

Blockcipher-based Authenticated Encryption: How Small Can We Go?

Avik Chakraborti¹, Tetsu Iwata², Kazuhiko Minematsu³, and Mridul Nandi⁴

¹ NTT Secure Platform Laboratories, Japan, chakraborti.avik@lab.ntt.co.jp

² Nagoya University, Japan, iwata@cse.nagoya-u.ac.jp

³ NEC Corporation, Japan, k-minematsu@ah.jp.nec.com

⁴ Applied Statistics Unit, Indian Statistical Institute, Kolkata,
mridul.nandi@gmail.com

Abstract. This paper presents a design of authenticated encryption (AE) focusing on minimizing the implementation size, i.e., hardware gates or working memory on software. The scheme is called COFB, for COmbined FeedBack. COFB uses an n -bit blockcipher as the underlying primitive, and relies on the use of a nonce for security. In addition to the state required for executing the underlying blockcipher, COFB needs only $n/2$ bits state as a mask. Till date, for all existing constructions in which masks have been applied, at least n bit masks have been used. Thus, we have shown the possibility of reducing the size of a mask without degrading the security level much. Moreover, it requires one blockcipher call to process one input block. We show COFB is provably secure up to $O(2^{n/2}/n)$ queries which is almost up to the standard birthday bound. We also present our hardware implementation results. Experimental implementation results suggest that our proposal has a good performance and the smallest footprint among all known blockcipher-based AE.

Keywords: COFB, AES, authenticated encryption, blockcipher.

1 Introduction

Authenticated encryption (AE) is a symmetric-key cryptographic primitive for providing both confidentiality and authenticity. Due to the recent rise in communication networks operated on small devices, the era of the so-called Internet of Things, AE is expected to play a key role in securing these networks.

In this paper, we study blockcipher modes for AE with primary focus on the hardware implementation size. Here, we consider the overhead in size, thus the state memory size beyond the underlying blockcipher itself (including the key schedule) is the criteria we want to minimize. We observe this direction has not received much attention until the launch of CAESAR competition (see below), while it would be relevant for future communication devices requiring ultra low-power operations. A general approach to reduce the entire hardware size of AE modes is to use a lightweight blockcipher [15, 17, 25, 48, 49] or to use standard AES implemented in a tiny, serialized core [37], where the latter is shown to be

effective for various schemes including popular CCM [5] or OCB [32] modes, as shown in [16] and [12]. Our approach is orthogonal to these directions.

In this paper, we propose a new blockcipher AE mode which utilizes both plaintext and ciphertext feedback. Our proposal is called COFB for COmbined FeedBack, and we show that this enables essentially AE using the minimum amount of state memory while keeping the security level similar to the previous schemes. Specifically, let n denote the block size in bits of the underlying blockcipher, then our proposal needs an $n/2$ -bit register for a mask in addition to the registers required for holding round keys and the internal state memory (i.e., n bits) for the blockcipher computation. Ignoring the state for the round keys, it requires $1.5n$ bit state. It has provable security up to $O(2^{n/2}/n)$ queries, based on the standard assumption that the blockcipher is a PRP (PseudoRandom Permutation). Our scheme is efficient in that the rate is 1, i.e., it makes one blockcipher call to process one input block, meaning that it is as fast as encryption-only modes.

CAESAR [3], started in 2012, attracted 57 AE schemes, and there are schemes that were designed to minimize the implementation size. The most relevant one is JAMBU [52], which can be implemented with $1.5n$ -bit state memory. However, the provable security result is not published for this scheme[‡], and the security claim about the confidentiality in the nonce misuse scenario was shown to be flawed [40]. We also point out that the rate of JAMBU is $1/2$, i.e., it makes two blockcipher calls to process one input block. This can be seen in our implementation results where COFB is more efficient, in terms of throughput per area, than JAMBU by a factor of two. CLOC and SILC [28, 29] have provable security results and were designed to minimize the implementation size, however, they do not allow the implementation with $1.5n$ -bit state and the rate is also $1/2$.

On the downside, COFB is completely serial both for encryption and decryption. However, we argue that this is a reasonable trade-off, as tiny devices are our primal target platform for COFB. We present Table 1 to show a comparison of blockcipher AE modes including COFB.

In order to instantiate our efficiency claim, we implemented COFB on hardware and evaluated it on FPGAs. The implementation results show the impressive performance figures of COFB both for size and speed. For the sake of completeness we also compare COFB with various schemes (not limited to blockcipher modes) listed in the hardware benchmark framework called ATHENa [1]. We have to warn that this is a rough comparison ignoring differences in several implementation factors (see Sect. 6). Nevertheless, we think this comparison implies a good performance of COFB among others even using the standard AES-128.

2 Preliminaries

Notation. We fix a positive integer n which is the block size in bits of the underlying blockcipher E_K . Typically, we consider $n = 128$ and AES-128 [7] is

[‡] The authenticity result was briefly presented in the latest specification [52].

Table 1. Comparison of AE modes, using an n -bit blockcipher with k -bit keys. An inverse-free mode is a mode that does not need the blockcipher inverse (decryption) function for both encryption and decryption. For JAMBU, the authenticity bound was briefly presented in [52].

Scheme	State Size	Rate	Parallel	Inverse-Free	Sec. Proof	Ref
COFB	$1.5n + k$	1	No	Yes	Yes	This work
JAMBU	$1.5n + k$	1/2	No	Yes	Partial	[52]
CLOC/ SILC	$2n + k$	1/2	No	Yes	Yes	[28, 29]
iFEED	$3n + k$	1	Only for Enc	Yes	Flawed [47]	[54]
OCB	$\geq 3n + k$	1	Yes	No	Yes	[32, 41, 42]

the underlying blockcipher, where K is the 128-bit AES key. The empty string is denoted by λ . For any $X \in \{0, 1\}^*$, where $\{0, 1\}^*$ is the set of all finite bit strings (including λ), we denote the number of bits of X by $|X|$. Note that $|\lambda| = 0$. For two bit strings X and Y , $X\|Y$ denotes the concatenation of X and Y . A bit string X is called a *complete* (or *incomplete*) block if $|X| = n$ (or $|X| < n$ respectively). We write the set of all complete (or incomplete) blocks as \mathcal{B} (or $\mathcal{B}^<$ respectively). Let $\mathcal{B}^{\leq} = \mathcal{B}^< \cup \mathcal{B}$ denote the set of all blocks. For $B \in \mathcal{B}^{\leq}$, we define \bar{B} as follows:

$$\bar{B} = \begin{cases} 0^n & \text{if } B = \lambda \\ B\|10^{n-1-|B|} & \text{if } B \neq \lambda \text{ and } |B| < n \\ B & \text{if } |B| = n \end{cases}$$

Given $Z \in \{0, 1\}^*$, we define the parsing of Z into n -bit blocks as

$$(Z[1], Z[2], \dots, Z[z]) \stackrel{n}{\leftarrow} Z, \quad (1)$$

where $z = \lceil |Z|/n \rceil$, $|Z[i]| = n$ for all $i < z$ and $1 \leq |Z[z]| \leq n$ such that $Z = (Z[1] \| Z[2] \| \dots \| Z[z])$. If $Z = \lambda$, we let $z = 1$ and $Z[1] = \lambda$. We write $\|Z\| = z$ (number of blocks present in Z). We similarly write $(Z[1], Z[2], \dots, Z[z]) \stackrel{m}{\leftarrow} Z$ to denote the parsing of the bit string Z into m bit strings $Z[1], Z[2], \dots, Z[z-1]$ and $1 \leq |Z[z]| \leq m$. Given any sequence $Z = (Z[1], \dots, Z[s])$ and $1 \leq a \leq b \leq s$, we represent the subsequence $(Z[a], \dots, Z[b])$ by $Z[a..b]$. Similarly, for integers $a \leq b$, we write $[a..b]$ for the set $\{a, a+1, \dots, b\}$. For two bit strings X and Y with $|X| \geq |Y|$, we define the extended xor-operation as

$$\begin{aligned} X \oplus Y &= X[1..|Y|] \oplus Y \text{ and} \\ X \bar{\oplus} Y &= X \oplus (Y\|0^{|X|-|Y|}), \end{aligned}$$

where $(X[1], X[2], \dots, X[x]) \stackrel{1}{\leftarrow} X$ and thus $X[1..|Y|]$ denotes the first $|Y|$ bits of X . When $|X| = |Y|$, both operations reduce to the standard $X \oplus Y$.

Let $\gamma = (\gamma[1], \dots, \gamma[s])$ be a tuple of equal-length strings. We define $\text{mcoll}(\gamma) = r$ if there exist distinct $i_1, \dots, i_r \in [1..s]$ such that $\gamma[i_1] = \dots = \gamma[i_r]$ and r is the maximum of such integer. We say that $\{i_1, \dots, i_r\}$ is an r -multi-collision set for γ .

Authenticated Encryption and Security Definitions. An authenticated encryption (AE) is an integrated scheme that provides both privacy of a plaintext $M \in \{0, 1\}^*$ and authenticity of M as well as associate data $A \in \{0, 1\}^*$. Taking a nonce N (which is a value never repeats at encryption) together with associated data A and plaintext M , the encryption function of AE, \mathcal{E}_K , produces a tagged-ciphertext (C, T) where $|C| = |M|$ and $|T| = t$. Typically, t is a fixed length and we assume $n = t$ throughout the paper. The corresponding decryption function, \mathcal{D}_K , takes (N, A, C, T) and returns a decrypted plaintext M when the verification on (N, A, C, T) is successful, otherwise returns the atomic error symbol denoted by \perp .

Privacy. Given an adversary \mathcal{A} , we define the *PRF-advantage* of \mathcal{A} against \mathcal{E} as $\mathbf{Adv}_{\mathcal{E}}^{\text{prf}}(\mathcal{A}) = |\Pr[\mathcal{A}^{\mathcal{E}_K} = 1] - \Pr[\mathcal{A}^{\$} = 1]|$, where $\$$ returns a random string of the same length as the output length of \mathcal{E}_K , by assuming that the output length of \mathcal{E}_K is uniquely determined by the query. The PRF-advantage of \mathcal{E} is defined as

$$\mathbf{Adv}_{\mathcal{E}}^{\text{prf}}(q, \sigma, t) = \max_{\mathcal{A}} \mathbf{Adv}_{\mathcal{E}}^{\text{prf}}(\mathcal{A}),$$

where the maximum is taken over all adversaries running in time t and making q queries with the total number of blocks in all the queries being at most σ . If \mathcal{E}_K is an encryption function of AE, we call it the *privacy advantage* and write as $\mathbf{Adv}_{\mathcal{E}}^{\text{priv}}(q, \sigma, t)$, as the maximum of all nonce-respecting adversaries (that is, the adversary can arbitrarily choose nonces provided all nonce values in the encryption queries are distinct).

Authenticity. We say that an adversary \mathcal{A} *forges* an AE scheme $(\mathcal{E}, \mathcal{D})$ if \mathcal{A} is able to compute a tuple (N, A, C, T) satisfying $\mathcal{D}_K(N, A, C, T) \neq \perp$, without querying (N, A, M) for some M to \mathcal{E}_K and receiving (C, T) , i.e. (N, A, C, T) is a non-trivial forgery.

In general, a forger can make q_f forging attempts without restriction on N in the decryption queries, that is, N can be repeated in the decryption queries and an encryption query and a decryption query can use the same N . The *forging advantage* for an adversary \mathcal{A} is written as $\mathbf{Adv}_{\mathcal{E}}^{\text{auth}}(\mathcal{A}) = \Pr[\mathcal{A}^{\mathcal{E}} \text{ forges}]$, and we write

$$\mathbf{Adv}_{\mathcal{E}}^{\text{auth}}((q, q_f), (\sigma, \sigma_f), t) = \max_{\mathcal{A}} \mathbf{Adv}_{\mathcal{E}}^{\text{auth}}(\mathcal{A})$$

to denote the maximum forging advantage for all adversaries running in time t , making q encryption and q_f decryption queries with total number of queried blocks being at most σ and σ_f , respectively.

Unified Security Notion for AE. The privacy and authenticity advantages can be unified into a single security notion as introduced in [23, 43]. Let \mathcal{A} be an adversary that only makes non-repeating queries to \mathcal{D}_K . Then, we define the AE-advantage of \mathcal{A} against \mathcal{E} as

$$\mathbf{Adv}_{\mathcal{E}}^{\text{AE}}(\mathcal{A}) = |\Pr[\mathcal{A}^{\mathcal{E}_K, \mathcal{D}_K} = 1] - \Pr[\mathcal{A}^{\$, \perp} = 1]|,$$

where \perp -oracle always returns \perp and $\$$ -oracle is as the privacy advantage. We similarly define $\mathbf{Adv}_{\mathcal{E}}^{\text{AE}}((q, q_f), (\sigma, \sigma_f), t) = \max_{\mathcal{A}} \mathbf{Adv}_{\mathcal{E}}^{\text{AE}}(\mathcal{A})$, where the maximum is taken over all adversaries running in time t , making q encryption and

q_f decryption queries with the total number of blocks being at most σ and σ_f , respectively.

Blockcipher Security. We use a blockcipher E as the underlying primitive, and we assume the security of E as a PRP (pseudorandom permutation). The *PRP-advantage* of a blockcipher E is defined as $\mathbf{Adv}_E^{\text{PRP}}(\mathcal{A}) = |\Pr[\mathcal{A}^{E_K} = 1] - \Pr[\mathcal{A}^{\mathbf{P}} = 1]|$, where \mathbf{P} is a random permutation uniformly distributed over all permutations over $\{0, 1\}^n$. We write

$$\mathbf{Adv}_E^{\text{PRP}}(q, t) = \max_{\mathcal{A}} \mathbf{Adv}_E^{\text{PRP}}(\mathcal{A}),$$

where the maximum is taken over all adversaries running in time t and making q queries. Here, σ does not appear as each query has a fixed length.

3 Combined Feedback Mode

Let E_K be the underlying primitive, a blockcipher, with key K . Depending on how the next input block of E_K is determined from the previous output of E_K , a plaintext block, or a ciphertext block, we can categorize different types of feedback modes. Some of the feedback modes are illustrated in Fig. 3.1. The first three modes are known as the *message feedback mode*, *ciphertext feedback mode*, and *output feedback mode*, respectively. The examples using the first three modes can be found in the basic encryption schemes [4] or AE schemes [5, 28, 29, 54]. The fourth mode, which uses additional (linear) operation $G : \mathcal{B} \rightarrow \mathcal{B}$, is new. We call it *combined feedback*. In the combined feedback mode, the next input block $X[i]$ of the underlying primitive E_K depends on at least two of the following three values: (i) previous output $E_K(X[i-1])$, (ii) plaintext $M[i]$, and (iii) ciphertext $C[i]$. With an appropriate choice of G , this feedback mode turns out to be useful for building small and efficient AE schemes. We provide a unified presentation of all types of feedback functions below.

Definition 1 (Feedback Function). A function $\rho : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B} \times \mathcal{B}$ is called a *feedback function* (for an encryption) if there exists a function $\rho' : \mathcal{B} \times \mathcal{B} \rightarrow \mathcal{B} \times \mathcal{B}$ (used for decryption) such that

$$\forall Y, M \in \mathcal{B}, \quad \rho(Y, M) = (X, C) \Rightarrow \rho'(Y, C) = (X, M). \quad (2)$$

ρ is called a *plaintext or output feedback* if X depends only on M or Y , respectively (e.g., the first and third mode in Fig. 3.1). Similarly, it is called *ciphertext feedback* if X depends only on C in the function ρ' (e.g., the second mode in Fig. 3.1). All other feedback functions are called *combined feedback*.

The condition stated in Eq. (2) is sufficient for inverting the feedback computation from the ciphertext. Given the previous output block $Y = E_K(X[i-1])$ and a ciphertext block $C = C[i-1]$, we are able to compute $(X, M) = (X[i], M[i])$ by using $\rho'(Y, C)$.

In particular, when G is not the zero function nor the identity function, the combined feedback mode using this G is not reduced to the remaining three modes. It can be described as $\rho(Y, M) = (X, C) = (G(Y) \oplus M, Y \oplus M)$.

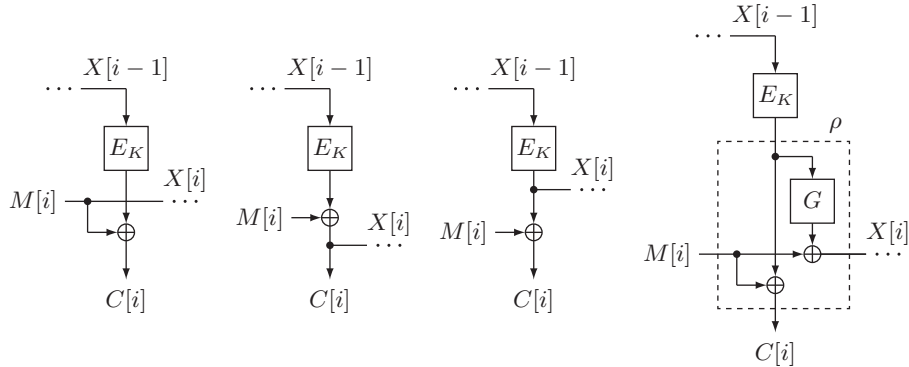


Fig. 3.1. Different types of feedback modes. We introduce the last feedback mode (called the combined feedback mode) in our construction.

4 COFB: a Small-State, Rate-1, Inverse-Free AE Mode

In this section, we present our proposal, COFB, which has rate-1 (i.e. needs one blockcipher call for one input block), and is inverse-free, i.e., it does not need a blockcipher inverse (decryption). In addition to these features, this mode has a quite small state size, namely $1.5n + k$ bits, in case the underlying blockcipher has an n -bit block and k -bit keys. We first specify the basic building blocks and parameters used in our construction.

Key and Blockcipher. The underlying cryptographic primitive is an n -bit blockcipher, E_K . We assume that n is a multiple of 4. The key of the scheme is the key of the blockcipher, i.e. K . As mentioned we typically assume that E_K is AES-128 with $n = k = 128$, however, COFB can be instantiated with any blockcipher of any n -bit block size by appropriately defining other components.

Masking Function. We define the masking function $\text{mask} : \{0, 1\}^{n/2} \times \mathbb{N}^2 \rightarrow \{0, 1\}^{n/2}$ as follows:

$$\text{mask}(\Delta, a, b) = \alpha^a \cdot (1 + \alpha)^b \cdot \Delta \quad (3)$$

We may write $\text{mask}_\Delta(a, b)$ to mean $\text{mask}(\Delta, a, b)$. Here, \cdot denotes the multiplication over $\text{GF}(2^{n/2})$, and α denotes the primitive element of the field. For the primitive polynomial defining the field, we choose the lexicographically first one, that is, $p(x) = x^{64} + x^4 + x^3 + x + 1$ following [6, 27]. Rogaway [41] showed that for all $(a, b) \in \{0, \dots, 2^{51}\} \times \{0, \dots, 2^{10}\}$, the values of $\alpha^a \cdot (1 + \alpha)^b$ are distinct[§]. For other values of n , we need to identify the primitive element α of the primitive polynomial and an integer L such that $\alpha^a \cdot (1 + \alpha)^b$ are distinct

[§] If we follow the notations of [41], the right hand side of Eq. (3) could be written as $2^a 3^b \Delta$.




Fig. 4.1. Encryption of COFB for 3-block associated data and plaintext.

for all $(a, b) \in \{0, \dots, L\} \times \{0, \dots, 4\}$. Then the total allowed size of a message and associated data would be at most nL bits. We need this condition to prove the security claim. In particular we have the following properties of the masking function.

Lemma 1. *For any $(a, b) \neq (a', b')$ chosen from the set $\{0, \dots, L\} \times \{0, \dots, 4\}$ (as described above), $c \in \{0, 1\}^{n/2}$ and a random $n/2$ bit string Δ , we have*

$$\Pr[\text{mask}_\Delta(a, b) \oplus \text{mask}_\Delta(a', b') = c] = \frac{1}{2^{n/2}}, \text{ and } \Pr[\text{mask}_\Delta(a, b) = c] = \frac{1}{2^{n/2}}.$$

Proof of the first equation trivially follows from the fact that $\alpha^a \cdot (1 + \alpha)^b$ are distinct for all $(a, b) \in \{0, \dots, L\} \times \{0, \dots, 4\}$.

Similar masking functions are frequently used in other modes, such as [9, 35, 41], however, the masks are full n bits. The use of n -bit masking function usually allows to redefine the AE scheme as a mode of XE or XEX tweakable blockcipher [41], which significantly reduces the proof complexity. In our case, to reduce the state size, we decided to use the $n/2$ -bit masking function, and as a result the proof is ad-hoc and does not rely on XE or XEX.

Feedback Function. Let $Y \in \{0, 1\}^n$ and $(Y[1], Y[2], Y[3], Y[4]) \stackrel{n/4}{\leftarrow} Y$, where $Y[i] \in \{0, 1\}^{n/4}$. We define $G : \mathcal{B} \rightarrow \mathcal{B}$ as $G(Y) = (Y[2], Y[3], Y[4], Y[4] \oplus Y[1])$. We also view G as the $n \times n$ non-singular matrix, so we write $G(Y)$ and $G \cdot Y$ interchangeably. For $M \in \mathcal{B}^{\leq}$ and $Y \in \mathcal{B}$, we define $\rho_1(Y, M) = G \cdot Y \oplus \overline{M}$. The feedback function ρ and its corresponding ρ' are defined as

$$\begin{aligned}\rho(Y, M) &= (\rho_1(Y, M), Y \underline{\oplus} M), \\ \rho'(Y, C) &= (\rho_1(Y, Y \underline{\oplus} C), Y \underline{\oplus} C).\end{aligned}$$

Note that when $(X, M) = \rho(Y, C)$ then $X = (G \oplus I)Y \underline{\oplus} C$. Our choice of G ensures that $I \oplus G$ is also invertible matrix. So when Y is chosen randomly for both computations of X (through ρ and ρ'), X also behaves randomly. We need this property when we bound probability of bad events later.

Tweak Value for The Last Block. Given $B \in \{0, 1\}^*$, we define $\delta_B \in \{1, 2\}$ as follows:

$$\delta_B = \begin{cases} 1 & \text{if } B \neq \lambda \text{ and } n \text{ divides } |B| \\ 2 & \text{otherwise.} \end{cases} \quad (4)$$

This will be used to differentiate the cases that the last block of B is n bits or shorter, for B being associated data or plaintext or ciphertext. We also define a formatting function Fmt for a pair of bit strings (A, Z) , where A is associated data and Z could be either a message or a ciphertext. Let $(A[1], \dots, A[a]) \stackrel{n}{\leftarrow} A$ and $(Z[1], \dots, Z[z]) \stackrel{n}{\leftarrow} Z$. We define $\mathfrak{t}[i]$ as follows:

$$\mathfrak{t}[i] = \begin{cases} (i, 0) & \text{if } i < a \\ (a - 1, \delta_A) & \text{if } i = a \\ (i - 1, \delta_A) & \text{if } a < i < a + z \\ (a + z - 2, \delta_A + \delta_Z) & \text{if } i = a + z \end{cases}$$

Now, the formatting function $\text{Fmt}(A, Z)$ returns the following sequence:

$$((A[1], \mathfrak{t}[1]), \dots, (\overline{A[a]}, \mathfrak{t}[a]), (Z[1], \mathfrak{t}[a + 1]), \dots, (\overline{Z[z]}, \mathfrak{t}[a + z])),$$

where the first coordinate of each pair specifies the input block to be processed, and the second coordinate specifies the exponents of α and $1 + \alpha$ to determine the constant over $\text{GF}(2^{n/2})$. Let $\mathbb{Z}_{\geq 0}$ be the set of non-negative integers and \mathcal{X} be some non-empty set. We say that a function $f : \mathcal{X} \rightarrow (\mathcal{B} \times \mathbb{Z}_{\geq 0} \times \mathbb{Z}_{\geq 0})^+$ is *prefix-free* if for all $X \neq X'$, $f(X) = (Y[1], \dots, Y[\ell])$ is not a prefix of $f(X') = (Y'[1], \dots, Y'[\ell'])$ (in other words, $(Y[1], \dots, Y[\ell]) \neq (Y'[1], \dots, Y'[\ell])$). Here, for a set \mathcal{S} , \mathcal{S}^+ means $\mathcal{S} \cup \mathcal{S}^2 \cup \dots$, and we have the following lemma.

Lemma 2. *The function $\text{Fmt}(\cdot)$ is prefix-free.*

The proof is more or less straightforward and hence we skip it.

<p>Module Mask-Gen(K, N)</p> <ol style="list-style-type: none"> 1. $Y[0] \leftarrow E_K(0^{n/2} \parallel N)$ 2. $(Y^1[0], \dots, Y^4[0]) \xleftarrow{n/4} Y[0]$ 3. $\Delta \leftarrow Y^2[0] \parallel Y^3[0]$ 4. return $(\Delta, Y[0])$ <p>Algorithm COFB-$\mathcal{E}_K(N, A, M)$</p> <ol style="list-style-type: none"> 1. $(\Delta, Y[0]) \leftarrow \text{Mask-Gen}(K, N)$ 2. $(A[1], \dots, A[a]) \xleftarrow{n} A$ 3. $(M[1], \dots, M[m]) \xleftarrow{n} M$ 4. $\ell \leftarrow a + m$ 5. $((B[1], t[1]), \dots, (B[\ell], t[\ell])) \leftarrow \text{Fmt}(A, M)$ 6. for $i = 1$ to ℓ 7. $X[i] \leftarrow (B[i] \oplus G \cdot Y[i-1]) \oplus \text{mask}_\Delta(t[i])$ 8. $Y[i] \leftarrow E_K(X[i])$ 9. if $i > a$ then 10. $C[i-a] \leftarrow Y[i-1] \oplus M[i-a]$ 11. $T \leftarrow Y[\ell]$ 12. return (C, T) 	<p>Algorithm COFB-$\mathcal{D}_K(N, A, C, T)$</p> <ol style="list-style-type: none"> 1. $(\Delta, Y[0]) \leftarrow \text{Mask-Gen}(K, N)$ 2. $(A[1], \dots, A[a]) \xleftarrow{n} A$ 3. $(C[1], \dots, C[c]) \xleftarrow{n} C$ 4. $\ell \leftarrow a + c$ 5. $((B[1], t[1]), \dots, (B[\ell], t[\ell])) \leftarrow \text{Fmt}(A, C)$ 6. for $i = 1$ to ℓ 7. if $i \leq a$ then 8. $X[i] \leftarrow (B[i] \oplus G \cdot Y[i-1]) \oplus \text{mask}_\Delta(t[i])$ 9. else $X[i] \leftarrow (B[i] \oplus Y[i-1] \oplus G \cdot Y[i-1]) \oplus \text{mask}_\Delta(t[i])$ 10. $Y[i] \leftarrow E_K(X[i])$ 11. for $i = 1$ to c 12. $M[i] \leftarrow Y[i+a-1] \oplus C[i]$ 13. $M \leftarrow (M[1], \dots, M[c])$ 14. $T' \leftarrow Y[\ell]$ 15. if $T' = T$ then return M 16. else return \perp
--	---

Fig. 4.2. The encryption and decryption algorithms of COFB.

Now we present the specifications of COFB in Fig. 4.2. The encryption and decryption algorithms are denoted by COFB- \mathcal{E}_K and COFB- \mathcal{D}_K . We remark that the nonce length is $n/2$ bits, which is enough for the security up to the birthday bound. The nonce is processed as $E_K(0^{n/2} \parallel N)$ to yield the first internal chaining value. The encryption algorithm takes non-empty A and non-empty M , and outputs C and T such that $|C| = |M|$ and $|T| = n$. The decryption algorithm takes (N, A, C, T) with $|A|, |C| \neq 0$ and outputs M or \perp .

We remark that the pseudocodes of Fig. 4.2 are for clarity and not necessarily memory-efficient due to (say) the use of Fmt and caching multiple $Y[i]$ values at decryption. In fact, the encryption and decryption of COFB can be done with keeping one input/output state for the blockcipher and sequentially updating the $n/2$ -bit mask. See Fig. 4.1 for an illustration.

5 Security of COFB

We present the security analysis of COFB in Theorem 1.

Theorem 1 (Main Theorem).

$$\begin{aligned} \text{Adv}_{\text{COFB}}^{\text{AE}}((q, q_f), (\sigma, \sigma_f), t) &\leq \text{Adv}_{\text{AES}}^{\text{prp}}(q', t') + \frac{0.5(q')^2}{2^n} + \frac{4\sigma + 0.5nq_f}{2^{n/2}} \\ &\quad + \frac{q_f + (q + \sigma + \sigma_f) \cdot \sigma_f}{2^n} \end{aligned}$$

where, $q' = q + q_f + \sigma + \sigma_f$, which corresponds to the total number of blockcipher calls through the game, and $t' = t + O(q')$.

Proof. Without loss of generality, we can assume $q' \leq 2^{\frac{n}{2}-1}$, since otherwise the bound obviously holds as the right hand side becomes more than one. The first transition we make is to use an n -bit (uniform) random permutation P instead of E_K , and then to use an n -bit (uniform) random function R instead of P . This two-step transition requires the first two terms of our bound, from the standard PRP-PRF switching lemma and from the computation to the information security reduction (e.g., see [13]). Then what we need is a bound for COFB using R , denoted by COFB-R. That is, we prove

$$\text{Adv}_{\text{COFB-R}}^{\text{AE}}((q, q_f), (\sigma, \sigma_f), \infty) \leq \frac{4\sigma + 0.5nq_f}{2^{n/2}} + \frac{q_f + (q + \sigma + \sigma_f) \cdot \sigma_f}{2^n}. \quad (5)$$

For $i = 1, \dots, q$, we write (N_i, A_i, M_i) and (C_i, T_i) to denote the i -th encryption query and response. Here, $A_i = (A_i[1], \dots, A_i[a_i])$, $M_i = (M_i[1], \dots, M_i[m_i])$, and $C_i = (C_i[1], \dots, C_i[m_i])$. Let $\ell_i = a_i + m_i$, which denotes the total input block length for the i -th encryption query. We write $X_i[j]$ (resp. $Y_i[j]$) for $i = 1, \dots, q$ and $j = 0, \dots, \ell_i$ to denote the j -th input (resp. output) of the internal R invoked at the i -th encryption query, where the order of invocation follows the specification shown in Fig. 4.2. We remark that $X_i[0] = 0^{n/2} \parallel N_i$ and $Y_i[\ell_i] = T_i$ for all $i = 1, \dots, q$. Similarly, we write Δ_i to denote $Y_i^2[0] \parallel Y_i^3[0]$ where $Y_i^1[0] \parallel \dots \parallel Y_i^4[0] \xleftarrow{n/4} Y_i[0]$.

We introduce the following relaxations in the game, which only gain the advantage. First, after completing all queries and forging attempts (i.e. decryption queries), let the adversary learn all the Y -values for all encryption queries only. We remark that any X -values computed at the message processing phase (not the AD processing phase) of the i -th encryption query are immediately determined by the i -th query-response tuple, $(N_i, A_i, M_i, C_i, T_i)$ and Y_i values from the property of feedback function, and Δ -values (it is a part of $Y[0]$).

In case of the ideal oracle, all these variables corresponding to Y will be chosen uniformly and independently, where at the plaintext encryption phase $Y_i[j]$ is randomly chosen and used to determine $C_i[j]$ as $C_i[j] = Y_i[j-1] \oplus M_i[j]$, and at AD processing phase it is a dummy and has no influence to the response (C_i, T_i) . For decryption queries, the ideal oracle always returns \perp (here we assume that the adversary makes only fresh queries).

Coefficients-H Technique. We outline the Coefficients-H technique developed by Patarin, which serves as a convenient tool for bounding the advantage (see [39, 50]). We will use this technique (without giving a proof) to prove our main theorem. Consider two oracles $\mathcal{O}_0 = (\$, \perp)$ (the ideal oracle for the relaxed game) and \mathcal{O}_1 (real, i.e. our construction in the same relaxed game). Let \mathcal{V} denote the set of all possible views an adversary can obtain. For any view $\tau \in \mathcal{V}$, we will denote the probability to realize the view as $\text{ip}_{\text{real}}(\tau)$ (or $\text{ip}_{\text{ideal}}(\tau)$) when it is interacting with the real (or ideal respectively) oracle. We call these *interpolation probabilities*. Without loss of generality, we assume that the adversary is deterministic and fixed. Then, the probability space for the interpolation probabilities is uniquely determined by the underlying oracle. As we deal with stateless oracles, these probabilities are independent of the order of query responses in the

view. Suppose we have a set of views, $\mathcal{V}_{\text{good}} \subseteq \mathcal{V}$, which we call *good* views, and the following conditions hold:

1. In the game involving the ideal oracle \mathcal{O}_0 (and the fixed adversary), the probability of getting a view in $\mathcal{V}_{\text{good}}$ is at least $1 - \epsilon_1$.
2. For any view $\tau \in \mathcal{V}_{\text{good}}$, we have $\text{ip}_{\text{real}}(\tau) \geq (1 - \epsilon_2) \cdot \text{ip}_{\text{ideal}}(\tau)$.

Then we have $|\Pr[\mathcal{A}^{\mathcal{O}_0} = 1] - \Pr[\mathcal{A}^{\mathcal{O}_1} = 1]| \leq \epsilon_1 + \epsilon_2$. The proof can be found at (say) [50]. Now we proceed with the proof of Theorem 1 by defining certain $\mathcal{V}_{\text{good}}$ for our games, and evaluating the bounds, ϵ_1 and ϵ_2 .

Views. In our case, a view τ is defined by the following tuple:

$$\tau = ((N_i, A_i, M_i, Y_i)_{i \in \{1, \dots, q\}}, (N_{i'}, A_{i'}, C_{i'}, T_{i'}, Z_{i'})_{i' \in \{1, \dots, q_f\}}),$$

where $Z_{i'}$ denotes the output of the decryption oracle \mathcal{D} (it is always \perp when we interact with the ideal oracle) for the i' -th decryption query $(N_{i'}, A_{i'}, C_{i'}, T_{i'})$. Note that Y_i denotes $(Y_i[0], \dots, Y_i[\ell_i]) = Y_i[0..\ell_i]$, where $\ell_i = a_i + m_i$, and a_i (resp. m_i) denotes the block length of A_i (resp. M_i). Here we implicitly use the fact that given a complete block $M_i[j]$, the mapping from $Y_i[j]$ to $C_i[j]$ is bijective and hence keeping those $Y_i[j]$ values instead of $C_i[j]$ is sufficient. Similarly we define $c_{i'}$ and $a_{i'}$, and write $\ell_{i'} = a_{i'} + c_{i'}$.

Let $(L_i[j], R_i[j]) \stackrel{n/2}{\leftarrow} X_i[j]$ for all $i \in [1..q]$ and $j \in [1..\ell_i]$. For any i , let p_i denote the length of the longest common prefix of $\text{Fmt}(A_i^*, C_i^*)$ and $\text{Fmt}(A_j, C_j)$ where $N_j = N_i^*$. If there is no such j , we define $p_i = -1$. Since Fmt is prefix-free, it holds that $p_i < \min\{\ell_i^*, \ell_j\}$. We observe that p_i is unique for all $i = 1, \dots, q_f$, as there is at most one encryption query that uses the same nonce as N_i^* .

Bad Views. Now we define a bad view. The complement of the set of bad views is defined to be the set of good views. A view is called bad if one of the following events occurs:

- B1:** $L_i[j] = 0^{n/2}$ for some $i \in [1..q]$ and $j > 0$.
- B2:** $X_i[j] = X_{i'}[j']$ for some $(i, j) \neq (i', j')$ where $j, j' > 0$.
- B3:** $\text{mcoll}(R) > n/2$, where R is the tuple of all $R_i[j]$ values. Recall that $(L_i[j], R_i[j]) \stackrel{n/2}{\leftarrow} X_i[j]$.
- B4:** $X_i^*[p_i + 1] = X_{i_1}[j_1]$ for some i, i_1, j_1 with p_i as defined above. Note that when $p_i \geq 0$, $X_i^*[p_i + 1]$ is determined from the values of Y .
- B5:** For some $Z_{i'} \neq \perp$. This clearly cannot happen for the ideal oracle case.

We add some intuitions on these events. When **B1** does not hold, then $X_i[j] \neq X_{i'}[0]$ for all i, i' , and $j > 0$. Hence Δ_i will be completely random. When **B2** does not hold, then all the inputs for the random function are distinct for encryption queries, which makes the responses from encryption oracle completely random in the “real” game. When **B3** does not hold, then at the right half of $X_i[j]$ we see at most $n/2$ multi-collisions. A successful forgery is to choose one of the $n/2$ multi-collision blocks and forge the left part so that the entire block collides. Forging

the left part has $2^{-n/2}$ probability due to randomness of masking. Finally, when **B4** does not hold, then the $(p_i + 1)$ -st input for the i -th forging attempt will be fresh with a high probability and so all the subsequent inputs will remain fresh with a high probability.

A view is called good if none of the above events hold. Let $\mathcal{V}_{\text{good}}$ be the set of all such good views. The following lemma bounds the probability of not realizing a good view while interacting with a random function (this will complete the first condition of the Coefficients-H technique).

Lemma 3.

$$\Pr_{\text{ideal}}[\tau \notin \mathcal{V}_{\text{good}}] \leq \frac{4\sigma + 0.5nq_f}{2^{n/2}}.$$

Proof (of Lemma 3). Throughout the proof, we assume all probability notations are defined over the ideal game. We bound all the bad events individually and then by using the union bound, we will obtain the final bound. We first develop some more notation. Let $(Y_i^1[j], Y_i^2[j], Y_i^3[j], Y_i^4[j]) \stackrel{n/4}{\leftarrow} Y_i[j]$. Similarly, we denote $(M_i^1[j], M_i^2[j]) \stackrel{n/2}{\leftarrow} M_i[j]$.

- (1) $\Pr[\mathbf{B1}] \leq \sigma/2^{n/2}$: We fix a pair of integers (i, j) for some $i \in [1..q]$ and $j \in [1..\ell_i]$. Now, $L_i[j]$ can be expressed as

$$(Y_i^2[j-1] \parallel Y_i^3[j-1]) \oplus (\alpha^a \cdot (1 + \alpha)^b \cdot \Delta_i) \oplus M_i^1[j]$$

for some a and b . Note that when $j > 1$, Δ_i and $Y_i[j-1]$ are independently and uniformly distributed, and hence for those j , we have $\Pr[L_i[j] = 0^{n/2}] = 2^{-n/2}$ (apply Lemma 1 after conditioning $Y_i[j-1]$). Now when $j = 1$, we have the following three possible choice: (i) $L_i[1] = (1 + \alpha) \cdot \Delta_i \oplus \text{Cons}$ if $a_i \geq 2$, (ii) $L_i[1] = \alpha \cdot \Delta_i \oplus \text{Cons}$ if $a_i = 1$ and the associated data block is full, and (iii) $L_i[1] = \alpha^2 \cdot \Delta_i \oplus \text{Cons}$ if $a_i = 1$ and the associated data block is not full, for some constant Cons . In all cases by applying Lemma 1, $\Pr[\mathbf{B1}] \leq \sigma/2^{n/2}$.

- (2) $\Pr[\mathbf{B2}] \leq \sigma/2^{n/2}$: For any $(i, j) \neq (i', j')$ with $j, j' \geq 1$, the equality event $X_i[j] = X_{i'}[j']$ has a probability at most 2^{-n} since this event is a non-trivial linear equation on $Y_i[j-1]$ and $Y_{i'}[j'-1]$ and they are independent to each other. Note that $\sigma^2/2^n \leq \sigma/2^{n/2}$ as we are estimating probabilities.
- (3) $\Pr[\mathbf{B3}] \leq 2\sigma/2^{n/2}$: The event **B3** is a multi-collision event for randomly chosen σ many $n/2$ -bit strings as Y values are mapped in a regular manner (see the feedback function) to R values. From the union bound, we have

$$\Pr[\mathbf{B3}] \leq \binom{\sigma}{n/2} \frac{1}{2^{(n/2) \cdot ((n/2)-1)}} \leq \frac{\sigma^{n/2}}{2^{(n/2) \cdot ((n/2)-1)}} \leq \left(\frac{\sigma}{2^{(n/2)-1}} \right)^{n/2} \leq \frac{2\sigma}{2^{n/2}},$$

where the last inequality follows from the assumption $(\sigma \leq 2^{(n/2)-1})$.

- (4) $\Pr[\mathbf{B4} \wedge \mathbf{B1}^c \wedge \mathbf{B3}^c] \leq 0.5nq_f/2^{n/2}$: We fix some i and want to bound the probability $\Pr[X_i^*[p_i + 1] = X_{i_1}[j_1] \wedge \mathbf{B1}^c \wedge \mathbf{B3}^c]$ for some i_1, j_1 . If $p_i = -1$ (i.e., N_i^* does not appear in encryption queries), then N_i^* is fresh as left

$n/2$ bits of all $X_i[j]$ is non-zero for all $j > 0$ (since we also consider **B1** does not hold). So the probability is zero. Now we consider $p_i \geq 0$. The event **B3**^c implies that at most $n/2$ possible values of (i_1, j_1) are possible for which $X_i^*[p_i + 1] = X_{i_1}[j_1]$ can hold. Fix any such (i_1, j_1) . Now it is sufficient to bound the probability for equality for the left $n/2$ bits. We first consider the case where $j_1 = p_i + 1$. Now from the definition of p_i , $(C_i^*[p_i + 1], \mathbf{t}_i^*[p_i + 1]) \neq (C_{i_1}[p_i + 1], \mathbf{t}_{i_1}[p_i + 1])$. If $\mathbf{t}_i[p_i + 1] = \mathbf{t}_{i_1}[p_i + 1]$ then the bad event cannot hold with probability one. Otherwise, we obtain a non-trivial linear equation in Δ_{i_1} and apply Lemma 1, and we also use the fact that $G + I$ is non singular. A similar argument holds for the other choices of j_1 . Therefore, the probability for the atomic case is at most $2^{-n/2}$, and because we have at most $q_f \cdot n/2$ chances, $\Pr[\mathbf{B4} \wedge \mathbf{B1}^c \wedge \mathbf{B3}^c]$ is at most $(n/2) \cdot q_f \cdot 1/2^{n/2}$.

Summarizing, we have

$$\begin{aligned} \Pr_{\text{ideal}}[\tau \notin \mathcal{V}_{\text{good}}] &\leq \Pr[\mathbf{B1}] + \Pr[\mathbf{B2}] + \Pr[\mathbf{B3}] + \Pr[\mathbf{B4} \wedge \mathbf{B1}^c \wedge \mathbf{B3}^c] \\ &\leq \frac{\sigma}{2^{n/2}} + \frac{\sigma}{2^{n/2}} + \frac{2\sigma}{2^{n/2}} + \frac{0.5nq_f}{2^{n/2}} = \frac{4\sigma + 0.5nq_f}{2^{n/2}}, \end{aligned}$$

which concludes the proof. \square

Lower Bound of $\text{ip}_{\text{real}}(\tau)$. We consider the ratio of $\text{ip}_{\text{real}}(\tau)$ and $\text{ip}_{\text{ideal}}(\tau)$. In this paragraph we assume that all the probability space, except for $\text{ip}_{\text{ideal}}(*)$, is defined over the real game. We fix a good view

$$\tau = ((N_i, A_i, M_i, Y_i)_{i \in \{1, \dots, q\}}, (N_{i'}^*, A_{i'}^*, C_{i'}^*, T_{i'}^*, Z_{i'}^*)_{i' \in \{1, \dots, q_f\}}),$$

where $Z_{i'}^* = \perp$. We separate τ into

$$\tau_e = (N_i, A_i, M_i, Y_i)_{i \in \{1, \dots, q\}} \text{ and } \tau_d = (N_{i'}^*, A_{i'}^*, C_{i'}^*, T_{i'}^*, Z_{i'}^*)_{i' \in \{1, \dots, q_f\}},$$

and we first see that for a good view τ , $\text{ip}_{\text{ideal}}(\tau)$ equals to $1/2^{n(q+\sigma)}$.

Now we consider the real case. Since **B1** and **B2** do not hold with τ , all inputs of the random function inside τ_e are distinct, which implies that the released Y -values are independent and uniformly random. The variables in τ_e are uniquely determined given these Y -values, and there are exactly $q + \sigma$ distinct input-output of R . Therefore, $\Pr[\tau_e]$ is exactly $2^{-n(q+\sigma)}$.

We next evaluate

$$\text{ip}_{\text{real}}(\tau) = \Pr[\tau_e, \tau_d] = \Pr[\tau_e] \cdot \Pr[\tau_d | \tau_e] = \frac{1}{2^{n(q+\sigma)}} \cdot \Pr[\tau_d | \tau_e]. \quad (6)$$

We observe that $\Pr[\tau_d | \tau_e]$ equals to $\Pr[\perp_{\text{all}} | \tau_e]$, where \perp_{all} denotes the event that $Z_i^* = \perp$ for all $i = 1, \dots, q_f$, as other variables in τ_d are determined by τ_e .

Let η denote the event that, for all $i = 1, \dots, q_f$, $X_i^*[j]$ for $p_i < j \leq \ell_i^*$ is not colliding to X -values in τ_e and $X_i^*[j']$ for all $j' \neq j$. For $j = p_i + 1$, the above condition is fulfilled by **B4**, and thus $Y_i^*[p_i + 1]$ is uniformly random, and hence

$X_i^*[p_i + 2]$ is also uniformly random, due to the property of feedback function (here, observe that the mask addition between the chain of $Y_i^*[j]$ to $X_i^*[j + 1]$ does not reduce the randomness).

Now we have $\Pr[\perp_{\text{all}}|\tau_e] = 1 - \Pr[(\perp_{\text{all}})^c|\tau_e]$, and we also have $\Pr[(\perp_{\text{all}})^c|\tau_e] = \Pr[(\perp_{\text{all}})^c, \eta|\tau_e] + \Pr[(\perp_{\text{all}})^c, \eta^c|\tau_e]$. Here, $\Pr[(\perp_{\text{all}})^c, \eta|\tau_e]$ is the probability that at least one T_i^* for some $i = 1, \dots, q_f$ is correct as a guess of $Y_i^*[\ell_i^*]$. Here $Y_i^*[\ell_i^*]$ is completely random from η , hence using the union bound we have

$$\Pr[(\perp_{\text{all}})^c, \eta|\tau_e] \leq \frac{q_f}{2^n}.$$

For $\Pr[(\perp_{\text{all}})^c, \eta^c|\tau_e]$ which is at most $\Pr[\eta^c|\tau_e]$, the above observation suggests that this can be evaluated by counting the number of possible bad pairs (i.e. a pair that a collision inside the pair violates η) among the all X -values in τ_e and all X^* -values in τ_d , as in the same manner to the collision analysis of e.g., CBC-MAC using R. For each i -th decryption query, the number of bad pairs is at most $(q + \sigma + \ell_i^*) \cdot \ell_i^* \leq (q + \sigma + \sigma_f) \cdot \ell_i^*$. Therefore, the total number of bad pairs is $\sum_{1 \leq i \leq q_f} (q + \sigma + \sigma_f) \cdot \ell_i^* \leq (q + \sigma + \sigma_f) \cdot \sigma_f$, and we have

$$\Pr[(\perp_{\text{all}})^c, \eta^c|\tau_e] \leq \frac{(q + \sigma + \sigma_f) \cdot \sigma_f}{2^n}.$$

Combining all, we have

$$\begin{aligned} \text{ip}_{\text{real}}(\tau) &= \frac{1}{2^{n(q+\sigma)}} \cdot \Pr[\tau_d|\tau_e] = \text{ip}_{\text{ideal}}(\tau) \cdot \Pr[\perp_{\text{all}}|\tau_e] \\ &\geq \text{ip}_{\text{ideal}}(\tau) \cdot (1 - (\Pr[(\perp_{\text{all}})^c, \eta|\tau_e] + \Pr[(\perp_{\text{all}})^c, \eta^c|\tau_e])) \\ &\geq \text{ip}_{\text{ideal}}(\tau) \cdot \left(1 - \frac{q_f + (q + \sigma + \sigma_f) \cdot \sigma_f}{2^n}\right). \end{aligned}$$

□

6 Hardware Implementation of COFB

6.1 Overview

COFB primarily aims to achieve a lightweight implementation on small hardware devices. For such devices, the hardware resource for implementing memory is often the dominant factor of the size of entire implementation, and the scalability by parallelizing the internal components is not needed. In this respect, COFB's small state size and completely serial operation is quite desirable.

For implementation aspects, COFB is simple, as it consists of a blockcipher and several basic operations (bitwise XOR, the feedback function, and the constant multiplications over $\text{GF}(2^{n/2})$). Combined with the small state size, this implies that the implementation size of COFB is largely dominated by the underlying blockcipher.

We also provide the number of clock cycles needed to process input bytes, as a conventional way to estimate the speed. Here, COFB taking a -block AD

Table 2. Clock cycles per message byte for COFB using a 128-bit blockcipher.

	Message length (Bytes)										
	16	32	64	128	256	512	1024	2048	4096	16384	32768
cpb	2.93	2.22	1.86	1.68	1.59	1.54	1.52	1.51	1.50	1.50	1.50

(associated data) and an m -block message needs $12(a + m) + 23$ cycles. Table 2 shows the number of average cycles per input message bytes, which we call cycles per byte (cpb), assuming AD has the same length as message and the underlying blockcipher has 128-bit block. That is, the table shows $(12 \cdot 2m + 23)/16m$.

6.2 Hardware Architecture

We describe the hardware implementation of COFB using AES-128. This is a basic implementation without any pipelining, and employs a module architecture. We primary focus on the encryption-only circuit, however, the combined encryption and decryption circuit should have very small amount of overhead thanks to the inverse-freeness (i.e. no AES decryption routine is needed) and simplicity of the mode. Due to the similarity between the associated data and the message processing phase, the same hardware modules are used in the both phases. A single bit switch is used to distinguish between the two types of input data. The main architecture consists of the modules described below. We remark that, there is also a Finite State Machine (FSM) which controls the flow by sending signal to these modules. The FSM has a rather simple structure, and will be described in the full version. Then, the overall hardware architecture is described in Fig. 6.1.

1. **State Registers:** The state registers are used to store the intermediate states after each iteration. We use a 128-bit **State** register to store the 128-bit AES block state, a 64-bit Δ register to store the 64-bit mask applied to each AES input, and a 128-bit **Key** register to store the 128-bit key. The round key of AES is stored in the additional 128-bit register (**Round Key**), however, this is included in the AES module.
2. **AES Round:** AES round function module runs one AES round computation and produces a 128-bit output, using two 128-bit inputs, one from the **State** and the other from (internal) **Round Key** registers. The latter register is initialized by loading the master key, stored in the **Key** register, each time the AES function is invoked. The output of AES module is stored into the **State** register, which is the input for the next round. The entire operation is serial, while the internal round computation and the round key generation run in parallel, and needs 11 cycles to perform full AES-128 encryption.
3. **Feedback Function ρ :** The ρ module is to compute the linear feedback function ρ on the 128-bit data block and the 128-bit intermediate state value (output from the AES computation). The output is a 128-bit ciphertext and a 128-bit intermediate state (to be masked and stored to the **State** register).

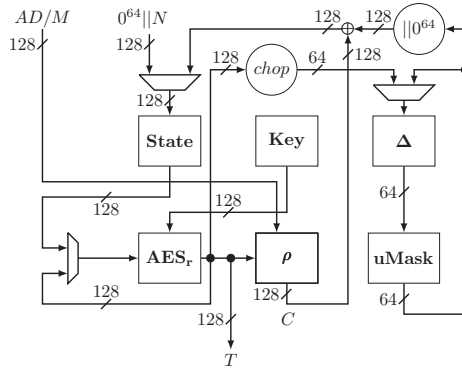


Fig. 6.1. Hardware circuit diagram

4. **Mask Update:** **uMask** module updates the mask stored in Δ register. **uMask** receives the current mask value and updates it by multiplying with α or $(1 + \alpha)$ or $(1 + \alpha)^2$ based on the signals generated by the FSM, where signals are to indicate the end of the message and the completeness of the final block process.

Basic Implementation: We describe a basic flow of our implementation of COFB, which generally follows the pseudocode of Fig. 4.2. Prior to the initialization, **State** register is loaded with $0^{64} \parallel N$. Once **State** register is initialized, the initialization process starts by encrypting the nonce ($0^{64} \parallel N$) with AES. Then, 64 bits of the encrypted nonce is chopped by the “chop” function as in Fig. 6.1, and this chopped value is stored into the Δ register (this is initialization of Δ). After these initializations, 128-bit associated data blocks are fetched and sent to the ρ module along with the previous AES output to produce a 128 bit intermediate state. This state is partially masked with 64-bit Δ for every AES call. After all the associated data blocks are processed, the message blocks are processed in the same manner, except that the ρ function produces 128-bit ciphertext blocks in addition to the intermediate state values. Finally, after the message processing is over, the tag is generated using an additional AES call.

Combined Encryption and Decryption: As mentioned earlier, we here focus on the encryption-only circuit. However, due to the similarity between the encryption and the decryption modes, the combined hardware for encryption and decryption can be built with a small increase in the area, with the same throughput. This can be done by adding a control flow to a binary signal for mode selection.

6.3 Implementation Results

We implemented COFB on Xilinx Virtex 6 and Virtex 7, using VHDL and Xilinx ISE 13.4. AES-128 is used as the internal blockcipher. Table 3 presents the im-

Table 3. Implementation results of COFB on FPGAs.

Platform	# Slice Registers	# LUTs	# Slices	Frequency (MHz)	Gbps	Mbps/LUT	Mbps/Slice
Virtex 6	722	1075	442	267.20	2.85	2.24	6.45
Virtex 7	722	1456	555	264.24	2.82	2.22	5.08

plementation results of COFB on Virtex 7 with the target device xc7vx330t and Virtex 6 with the target device xc6vx760. We employ RTL approach and a basic iterative type architecture. The areas are listed in the number of Slice Registers, Slice LUTs and Occupied Slices. We also report frequency (MHz), Throughput (Gbps), and throughput-area efficiency. In the full version, we will show the area utilization for this basic AES-based implementation.

For AES, we use the implementation available from Athena [1] maintained by George Mason University. This implementation stores all the round subkeys in a single register to make the AES implementation faster and parallelizable. However, the main motivation of COFB is to reduce hardware footprint. Hence, we change the above implementation to a sequential one such that it processes only one AES round in a single clock cycle. This in turn eliminates the need to store all the round subkeys in a single register and reduces the hardware area consumed by the AES module.

6.4 Comparison with ATHENA Database

We compare our implementation of COFB with the results published in ATHENA Database [2], taking Virtex 6 and Virtex 7 as our target platforms. We first warn that this is a rough comparison. Here, we ignore the overhead to support the GMU API and the fact that ours is encryption-only while the others are (to the best of our knowledge) supporting both encryption and decryption, and the difference in the achieved security level, both quantitative and qualitative. We acknowledge that supporting GMU API will require some additional overhead to the current figures of COFB. Nevertheless, we think the current figures of COFB suggest that small hardware implementations are possible compared with other blockcipher AE modes shown in the table, using the same AES-128, even if we add a circuit for supporting GMU API and decryption.

We also remark that it is basically hard to compare COFB using AES-128 with other non-block-cipher-based AE schemes in the right way, because of the difference in the primitives and the types of security guarantee. For example, ACORN is built from scratch and does not have any provable security result, and is subjected to several cryptanalysis [20, 45, 44, 34]. Joltik and JAMBU-SIMON employ lightweight (tweakable) blockciphers allowing smaller implementation than AES, and Sponge AE schemes (ASCON, Ketje, NORX, and PRIMATES-HANUMAN) use a keyless permutation of a large block size to avoid key scheduling circuit and have the provable security relying on the random permutation model. In Table 4, we provide the comparison table only on the Vertex 6 platform. The comparison table on the Vertex 7 platform will be provided in the full version.

Table 4. Comparison on Virtex 6 [2]. In the “Primitive” column, SC denotes Stream cipher, (T)BC denotes (Tweakable) blockcipher, and BC-RF denotes the blockcipher’s round function.

Scheme	Primitive	#LUT	#Slices	Gbps	Mbps / LUT	Mbps / Slices
ACORN [51]	SC	455	135	3.112	6.840	23.052
AEIS [53]	BC-RF	7592	2028	70.927	9.342	34.974
AES-COPA [10]	BC	7754	2358	2.500	0.322	1.060
AES-GCM [22]	BC	3175	1053	3.239	1.020	3.076
AES-OTR [36]	BC	5102	1385	2.741	0.537	1.979
AEZ [26]	BC-RF	4597	1246	8.585	0.747	2.756
ASCON [21]	Sponge	1271	413	3.172	2.496	7.680
CLOC [29]	BC	3145	891	2.996	0.488	1.724
DEOXY [31]	TBC	3143	951	2.793	0.889	2.937
ELmD [19]	BC	4302	1584	3.168	0.736	2.091
JAMBU-AES [52]	BC	1836	652	1.999	1.089	3.067
JAMBU-SIMON [52]	BC (non-AES)	1222	453	0.363	0.297	0.801
Joltik [30]	TBC	1292	442	0.853	0.660	0.826
Ketje [14]	Sponge	1270	456	7.345	5.783	16.107
Minalpher [46]	BC (non-AES)	2879	1104	1.831	0.636	1.659
NORX [11]	Sponge	2964	1016	11.029	3.721	10.855
PRIMATES-HANUMAN [8]	Sponge	1012	390	0.964	0.953	2.472
OCB [33]	BC	4249	1348	3.122	0.735	2.316
SCREAM [24]	TBC	2052	834	1.039	0.506	1.246
SILC [29]	BC	3066	921	4.040	1.318	4.387
Tiaoxin [38]	BC-RF	7123	2101	52.838	7.418	25.149
TrivA-ck [18]	SC	2118	687	15.374	7.259	22.378
COFB	BC	1075	442	2.850	2.240	6.450

7 Conclusion

This paper presents COFB, a blockcipher mode for AE focusing on the state size. When instantiated with an n -bit blockcipher, COFB operates at rate-1, and requires state size of $1.5n$ bits, and is provable secure up to $O(2^{n/2}/n)$ queries based on the standard PRP assumption on the blockcipher. In fact this is the first scheme fulfilling these features at once. A key idea of COFB is a new type of feedback function combining both plaintext and ciphertext blocks. We have also presented the hardware implementation results, which demonstrate the effectiveness of our approach.

Acknowledgements. The authors thank the anonymous reviewers for helpful feedback. The work by Tetsu Iwata was supported in part by JSPS KAKENHI, Grant-in-Aid for Scientific Research (B), Grant Number 26280045, and was partially carried out while visiting Nanyang Technological University, Singapore.

References

1. ATHENA: Automated Tool for Hardware Evaluation. <https://cryptography.gmu.edu/athena/>.

2. Authenticated Encryption FPGA Ranking. https://cryptography.gmu.edu/athenadb/fpga_auth_cipher/rankings_view.
3. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. <http://competitions.cr.jp.to/caesar.html/>.
4. Recommendation for Block Cipher Modes of Operation: Methods and Techniques. NIST Special Publication 800-38A, 2001. National Institute of Standards and Technology.
5. Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality . NIST Special Publication 800-38C, 2004. National Institute of Standards and Technology.
6. Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication. NIST Special Publication 800-38B, 2005. National Institute of Standards and Technology.
7. NIST FIPS 197. Advanced Encryption Standard (AES). *Federal Information Processing Standards Publication*, 197, 2001.
8. Elena Andreeva, Begül Bilgin, Andrey Bogdanov, Atul Luykx, Florian Mendel, Bart Mennink, Nicky Mouha, Qingju Wang, and Kan Yasuda. PRIMATES v1.02. Submission to CAESAR. 2016. <https://competitions.cr.jp.to/round2/primatesv102.pdf>.
9. Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. Parallelizable and Authenticated Online Ciphers. In *ASIACRYPT (1)*, volume 8269 of *LNCS*, pages 424–443. Springer, 2013.
10. Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. AES-COPA v.2. Submission to CAESAR. 2015. <https://competitions.cr.jp.to/round2/aescopav2.pdf>.
11. Jean-Philippe Aumasson, Philipp Jovanovic, and Samuel Neves. NORX v3.0. Submission to CAESAR. 2016. <https://competitions.cr.jp.to/round3/norxv30.pdf>.
12. Subhadeep Banik, Andrey Bogdanov, and Kazuhiko Minematsu. Low-Area Hardware Implementations of CLOC, SILC and AES-OTR. DIAC 2015.
13. Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of the cipher block chaining message authentication code. *J. Comput. Syst. Sci.*, 61(3):362–399, 2000.
14. Guido Bertoni, Michaël Peeters, Joan Daemen, Gilles Van Assche, and Ronny Van Keer. Ketje v2. Submission to CAESAR. 2016. <https://competitions.cr.jp.to/round3/ketjev2.pdf>.
15. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: An Ultra-Lightweight Block Cipher. In *CHES 2007*, pages 450–466, 2007.
16. Andrey Bogdanov, Florian Mendel, Francesco Regazzoni, Vincent Rijmen, and Elmar Tischhauser. ALE: AES-Based Lightweight Authenticated Encryption. In *FSE 2013*, pages 447–466, 2013.
17. Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knezevic, Lars R. Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, Peter Rombouts, Søren S. Thomsen, and Tolga Yalçın. PRINCE - A Low-Latency Block Cipher for Pervasive Computing Applications - Extended Abstract. In *ASIACRYPT 2012*, pages 208–225, 2012.
18. Avik Chakraborti and Mridul Nandi. TriviA-ck-v2. Submission to CAESAR. 2015. <https://competitions.cr.jp.to/round2/triviackv2.pdf>.
19. Nilanjan Datta and Mridul Nandi. Proposal of ELmD v2.1. Submission to CAESAR. 2015. <https://competitions.cr.jp.to/round2/elmdv21.pdf>.

20. Prakash Dey, Raghvendra Singh Rohit, and Avishek Adhikari. Full key recovery of ACORN with a single fault. *J. Inf. Sec. Appl.*, 29:57–64, 2016.
21. Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer. Ascon v1.2. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/asconv12.pdf>.
22. Morris Dworkin. Recommendation for block cipher modes of operation: Galois/counter mode (GCM) and GMAC. NIST Special Publication 800-38D, 2011. csrc.nist.gov/publications/nistpubs/800-38D/SP-800-38D.pdf.
23. Ewan Fleischmann, Christian Forler, and Stefan Lucks. McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes. In *FSE 2012*, pages 196–215, 2012.
24. Vincent Grosso, Ga etan Leurent, Francois-Xavier Standaert, Kerem Varici, Anthony Journault, Francois Durvaux, Lubos Gaspar, and St ephanie Kerckhof. SCREAM Side-Channel Resistant Authenticated Encryption with Masking. Submission to CAESAR. 2015. <https://competitions.cr.yp.to/round2/screamv3.pdf>.
25. Jian Guo, Thomas Peyrin, Axel Poschmann, and Matthew J. B. Robshaw. The LED Block Cipher. In *CHES 2011*, pages 326–341, 2011.
26. Viet Tung Hoang, Ted Krovetz, and Philip Rogaway. AEZ v4.2: Authenticated Encryption by Enciphering. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/aezv42.pdf>.
27. Tetsu Iwata and Kaoru Kurosawa. OMAC: One-Key CBC MAC. In *FSE*, pages 129–153, 2003.
28. Tetsu Iwata, Kazuhiko Minematsu, Jian Guo, and Sumio Morioka. CLOC: Authenticated Encryption for Short Input. In *FSE 2014*, pages 149–167, 2014.
29. Tetsu Iwata, Kazuhiko Minematsu, Jian Guo, Sumio Morioka, and Eita Kobayashi. CLOC and SILC. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/clocsilcv3.pdf>.
30. J er emy Jean, Ivica Nikoli c, and Thomas Peyrin. Joltik v1.3. Submission to CAESAR. 2015. <https://competitions.cr.yp.to/round2/joltikv13.pdf>.
31. J er emy Jean, Ivica Nikoli c, and Thomas Peyrin. Deoxys v1.41. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/deoxysv141.pdf>.
32. Ted Krovetz and Phillip Rogaway. The Software Performance of Authenticated-Encryption Modes. In *FSE*, pages 306–327, 2011.
33. Ted Krovetz and Phillip Rogaway. OCB(v1.1). Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/ocbv11.pdf>.
34. Fr ed eric Lafitte, Liran Lerman, Olivier Markowitch, and Dirk Van Heule. SAT-based cryptanalysis of ACORN. *IACR Cryptology ePrint Archive*, 2016:521, 2016.
35. Kazuhiko Minematsu. Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In *EUROCRYPT*, volume 8441 of *LNCS*, pages 275–292. Springer, 2014.
36. Kazuhiko Minematsu. AES-OTR v3.1. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/aesotrv31.pdf>.
37. Amir Moradi, Axel Poschmann, San Ling, Christof Paar, and Huaxiong Wang. Pushing the Limits: A Very Compact and a Threshold Implementation of AES. In *EUROCRYPT 2011*, pages 69–88, 2011.
38. Ivica Nikoli c. Tiaoxin – 346. Submission to CAESAR. 2016. <https://competitions.cr.yp.to/round3/tiaoxinv21.pdf>.
39. J. Patarin. Etude des G en erateurs de Permutations Bas es sur le Sch ema du D.E.S. Phd Th esis de Doctorat de l’Universit e de Paris 6, 1991.

40. Thomas Peyrin, Siang Meng Sim, Lei Wang, and Guoyan Zhang. Cryptanalysis of JAMBU. In *FSE 2015*, pages 264–281, 2015.
41. Phillip Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In *ASIACRYPT*, pages 16–31, 2004.
42. Phillip Rogaway, Mihir Bellare, and John Black. OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Trans. Inf. Syst. Secur.*, 6(3):365–403, 2003.
43. Phillip Rogaway and Thomas Shrimpton. A Provable-Security Treatment of the Key-Wrap Problem. In *EUROCRYPT*, pages 373–390, 2006.
44. Md. Iftexhar Salam, Harry Bartlett, Ed Dawson, Josef Pieprzyk, Leonie Simpson, and Kenneth Koon-Ho Wong. Investigating cube attacks on the authenticated encryption stream cipher ACORN. In *ATIS 2016*, pages 15–26, 2016.
45. Md. Iftexhar Salam, Kenneth Koon-Ho Wong, Harry Bartlett, Leonie Ruth Simpson, Ed Dawson, and Josef Pieprzyk. Finding state collisions in the authenticated encryption stream cipher ACORN. In *Proceedings of the Australasian Computer Science Week Multiconference*, page 36, 2016.
46. Yu Sasaki, Yosuke Todo, Kazumaro Aoki, Yusuke Naito, Takeshi Sugawara, Yumiko Murakami, Mitsuru Matsui, and Shoichi Hirose. Minalpher v1.1. Submission to CAESAR. 2015. <https://competitions.cr.yj.to/round2/minalpherv11.pdf>.
47. Willem Schro e, Bart Mennink, Elena Andreeva, and Bart Preneel. Forgery and Subkey Recovery on CAESAR Candidate iFeed. In *SAC*, volume 9566 of *LNCS*, pages 197–204. Springer, 2015.
48. Kyoji Shibutani, Takanori Isobe, Harunaga Hiwatari, Atsushi Mitsuda, Toru Akishita, and Taizo Shirai. Piccolo: An Ultra-Lightweight Blockcipher. In *CHES 2011*, pages 342–357, 2011.
49. Tomoyasu Suzaki, Kazuhiko Minematsu, Sumio Morioka, and Eita Kobayashi. TWINE : A Lightweight Block Cipher for Multiple Platforms. In *SAC 2012*, pages 339–354, 2012.
50. Serge Vaudenay. Decorrelation: A Theory for Block Cipher Security. *J. Cryptology*, 16(4):249–286, 2003.
51. Hongjun Wu. ACORN: A Lightweight Authenticated Cipher (v3). Submission to CAESAR. 2016. <https://competitions.cr.yj.to/round3/acornv3.pdf>.
52. Hongjun Wu and Tao Huang. The JAMBU Lightweight Authentication Encryption Mode (v2.1). Submission to CAESAR. 2016. <https://competitions.cr.yj.to/round3/jambuv21.pdf>.
53. Hongjun Wu and Bart Preneel. AEGIS : A Fast Authenticated Encryption Algorithm (v1.1). Submission to CAESAR. 2016. <https://competitions.cr.yj.to/round3/aegisv11.pdf>.
54. Liting Zhang, Wenling Wu, Han Sui, and Peng Wang. iFeed[AES] v1. Submission to CAESAR. 2014. <https://competitions.cr.yj.to/round1/ifeedaesv1.pdf>.