its Application to Fast Message Authentication. In: Coppersmith, D. (ed.) CRYPTO. Lecture Notes in Computer Science, vol. 963, pp. 29–42. Springer (1995)
25. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT. Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer (2004)
26. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: Reiter, M.K., Samarati, P. (eds.) ACM Conference on Computer and Communications Security. pp. 196–205. ACM (2001)
27. Rogaway, P., Zhang, H.: Online Ciphers from Tweakable Blockciphers. In: CT-RSA 2011. Lecture Notes in Computer Science, vol. 6558, pp. 237–249. Springer, Heidelberg (2011)
28. Wu, H., Preneel, B.: AEGIS: A Fast Authenticated Encryption Algorithm. Directions in Authenticated Ciphers (July 2012)

A Handling Arbitrary-Length Messages

A.1 COPE for Arbitrary-Length Messages

Our solution relies on the XLS construction [23] of VIL tweakable ciphers. XLS makes only three block-cipher calls and requires only simple bit operations outside block-cipher calls.

Let $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a tweakable cipher and $E : \mathcal{K}' \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ a block cipher. Then $\text{XLS}[\tilde{E}, E]$ yields a VIL permutation on $\{0, 1\}^{n+*}$, the set of string whose length is between n bits and $2n - 1$ bits. Specifically, we get $\text{XLS}[\tilde{E}, E] : \mathcal{K} \times \mathcal{K}' \times \mathcal{T} \times \{0, 1\}^{n+*} \rightarrow \{0, 1\}^{n+*}$. Using appropriate choice of (α, β, γ) , we can realize the ciphers used in XLS by the underlying block cipher in COPE encryption scheme \mathcal{E} , dependent on the message length d . So we write $\text{XLS}_{k,d}$ to denote the XLS invocation in COPE.

Let M be a message of at least n bits. Divide it into blocks as $M[1]M[2] \cdots M[d-1]M[d] \leftarrow M$, and assume that we have $1 \leq |M[d]| \leq n-1$. Then we can define $C \leftarrow \mathcal{E}_k(M)$ as

$$\begin{aligned} C[1]C[2] \cdots C[d-2], S &\leftarrow \mathcal{E}_k(M[1]M[2] \cdots M[d-2]) \text{ (let } \mathcal{E}_k \text{ output } S \text{ for now)} \\ C[d-1]C[d] &\leftarrow \text{XLS}_{k,d}((M[d-1] \oplus S) \| M[d]) \\ C &\leftarrow C[1]C[2] \cdots C[d]. \end{aligned}$$

The IND-CPA proof of COPE carries over with minor modifications. Note that we have to “wait” the processing of $M[d-1]$ till receiving $M[d]$ (or “redo” after receiving), making the scheme less online. Yet, we make only three calls to the block cipher to process these two blocks.

We require $|M| \geq n$. As pointed out by [27], it seems a challenging problem to handle the case $|M| < n$ with encryption-only online ciphers in a secure manner.

A.2 COPA for Arbitrary-Length Messages

There are solutions of arbitrary-length messages for COPA also. This time we can take the advantage of the tag to handle even the case $|M| < n$. The solution for the case $|M| > n$ also becomes more efficient owing to the presence of tags.

Tag Splitting for $|M| < n$. We can do a trick similar to tag splitting [11] if $|M| < n$. We first choose appropriate parameters (α, β, γ) to make it independent of the ordinary COPA encryption algorithm \mathcal{E} . Write it \mathcal{E}_k^* (which will be used only for fractional one-block messages). Given M such that $|M| = s < n$, we can define $(C, T) \leftarrow \mathcal{E}_k(M)$ as

$$\begin{aligned} (C', T') &\leftarrow \mathcal{E}_k^*(M10^*) \\ C &\leftarrow \lceil C' \rceil_s \text{ (leftmost } s \text{ bits)} \\ T &\leftarrow \lceil C' \rceil_{n-s} \lceil T' \rceil_s. \end{aligned}$$

One can directly verify the security of this extension. Note that the integrity relies on the 10^* padding as well as on the “partial” tag $\lceil T' \rceil_s$.

XLS for $|M| > n$. Our solution for this case is similar to that of COPE but is more efficient, in that COPA still remains fully online. Again, let M be a message whose length is more than n bits. Divide it into blocks as $M[1]M[2] \cdots M[d-1]M[d] \leftarrow M$, and assume that we have $1 \leq |M[d]| \leq n-1$. Then we can define $(C, T) \leftarrow \mathcal{E}_k(M)$ as

$$\begin{aligned} (C', T') &\leftarrow \mathcal{E}_k(M[1]M[2] \cdots M[d-1]) \\ C[d]T &\leftarrow \text{XLS}_{k,d}(M[d]T') \\ C &\leftarrow C'C[d], \end{aligned}$$

where $\text{XLS}_{k,d}$ is defined similarly to the case of COPE. Given the security of COPA and XLS, it is straightforward to verify that this extension is also secure.