

Cryptanalysis of OCB2

Akiko Inoue and Kazuhiko Minematsu

NEC Corporation, Japan

a-inoue@cj.jp.nec.com, k-minematsu@ah.jp.nec.com

Abstract. We present practical attacks against OCB2, an ISO-standard authenticated encryption (AE) scheme. OCB2 is a highly-efficient blockcipher mode of operation. It has been extensively studied and widely believed to be secure thanks to the provable security proofs. Our attacks allow the adversary to create forgeries with single encryption query of almost-known plaintext. This attack can be further extended to powerful almost-universal and universal forgeries using more queries. The source of our attacks is the way OCB2 implements AE using a tweakable blockcipher, called XEX*. We have verified our attacks using a reference code of OCB2. Our attacks do not break the privacy of OCB2, and are not applicable to the others, including OCB1 and OCB3.

Keywords: OCB, Authenticated Encryption, Cryptanalysis, Forgery, XEX

1 Introduction

Authenticated encryption (AE) is a form of symmetric-key encryption that provides both confidentiality and authenticity of messages. Now it is widely accepted that AE is a fundamental security tool for many practical applications, such as TLS.

OCB is a blockcipher mode of operation for AE. It is one of the most celebrated schemes in the cryptography for its beautiful and innovative architecture. OCB is very efficient. In fact, it is as efficient as encryption-only modes, and is parallelizable. There are three versions of OCB. The first version (OCB1) was proposed at ACM CCS 2001 by Rogaway et. al [RBBK01]. The second one (OCB2) was proposed at Asiacrypt 2004 by Rogaway [Rog04a] (hereafter Rog04), and the latest one (OCB3) was proposed at FSE 2011 by Krovetz and Rogaway [KR11]. Each version of OCB has received significant attentions from theory and practice. OCB1 was considered for the security mechanism of Wireless LAN (IEEE 802-11), OCB2 is standardized in ISO/IEC 19772:2009 [ISO09], and OCB3 is in RFC 7253 [RFC14]. Moreover, OCB3 is a finalist of CAESAR competition¹. Various versions of OCB have been implemented in the cryptographic libraries, such as Botan, BouncyCastle, LibTomCrypt, OpenSSL, SJCL etc.

¹ <https://competitions.cr.yp.to/caesar.html>

The security of (all versions of) OCB has been extensively studied. For each version, the author(s) provided the security reductions, that is, a break of OCB implies the break of the underlying blockcipher. If the blockcipher is secure, the security bounds (roughly, the success probability of attacks) are $O(\sigma^2/2^n)$ for σ processed blocks, called “birthday bounds” named after the birthday paradox. From the attacker point of view, Ferguson [Fer02] and Sun et.al [SWZ12] showed collision attacks that need $2^{n/2}$ processed blocks, implying the security bounds are tight. Attacks in the misuse scenarios are proposed by Andreeva et. al [ABL⁺14] and Ashur et. al [ADL17]. Robustness of the CAESAR candidates including OCB was studied by Damian and Vaudenay [VV18]. From the provable security point of view, Aoki and Yasuda [AY13] showed a relaxation on the assumption of the underlying blockcipher in a modified OCB, and Ritam and Nandi [BN17] showed an improvement on the security bound of OCB3.

Despite aforementioned attacks and studies on provable security, it is widely believed that the security proofs are correct, thus all versions of OCB are provably secure in the standard, nonce-respecting scenario, up to the birthday bounds.

In this paper, we first invalidate this belief by presenting an attack against OCB2. The attack is practical: it needs only one encryption query to create a forgery with success probability one. No heavy computation or large memory is needed, we just take a sum of some blocks to create a forgery. The attack works for any n -bit blockcipher E , and works for any n for which OCB2 can be defined. The forgery can be any long ciphertext depending on the first encryption query, and the encryption query is almost a known-plaintext query. In addition, while the first forgery is existential, the subsequent forgeries have much more freedom, as the adversary learns the mask applied to E from the first forgery. This enables to learn multiple blockcipher inputs and outputs, and it results in an almost-universal forgery attack in the reforgeability setting [BC09, FLLW17] (see Section 4.3) using few more queries. Moreover, using more queries (roughly $a + m$ when the target AD has a blocks and the target plaintext has m blocks), a universal forgery attack is finally possible.

Our attacks point out a critical flaw of OCB2, and show that the authenticity bound of OCB2 turns out to be one. We emphasize that this is not a flaw of the global structure of OCB shared by the all versions, but is about the details of OCB2 that have been overlooked for more than a decade.

The root of our attacks is in the tweakable blockcipher (TBC) [LRW02] inside OCB2, called XEX*. More precisely, it is about how OCB2 uses XEX* to implement AE. Here, an encryption of XEX* is either $C = \Delta \oplus E_K(\Delta \oplus M)$ or $C = E_K(\Delta \oplus M)$, where Δ is derived from the encryption of nonce multiplied by a constant in $\text{GF}(2^n)$. The former is called XEX, and the latter is XE, hence XEX* combines XEX and XE.

Original security proof of Rog04 showed that XEX* implements a provably secure TBC and OCB2 can be seen as a mode of XEX*. Rog04 concluded that OCB2 is provably secure since if XEX* was replaced with a perfect tweakable random permutation, the resulting mode is ideally secure. This is an elegant hybrid argument, which significantly reduces the complexity of the proof. In fact,

the proof strategy of Rog04 and the idea of using TBC as a mode component was followed by many schemes. In the ongoing CAESAR competition, the most submissions of parallelizable AE based on (tweakable) blockcipher or cryptographic permutation, including OCB3 itself, employed the strategy of Rog04 or OCB in general. However, since XEX^* is not a plain TBC (called *tagged* TBC), there are some access rules that the adversary must follow in the game of distinguishing XEX^* from the ideal. We found that, even though OCB2 is a mode of XEX^* , the authenticity game for OCB2 forces the adversary to violate the access rule of XEX^* . This implies that, even if XEX^* is a secure (tagged) TBC, its security can not imply the security of OCB2. When OCB2 is invoked, this violated access occurs around the processing of the last two message blocks, where XEX is used for the last-but-one block and XE is used for the last one. The attack exploits this “violated” access and a relationship between the different masks applied to the blockcipher.

We have confirmed our attacks using the reference code of OCB2 by Krovetz². Appendix A shows an example.

We also show how to fix OCB2, say by using XEX for the last block. Our attacks are not applicable to OCB1 and OCB3, and we have not found other schemes based on OCB that could be affected.

2 Preliminaries

2.1 Notations

Let $\{0,1\}^*$ be the set of all binary strings, including the empty string ε . For $X \in \{0,1\}^*$, its bit length is denoted by $|X|$, and $\text{msb}_c(X)$ denotes the first $c \leq |X|$ bits of X . Here, $|\varepsilon| = 0$. Let 0^c be the sequence of c zeros, where $0^0 = \varepsilon$. The concatenation of two bit strings X and Y is written $X \parallel Y$, or XY when no confusion is possible.

When $0 \leq |X| \leq n$, let $\text{len}(X)$ be the n -bit encoding of $|X|$. We also define the zero padding, $X \parallel 0^*$, and the one-zero padding, $X \parallel 10^*$. Both are X when $|X| = n$ and $X \parallel 0^* = X \parallel 0^{n-|X|}$ and $X \parallel 10^* = X \parallel 10^{n-|X|-1}$ when $0 \leq |X| < n$.

For $X \in \{0,1\}^*$, we also define the parsing into n -bit blocks denoted by

$$(X[1], X[2], \dots, X[m]) \stackrel{n}{\leftarrow} X,$$

where $m = |X|_n \stackrel{\text{def}}{=} \lceil |X|/n \rceil$, $X[1] \parallel X[2] \parallel \dots \parallel X[m] = X$, $|X[i]| = n$ for $1 \leq i < m$ and $0 < |X[m]| \leq n$ when $|X| > 0$. When $|X| = 0$, we let $X[1] \stackrel{n}{\leftarrow} X$, where $X[1] = \varepsilon$.

If X is a random variable distributed uniformly over \mathcal{X} , we write $X \stackrel{\$}{\leftarrow} \mathcal{X}$.

² <http://web.cs.ucdavis.edu/~rogaway/ocb/code-2.0.htm>.

(TWEAKABLE) BLOCKCIPHER. A tweakable blockcipher (TBC) [LRW02] is a keyed function $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$ such that for each $(K, T) \in \mathcal{K} \times \mathcal{T}$, $\tilde{E}(K, T, \cdot)$ is a permutation over \mathcal{M} . Here, K is the key and T is a public value called tweak. A conventional block cipher is a TBC with \mathcal{T} being a singleton, and specifically written as $E : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{M}$. The encryption of $X \in \mathcal{M}$ under key $K \in \mathcal{K}$ and tweak $T \in \mathcal{T}$ is $\tilde{E}(K, T, X)$ and is also written as $\tilde{E}_K(T, X)$ or $\tilde{E}_K^T(X)$. For blockcipher we write as $E_K(X)$. The decryption is written as $\tilde{E}_K^{-1, T}(Y)$ for TBCs and $E_K^{-1}(Y)$ for blockciphers. For any $T \in \mathcal{T}$ and $K \in \mathcal{K}$, when $Y = \tilde{E}_K^T(X)$ we have $\tilde{E}_K^{-1, T}(Y) = X$.

We may omit the subscript K when being keyed with K is obvious. Moreover, for a mode of operation using E_K , we may treat E_K as a key and write as Mode_E . We typically assume $n = 128$.

GALOIS FIELD. For $a \in \{0, 1\}^n$, we may see it as an element of $\text{GF}(2^n)$ by taking it as a coefficient vector. Following [Rog04a][IK03], by writing $2 \cdot a$ or $2a$, we mean the *doubling* over $\text{GF}(2^n)$, i.e., the multiplication by the generator of the field (the polynomial \mathbf{x}) and a .

In a similar fashion, $2^c a$ denotes c -times doublings of a , and $3a$ denotes $2a \oplus a$. Combined expressions such as $2^i \cdot 3^j a$ are similarly defined. For example, $2^2 \cdot 3a = 3(2^2 a) = 2^3 a \oplus 2^2 a$.

For $n = 128$, it is common to define the field $\text{GF}(2^n)$ by the lexicographically-first primitive polynomial, $\mathbf{x}^{128} + \mathbf{x}^7 + \mathbf{x}^2 + \mathbf{x} + 1$. With this, $2a$ is $(a \ll 1)$ if $\text{msb}_1(a) = 0$ and $(a \ll 1) \oplus (0^{120}10000111)$ if $\text{msb}_1(a) = 1$, where $(a \ll 1)$ denotes the left-shift of a by one bit.

2.2 AEAD

In most cases, we say AE to mean Authenticated Encryption with Associated Data (AEAD) [Rog02]. Let Π be an AEAD scheme defined over key space \mathcal{K} , nonce space \mathcal{N} , associated data (AD) space \mathcal{A} , message space \mathcal{M} and tag space $\mathcal{T} = \{0, 1\}^\tau$ for some fixed τ . Here, AD is a part of message that should not be encrypted but authenticated. For example AD can be header information of a network protocol. The encryption of Π is a function $\Pi.\mathcal{E} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{M} \times \mathcal{T}$, and the decryption (or verification) is a function $\Pi.\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{M} \cup \{\perp\}$, where \perp denotes the verification failure. To encrypt plaintext M with nonce N and AD A and key K , we compute $\Pi.\mathcal{E}_K(N, A, M) = (C, T)$ to produce ciphertext C and tag T . The tuple (N, A, C, T) is sent to the receiver. For decryption, we compute $\Pi.\mathcal{D}_K(N', A', C', T')$ and it returns the plaintext M' if the verification succeeds, and \perp if the verification fails. We write Π_K to mean the tuple $(\Pi.\mathcal{E}_K, \Pi.\mathcal{D}_K)$.

SECURITY NOTIONS. The security of AEAD can be defined by two notions, privacy and authenticity [BRW04, Rog04c]. The privacy notion captures the indistinguishability from random for a nonce-respecting adversary in a chosen-plaintext-attack setting. Let $\mathcal{A}^{O_1, O_2, \dots, O_c}$ denote the adversary \mathcal{A} accessing c

oracles, O_1, \dots, O_c , in an arbitrarily order. The privacy notion is defined over the game where \mathcal{A}^{O_1} returns a binary decision at the end of the game, and O_1 is either $\Pi.\mathcal{E}_K$ or the random-bits oracle $\$$. Here, $\$$ returns $(C, T) \xleftarrow{\$} \{0, 1\}^{|M|} \times \mathcal{T}$ for any encryption query (N, A, M) . The privacy advantage is defined as

$$\mathbf{Adv}_{\Pi_K}^{\text{priv}}(\mathcal{A}) \stackrel{\text{def}}{=} \Pr \left[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\Pi.\mathcal{E}_K(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\$(\cdot, \cdot)} \Rightarrow 1 \right],$$

where $\mathcal{A}^{O_1} \Rightarrow 1$ denotes the event that the final decision of \mathcal{A} is 1. We assume that \mathcal{A} in the privacy game is nonce-respecting, that is, \mathcal{A} does not make two queries with the same nonce.

For the authenticity notion, we consider \mathcal{A} accessing $\Pi.\mathcal{E}_K$ and $\Pi.\mathcal{D}_K$, where \mathcal{A} is considered to win when it creates a successful forgery. Formally, the authenticity advantage is defined as

$$\mathbf{Adv}_{\Pi_K}^{\text{auth}}(\mathcal{A}) \stackrel{\text{def}}{=} \Pr \left[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\Pi.\mathcal{E}_K, \Pi.\mathcal{D}_K} \text{ forges} \right],$$

where \mathcal{A} forges if it receives any $M' \neq \perp$ from $\Pi.\mathcal{D}_K$. To avoid a trivial win, \mathcal{A} is assumed not to query (N, A, C, T) to $\Pi.\mathcal{D}_K$ when there is a previous encryption query (N, A, M) for some M and the corresponding response (C, T) . We also assume \mathcal{A} is nonce-respecting with respect to encryption queries.

2.3 Security of (Tweakable) Blockciphers

Let $\tilde{E}_K : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$ be a TBC. Let \tilde{P} be the tweakable uniform random permutation (TURP). This is an information-theoretic TBC with key uniformly distributed over all the tweakable permutations, i.e. the set of $f(T, X) : \mathcal{T} \times \mathcal{M} \rightarrow \mathcal{M}$ such that $f(T, *)$ is a permutation over \mathcal{M} for all $T \in \mathcal{T}$. The decryption is denoted by \tilde{P}^{-1} . We define Tweakable Pseudorandom Permutation (TPRP) advantage and Tweakable Strong PRP (TSPRP) advantage:

$$\mathbf{Adv}_{\tilde{E}_K}^{\text{tprp}}(\mathcal{A}) \stackrel{\text{def}}{=} \Pr \left[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\tilde{E}_K} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\tilde{P}} \Rightarrow 1 \right] \quad (1)$$

$$\mathbf{Adv}_{\tilde{E}_K}^{\text{tsprp}}(\mathcal{A}_{\pm}) \stackrel{\text{def}}{=} \Pr \left[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\tilde{E}_K, \tilde{E}_K^{-1}} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\tilde{P}, \tilde{P}^{-1}} \Rightarrow 1 \right]. \quad (2)$$

Here, \mathcal{A} performs Chosen-plaintext attack with chosen tweaks, and \mathcal{A}_{\pm} performs Chosen-ciphertext attack with chosen tweaks (i.e. it can query any (T, M) to encrypt and any (T, C) to decrypt).

For blockcipher $E : \mathcal{K} \times \mathcal{M} \rightarrow \mathcal{M}$, we analogously define PRP advantage $\mathbf{Adv}_{E_K}^{\text{prp}}(\mathcal{A})$ and SPRP advantage $\mathbf{Adv}_{E_K}^{\text{sprp}}(\mathcal{A}_{\pm})$, using URP P (which uniformly distributes over all permutations over \mathcal{M}) as an ideal primitive.

3 OCB2

As described, OCB2 is a blockcipher mode for AEAD proposed at [Rog04a]. It is parallelizable, and is rate-1 as it needs one blockcipher call to process one

<p>Algorithm OCB2.$\mathcal{E}_E(N, A, M)$</p> <ol style="list-style-type: none"> 1. $L \leftarrow E(N)$ 2. $(M[1], \dots, M[m]) \xleftarrow{r} M$ 3. for $i = 1$ to $m - 1$ 4. $C[i] \leftarrow 2^i L \oplus E(2^i L \oplus M[i])$ 5. $\text{Pad} \leftarrow E(2^m L \oplus \text{len}(M[m]))$ 6. $C[m] \leftarrow M[m] \oplus \text{msb}_{ M[m] }(\text{Pad})$ 7. $\Sigma \leftarrow C[m] \parallel 0^* \oplus \text{Pad}$ 8. $\Sigma \leftarrow M[1] \oplus \dots \oplus M[m-1] \oplus \Sigma$ 9. $T \leftarrow E(2^m 3L \oplus \Sigma)$ 10. if $A \neq \varepsilon$ then $T \leftarrow T \oplus \text{PMAC}_E(A)$ 11. $T \leftarrow \text{msb}_\tau(T)$ 12. return (C, T) 	<p>Algorithm OCB2.$\mathcal{D}_E(N, A, C, T)$</p> <ol style="list-style-type: none"> 1. $L \leftarrow E(N)$ 2. $(C[1], \dots, C[m]) \xleftarrow{r} C$ 3. for $i = 1$ to $m - 1$ 4. $M[i] \leftarrow 2^i L \oplus E^{-1}(2^i L \oplus C[i])$ 5. $\text{Pad} \leftarrow E(2^m L \oplus \text{len}(C[m]))$ 6. $M[m] \leftarrow C[m] \oplus \text{msb}_{ C[m] }(\text{Pad})$ 7. $\Sigma \leftarrow C[m] \parallel 0^* \oplus \text{Pad}$ 8. $\Sigma \leftarrow M[1] \oplus \dots \oplus M[m-1] \oplus \Sigma$ 9. $T^* \leftarrow E(2^m 3L \oplus \Sigma)$ 10. if $A \neq \varepsilon$ then $T^* \leftarrow T^* \oplus \text{PMAC}_E(A)$ 11. $T^* \leftarrow \text{msb}_\tau(T^*)$ 12. if $T = T^*$ return M 13. else return \perp
<p>Algorithm PMAC$_E(A)$</p> <ol style="list-style-type: none"> 1. $S \leftarrow 0^n$ 2. $V \leftarrow 3^2 E(0^n)$ 3. $(A[1], \dots, A[a]) \xleftarrow{r} A$ 4. for $i = 1$ to $a - 1$ 5. $S \leftarrow S \oplus E(2^i V \oplus A[i])$ 6. $S \leftarrow S \oplus A[a] \parallel 10^*$ 7. if $A[a] = n$ 8. $Q \leftarrow E(3V \oplus S)$ 9. else $Q \leftarrow E(3^2 V \oplus S)$ 10. return Q 	

Fig. 1. Algorithms of OCB2.

message block. The pseudocode of OCB2 is shown in Fig. 1. The basic form of OCB2 presented in Rog04 (Fig. 2) is a plain AE³, where there is no AD. Section 11 of Rog04 described that the plain OCB2 can be converted into an AEAD, by computing PMAC for AD and taking a sum (XOR) of PMAC output and the tag of plain OCB2 *only if* AD is non-empty. This scheme is referred to as AEM in Rog04. The specification of PMAC in OCB2 is shown in the full version of Rog04 [Rog04b]. It is different from the initial version [BR02] in that it uses doublings for mask generation and is designed so that it is computationally independent from the plain AE part. We remark that our attack is independent of the specification of PMAC.

³ It was named OCB1 in the paper.

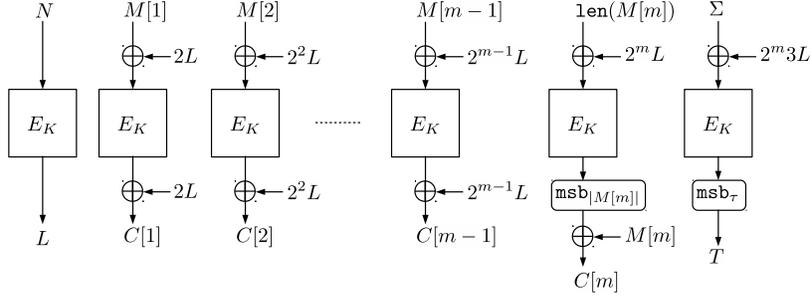


Fig. 2. OCB2 for the case of empty AD.

4 Attacks

4.1 Minimal Example of Forgery

We describe the minimal example of our forgery attack against OCB2. For simplicity, let $\tau = n$.

1. Encrypt (N, A, M) , where A is empty, M is a $2n$ -bit message $M = M[1] \parallel M[2]$ specified as

$$M[1] = \mathbf{1en}(0^n),$$

and $M[2]$ is any n -bit block. Encryption oracle returns the pair of $2n$ -bit ciphertext and tag, $C = C[1] \parallel C[2]$ and T .

2. Decrypt (N', A', C', T') with $|C'| = n$ such that

$$N' = N, \tag{3}$$

$$A' = \varepsilon, \tag{4}$$

$$C' = C[1] \oplus \mathbf{1en}(0^n), \tag{5}$$

$$T' = M[2] \oplus C[2] \tag{6}$$

and T' will be always accepted as valid, while clearly $C' \neq C$ (different lengths), implying a successful forgery.

Why this works? Because

$$C[1] = 2L \oplus E(2L \oplus \mathbf{1en}(0^n)) \tag{7}$$

$$C[2] = M[2] \oplus \text{Pad}, \tag{8}$$

where $L = E(N)$ and $\text{Pad} = E(2^2L \oplus \text{len}(M[2])) = M[2] \oplus C[2]$. Let Pad' and Σ' be the true values for the decryption query. Then C' is decrypted to

$$M' = C' \oplus \text{Pad}' \quad (9)$$

$$= C' \oplus \underbrace{E(2L \oplus \text{len}(0^n))}_{\text{Pad}'} \quad (10)$$

$$= C[1] \oplus \text{len}(0^n) \oplus E(2L \oplus \text{len}(0^n)) \quad (11)$$

$$= \underbrace{2L \oplus E(2L \oplus \text{len}(0^n))}_{C[1]} \oplus \text{len}(0^n) \oplus E(2L \oplus \text{len}(0^n)) \quad (12)$$

$$= 2L \oplus \text{len}(0^n), \quad (13)$$

and the tag is generated as

$$T^* = E(2 \cdot 3L \oplus \Sigma') \quad (14)$$

$$= E(2 \cdot 3L \oplus \underbrace{C' \oplus \text{Pad}'}_{\Sigma'}) \quad (15)$$

$$= E(2 \cdot 3L \oplus M') \quad (16)$$

$$= E(2 \cdot 3L \oplus \underbrace{2L \oplus \text{len}(0^n)}_{M'}) \quad (17)$$

$$= E(2^2L \oplus \text{len}(0^n)) \quad (18)$$

$$= \text{Pad}, \quad (19)$$

where (18) follows from $2 \cdot 3L = 2^2L \oplus 2L$. Thus, T^* is determined by the first encryption query. See Fig. 3 for the graphical representation of the attack. As we mentioned, the attack is independent of the AD processing function (PMAC), and of the specification of $\text{len}(\cdot)$.

In the reference code, the input message is assumed to be a byte string and $\text{len}(X)$ is specified as $0^{n-8} \parallel \ell_X$, where ℓ_X denotes the standard binary encoding of $|X|$. Hence $\text{len}(0^n)$ for $n = 128$ is $0^{120}10^7$.

EXTENSIONS. Since we do not need the information on T in the first encryption query, the above attack works if AD is non-empty at the encryption query. The decryption query needs the empty AD. For the same reason, it also works even when the tag is truncated to $\tau < n$ bits; we set $T' = \text{msb}_\tau(M[2] \oplus C[2])$ and it will be accepted with probability one.

4.2 Forgery of Longer Messages

The attack of Section 4.1 can be extended to arbitrarily long messages:

1. Encrypt (N, A, M) , where A is arbitrary, M is a message of m blocks $M = M[1] \parallel \dots \parallel M[m-1] \parallel M[m]$ satisfying

$$M[m-1] = \text{len}(0^n),$$

and $M[m]$ is any s -bit string such that $\tau \leq s \leq n$. Encryption oracle returns the pair of ciphertext and tag, $C = C[1] \parallel \dots \parallel C[m-1] \parallel C[m]$ and T .

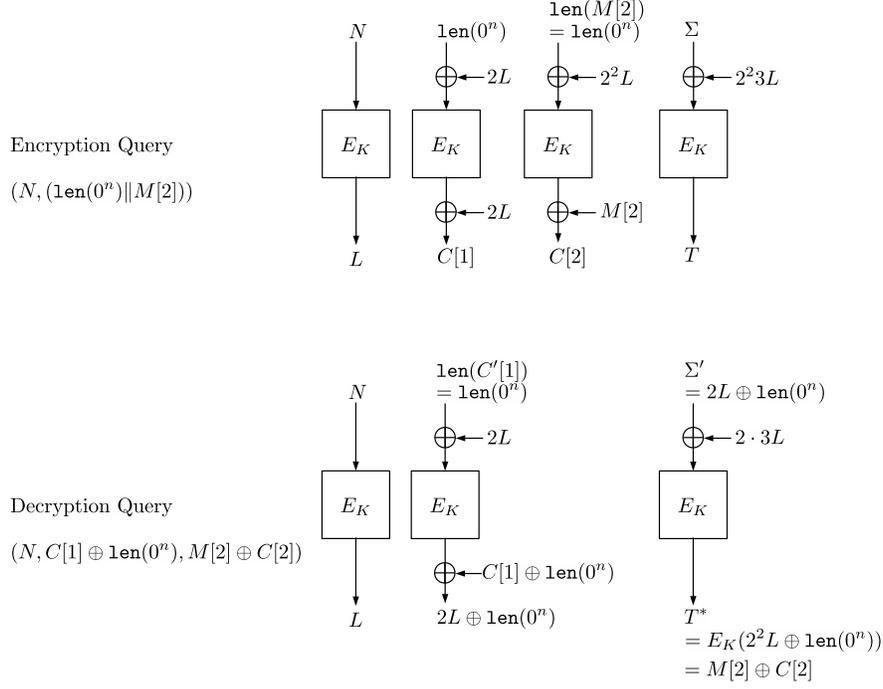


Fig. 3. Attack described at Section 4.1.

2. Decrypt (N', A', C', T') , such that $N' = N$, $A' = \varepsilon$ and C' has $m - 1$ blocks, $C' = C'[1] \parallel \dots \parallel C'[m - 2] \parallel C'[m - 1]$, specified as

$$C'[i] = C[i] \text{ for } 1 \leq i \leq m - 2, \quad (20)$$

$$C'[m - 1] = \sum_{i=1}^{m-2} M[i] \oplus C[m - 1] \oplus \mathbf{1en}(M[m]) \quad (21)$$

$$T' = \mathbf{msb}_\tau(M[m] \oplus C[m]). \quad (22)$$

This forgery attack again works. Let \bar{T}' be the true string of untruncated tag in the decryption query. Then we have

$$\bar{T}' = E(\Sigma' \oplus 3 \cdot 2^{m-1}L) \quad (23)$$

$$= E\left(\left(\sum_{i=1}^{m-2} M'[i] \oplus C'[m - 1] \oplus \text{Pad}'\right) \oplus 3 \cdot 2^{m-1}L\right) \quad (24)$$

$$= E\left(\sum_{i=1}^{m-2} M[i] \oplus C'[m - 1] \oplus C[m - 1] \oplus 2^{m-1}L \oplus 3 \cdot 2^{m-1}L\right), \quad (25)$$

where $M'[i] = M[i]$ is the i -th decrypted plaintext block, and $\text{Pad}' = C[m-1] \oplus 2^{m-1}L$. Since $2^{m-1}L \oplus 3 \cdot 2^{m-1}L = 2^mL$, the last term of (25) is further expanded as

$$E \left(\sum_{i=1}^{m-2} M[i] \oplus C'[m-1] \oplus C[m-1] \oplus 2^mL \right) \quad (26)$$

$$= E \left(\sum_{i=1}^{m-2} M[i] \oplus \left(\sum_{i=1}^{m-2} M[i] \oplus C[m-1] \oplus \mathbf{1en}(M[m]) \right) \oplus C[m-1] \oplus 2^mL \right) \quad (27)$$

$$= E(\mathbf{1en}(M[m]) \oplus 2^mL) \quad (28)$$

$$= \text{Pad} \quad (29)$$

Finally, we have

$$T^* = \text{msb}_\tau(\bar{T}') \quad (30)$$

$$= \text{msb}_\tau(\text{Pad}) \quad (31)$$

$$= \text{msb}_\tau(M[m] \oplus C[m]) \quad (\because \tau \leq |M[m]| \leq n) \quad (32)$$

$$= T' \quad (33)$$

Meanwhile, in the case that $\mathbf{1en}(M[m]) < \tau$, the adversary can forge T^* with probability $1/2^{\tau - \mathbf{1en}(M[m])}$, since the adversary only knows $\text{msb}_{\mathbf{1en}(M[m])}(\text{Pad})$ and has to guess the remaining $(\tau - \mathbf{1en}(M[m]))$ bits.

4.3 Almost Universal Forgery, Variant 1

We consider what will happen after the attack of Section 4.1 or 4.2, i.e. attacks in the reforgeability scenario. Here, the adversary learns $L = E(N)$ by the first forgery, since

$$\begin{aligned} M'[m-1] &= C'[m-1] \oplus \text{Pad}' \\ &= \sum_{i=1}^{m-2} M[i] \oplus \mathbf{1en}(M[m]) \oplus 2^{m-1}L, \end{aligned}$$

is returned by the decryption oracle, for the case of attack of Section 4.2. For $1 \leq i \leq m-1$, let $X[i]$ and $Y[i]$ be $X[i] = M[i] \oplus 2^iL$ and $Y[i] = C[i] \oplus 2^iL$, respectively. In addition, let $X[m] = \mathbf{1en}(M[m]) \oplus 2^mL$ and $Y[m] = \text{Pad}$. With the knowledge of L , the adversary learns the m pairs of E 's input and output, $\mathcal{S} = \{(X[i], Y[i]) : i = 1, \dots, m\}$, which allows the second forgery with more freedom than the first. For example, if the adversary wants to disturb only the j -th and k -th plaintext blocks for some $1 \leq j < k \leq m-1$, the second decryption

query $(N'', A'', C''[1] \parallel \dots \parallel C''[m], T'')$ can be set as $N'' = N$, $A'' = \varepsilon$, and

$$C''[i] = C[i] \quad (i \in \{1, \dots, m\} \setminus \{j, k\}), \quad (34)$$

$$C''[j] = Y[k] \oplus 2^j L, \quad (35)$$

$$C''[k] = Y[j] \oplus 2^k L, \quad (36)$$

$$T'' = T. \quad (37)$$

Then, the unverified plaintext will be

$$M''[i] = M[i] \quad (i \in \{1, \dots, m\} \setminus \{j, k\}) \quad (38)$$

$$M''[j] = X[k] \oplus 2^j L \quad (39)$$

$$M''[k] = X[j] \oplus 2^k L \quad (40)$$

and the valid checksum Σ'' is

$$\Sigma'' = \bigoplus_{i=1}^{m-1} M''[i] \oplus Y[m] \oplus C''[m] \quad (41)$$

$$= \bigoplus_{i \in \{1, \dots, m-1\} \setminus \{j, k\}} M[i] \oplus M''[j] \oplus M''[k] \oplus \text{Pad} \oplus C[m] \quad (42)$$

$$= \bigoplus_{i \in \{1, \dots, m-1\} \setminus \{j, k\}} M[i] \oplus M[j] \oplus M[k] \oplus \text{Pad} \oplus C[m] \quad (43)$$

$$= \bigoplus_{i=1}^{m-1} M[i] \oplus \text{Pad} \oplus C[m] \quad (44)$$

$$= \Sigma, \quad (45)$$

since $M''[j] \oplus M''[k] = X[k] \oplus 2^j L \oplus X[j] \oplus 2^k L = M[j] \oplus M[k]$. Therefore, the decryption query will be accepted.

Generally, the adversary can arbitrarily choose the i -th plaintext blocks in the second forgery from $\mathcal{I}_i = \{X[j] \oplus 2^j L : j = 1, \dots, m-1\}$, and by choosing the last ciphertext block, the checksum can be set to (say) the previous checksum value in the encryption query so that the tags are matched. This is a powerful, almost-universal forgery.

However, this attack will not work when the receiver is stateful and detects a replay in nonce, as the first and second forgeries share the same nonce.

4.4 Almost Universal Forgery, Variant 2

To bypass the problem of stateful receiver mentioned above, we consider another variant of attack. This will create an almost universal forgery for a nonce that is never queried with a help of few additional queries.

Suppose the adversary \mathcal{A} performs the forgery attack of Section 4.2 using ℓ blocks. Then she obtains ℓ input-output pairs of E , $\mathcal{S} = \{(X[i], Y[i]) : i = 1, \dots, \ell\}$, where $E(X[i]) = Y[i]$ for $1 \leq i \leq \ell$. Let $\mathcal{X} = \{X[1], \dots, X[\ell]\}$ and

$\mathcal{Y} = \{Y[1], \dots, Y[\ell]\}$. We assume \mathcal{A} wants to create a forgery of the form (N'', A'', C'', T'') with $A'' = \varepsilon$, $|C''|_n = m$, and $N'' = X[k]$ for some $k \in \{1, \dots, \ell\}$, where N'' has never been used as a nonce for encryption and decryption queries. Such N'' avoids the detectable replay at the receiver. Since she already knows the initial mask $E(N'') = Y[k]$ and \mathcal{S} , the most blocks of C'' can be set so that the decrypted plaintext blocks are predictable. The problem is in the last block ($C''[m]$), where we need to make sure that $\mathbf{1en}(C''[m]) \oplus 2^m E(N'') \in \mathcal{X}$ so that the corresponding Pad'' is predictable. This is possible by making another encryption query with nonce $N^\S = \mathbf{1en}(C''[m]) \oplus 2^m E(N'')$ and performing the forgery with nonce N^\S to recover $E(N^\S)$. The detailed attack procedure is as follows. For simplicity, we fix $|C''[m]| = n$.

1. Perform the forgery attack of Section 4.2 to obtain \mathcal{X} and \mathcal{Y} of size ℓ .
2. Choose $N'' = X[k] \in \mathcal{X}$ for some k as the target nonce, and set $N^\S = \mathbf{1en}(C''[m]) \oplus 2^m E(N'') = \mathbf{1en}(0^n) \oplus 2^m Y[k]$.
3. Perform the forgery of Section 4.1 or 4.2 with nonce N^\S . Recover $E(N^\S)$.
4. Arbitrarily choose m elements of \mathcal{Y} , $Y[j_1], \dots, Y[j_{m-1}]$ and $Y[j_{m+1}]$, possibly with duplications. Let $Y[j_m] = E(N^\S)$.
5. Decrypt (N'', A'', C'', T'') such that

$$\begin{aligned} C''[i] &= Y[j_i] \oplus 2^i E(N'') = Y[j_i] \oplus 2^i Y[k], \text{ for } i = 1, \dots, m-1, \\ C''[m] &= M''[m] \oplus E(N^\S), \\ T'' &= Y[j_{m+1}] \end{aligned}$$

where

$$\begin{aligned} M''[i] &= X[j_i] \oplus 2^i E(N'') = X[j_i] \oplus 2^i Y[k], \\ M''[m] &= X[j_{m+1}] \oplus \bigoplus_{i=1}^{m-1} M''[i] \oplus 3 \cdot 2^m E(N''). \end{aligned}$$

This will be accepted as valid since the true tag is

$$E(\Sigma'' \oplus 3 \cdot 2^m E(N'')) = E\left(\bigoplus_{i=1}^{m-1} M''[i] \oplus M''[m] \oplus 3 \cdot 2^m E(N'')\right) \quad (46)$$

$$= E(X[j_{m+1}]) = Y[j_{m+1}]. \quad (47)$$

Here, Step 3 of our attack can be omitted if N^\S is already in \mathcal{X} by chance.

5 Universal Forgery with More Queries

In fact, the attacks of Sections 4.1 and 4.2 allow to recover $E(X)$ for any $X \in \{0, 1\}^n$, as N can be of any value in $\{0, 1\}^n$. This also implies the mask values of PMAC is also computable by mounting the forgery of with $N = 0^n$.

This allows us to mount a universal forgery attack, i.e., we first arbitrarily choose (N^*, A^*, M^*) and then create a forgery (N^*, A^*, C^*, T^*) that is accepted and the decryption oracle returns M^* . Since this is a simple extension of the almost-universal forgery attacks, we briefly describe the procedure.

1. Choose the target (N^*, A^*, M^*) .
2. Perform the attack of Section 4.1 or 4.2 (which we call the basic forgery attack) with nonce N^* , which involves an encryption query with nonce N^* but the corresponding pair of AD and plaintext is different from (A^*, M^*) . We obtain a set of I/O pairs of E including $(N^*, L^* = E(N^*))$.
3. If $A^* \neq \varepsilon$, also perform the basic forgery attack with nonce 0^n . We obtain a pair $(0^n, L^\$ = E(0^n))$.
4. With the knowledge of L^* and $L^\$$, we can identify all the inputs and outputs of E necessary to construct (C^*, T^*) . Let \mathcal{X}^* (\mathcal{Y}^*) be the set of such blockcipher inputs (outputs). Here, \mathcal{X}^* is determined by A^* , M^* , L^* and $L^\$$.
5. For each element $X \in \mathcal{X}^*$, if $E(X)$ is not already recovered, Perform the basic forgery attack with nonce X and recover $E(X)$. This will give \mathcal{Y}^* .
6. Build (C^*, T^*) using \mathcal{X} and \mathcal{Y} and perform a decryption query (N^*, A^*, C^*, T^*) .

To realize a universal choice of N^* , this attack assumes the stateless receiver described at Section 4.3, as the attacker uses N^* as a nonce in a successful forgery (Step 2) before the final forgery. However, by slightly relaxing the universality for nonce, we can cope with a stateful receiver. That is, we replace Step 2 with any forgery attacks that give certain input and output sets of E (\mathcal{X} and \mathcal{Y}), and N^* is chosen from \mathcal{X} except those used as nonce. Then, N^* will not appear as a nonce in the queries except the final decryption query.

The above attack requires around $a + m$ basic forgery attacks if $|A^*|_n = a$ and $|M^*|_n = m$, however, depending on the target values, a further optimization in terms of the number of queries may be possible.

6 Design Flaw of OCB2

The root of the flaw in OCB2 is in the instantiation of AE using XEX^* . For blockcipher E_K , let

$$\text{XEX}_E^{N,i,j}(X) \stackrel{\text{def}}{=} E(2^i L \oplus X) \oplus 2^i L, \quad (48)$$

$$\text{XE}_E^{N,i,j}(X) \stackrel{\text{def}}{=} E(2^i 3^j L \oplus X), \quad (49)$$

where $L = E(N)$ for nonce N , for $i = 1, 2, \dots$ and $j = 0, 1, \dots$. Here, j is always set to 0 for XEX . XEX^* unifies them by introducing one bit b to the tweak. That is,

$$\text{XEX}_E^{*,b,N,i,j}(X) = \begin{cases} \text{XEX}_E^{N,i,j}(X) & \text{if } b = 1; \\ \text{XE}_E^{N,i,j}(X) & \text{if } b = 0. \end{cases} \quad (50)$$

Decryption is trivially defined, and is never invoked when $b = 0$. Rog04 refers b to *tag*; not to be confused with the tag in the global interface of AE.

Suppose an encryption query (N, A, M) , where $A = \varepsilon$ and M is parsed as $(M[1], \dots, M[m])$, is given to OCB2. It encrypts M by using $\text{XEX}_E^{*,1,N,i,0}$ for $M[i]$ with $i = 1, \dots, m-1$, and $\text{XEX}_E^{*,0,N,m,0}$ for $M[m]$. The checksum, Σ , is encrypted by $\text{XEX}_E^{*,0,N,m,1}$ to create the (untruncated) tag.

In the proof of OCB2, we first apply the standard conversion from computational to information theoretic security [BDJR97] and focus on the security of OCB2 instantiated by an n -bit uniform random permutation (URP), \mathbb{P} , denoted by $\text{OCB2}_{\mathbb{P}}$. Then, the proof of $\text{OCB2}_{\mathbb{P}}$ has two main steps: the indistinguishability of $\text{XEX}_{\mathbb{P}}^*$, and the privacy and authenticity of AE^4 which replaces $\text{XEX}_{\mathbb{P}}^*$ in $\text{OCB2}_{\mathbb{P}}$ with an ideal primitive, a tweakable random permutation $\tilde{\mathbb{P}}$. The latter step is not relevant to our attacks.

For the first step, Rog04 proved that $\text{XEX}_{\mathbb{P}}^*$ is indistinguishable from $\tilde{\mathbb{P}}$ for any adversary who queries to both encryption and decryption of $\text{XEX}_{\mathbb{P}}^*$ and respects the semantics of tag b . More precisely, the conditions for the adversary are as follows.

Definition 1. *We say an adversary querying XEX^* is tag-respecting when*

1. $\text{XEX}^{*,0,N,i,j}$ is only queried in encryption queries for any (N, i, j) ;
2. Once $\text{XEX}^{*,b,N,i,j}$ is queried in either encryption or decryption, then it is not allowed to query $\text{XEX}^{*,1-b,N,i,j}$, for any (N, i, j) .

Let $\Theta\text{CB2}_{\tilde{E}}$ be the mode of operations of TBC \tilde{E}_K which has the same interface as XEX_E^* . The pseudocode is shown in Fig. 4. Then, $\Theta\text{CB2}_{\text{XEX}_E^*}$ is equivalent to OCB2_E .

Let $\tilde{\mathbb{P}}$ be TURP which has the same interface as XEX^* . Rog04 showed that, for any privacy-adversary \mathcal{A} and authenticity-adversary \mathcal{A}_{\pm} ,

$$\text{Adv}_{\text{OCB2}_{\mathbb{P}}}^{\text{priv}}(\mathcal{A}) = \text{Adv}_{\Theta\text{CB2}_{\text{XEX}_{\mathbb{P}}^*}}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tprp}}(\mathcal{B}) + \text{Adv}_{\Theta\text{CB2}_{\tilde{\mathbb{P}}}}^{\text{priv}}(\mathcal{A}), \quad (51)$$

$$\text{Adv}_{\text{OCB2}_{\mathbb{P}}}^{\text{auth}}(\mathcal{A}_{\pm}) = \text{Adv}_{\Theta\text{CB2}_{\text{XEX}_{\mathbb{P}}^*}}^{\text{auth}}(\mathcal{A}_{\pm}) \leq \text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tsprp}}(\mathcal{B}_{\pm}) + \text{Adv}_{\Theta\text{CB2}_{\tilde{\mathbb{P}}}}^{\text{auth}}(\mathcal{A}_{\pm}) \quad (52)$$

hold for some CPA-adversary \mathcal{B} and CCA-adversary \mathcal{B}_{\pm} , which are tag-respecting and can simulate the privacy and the authenticity games involving $\Theta\text{CB2}_{\text{XEX}_{\mathbb{P}}^*}$ and \mathcal{A} and \mathcal{A}_{\pm} , respectively. From Rog04, we have

$$\text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tprp}}(\mathcal{B}) \leq \frac{4.5q^2}{2^n}, \text{ and } \text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tsprp}}(\mathcal{B}_{\pm}) \leq \frac{9.5q^2}{2^n} \quad (53)$$

for any \mathcal{B} and \mathcal{B}_{\pm} that are tag-respecting and use at most q queries. The last terms of (51) and (52) are proved to be almost ideally small: zero for privacy and $2^{n-\tau}/(2^n - 1)$ for authenticity with single decryption query.

⁴ An equivalent mode for OCB3 is called ΘCB3 [KR11].

Algorithm $\Theta\text{CB2}.\mathcal{E}_{\tilde{E}}(N, A, M)$	Algorithm $\Theta\text{CB2}.\mathcal{D}_{\tilde{E}}(N, A, C, T)$
<ol style="list-style-type: none"> 1. $(M[1], \dots, M[m]) \xleftarrow{n} M$ 2. for $i = 1$ to $m - 1$ 3. $C[i] \leftarrow \tilde{E}^{*,1,N,i,0}(M[i])$ 4. $\text{Pad} \leftarrow \tilde{E}^{*,0,N,m,0}(\text{len}(M[m]))$ 5. $C[m] \leftarrow M[m] \oplus \text{msb}_{ M[m] }(\text{Pad})$ 6. $\Sigma \leftarrow C[m] \parallel 0^* \oplus \text{Pad}$ 7. $\Sigma \leftarrow M[1] \oplus \dots \oplus M[m-1] \oplus \Sigma$ 8. $T \leftarrow \tilde{E}^{*,0,N,m,1}(\Sigma)$ 9. return (C, T) 	<ol style="list-style-type: none"> 1. $(C[1], \dots, C[m]) \xleftarrow{n} C$ 2. for $i = 1$ to $m - 1$ 3. $M[i] \leftarrow (\tilde{E}^{*,1,N,i,0})^{-1}(C[i])$ 4. $\text{Pad} \leftarrow \tilde{E}^{*,0,N,m,0}(\text{len}(C[m]))$ 5. $M[m] \leftarrow C[m] \oplus \text{msb}_{ C[m] }(\text{Pad})$ 6. $\Sigma \leftarrow C[m] \parallel 0^* \oplus \text{Pad}$ 7. $\Sigma \leftarrow M[1] \oplus \dots \oplus M[m-1] \oplus \Sigma$ 8. $T^* \leftarrow \tilde{E}^{*,0,N,m,1}(\Sigma)$ 9. if $T = T^*$ return M 10. else return \perp

Fig. 4. Algorithms of ΘCB2 . For simplicity, $\tau = n$ and $A = \varepsilon$.

The privacy bound is obtained from (53) and (51). However, to derive the authenticity bound, we need to identify \mathcal{B}_{\pm} that can simulate \mathcal{A}_{\pm} , where \mathcal{A}_{\pm} must compute the decryption of ΘCB2 , even with single decryption query⁵. Depending on \mathcal{A}_{\pm} , there are cases that no tag-respecting \mathcal{B}_{\pm} can simulate \mathcal{A}_{\pm} . For example, let us assume that \mathcal{A}_{\pm} first queries (N, A, M) of $|M| = 2n$ to the encryption oracle and then queries (N', A', C', T') to the decryption oracle, where $N' = N$, $A' = \varepsilon$ and $|C'| = n$, as well as the attack of Section 4.1. Then, \mathcal{B}_{\pm} who simulates \mathcal{A}_{\pm} first queries to $\text{XEX}^{*,1,N,1,0}$ and $\text{XEX}^{*,0,N,2,0}$ and $\text{XEX}^{*,0,\bar{N},2,1}$. For the second query, it queries to $\text{XEX}^{*,0,N,1,0}$ and $\text{XEX}^{*,0,N,1,1}$. Thus both $\text{XEX}^{*,1,N,1,0}$ and $\text{XEX}^{*,0,N,1,0}$ are queried, which implies a violation of the second condition of Definition 1. Consequently, the authenticity proof of Rog04 does not work, hence our attacks. At the same time, this also implies that the privacy attack, i.e. distinguishing the ciphertext from random using only encryption queries, is not possible.

7 Practicality, and Applicability to Others

ALMOST KNOWN-PLAINTEXT QUERY. The first encryption query is almost a known-plaintext query in that it works with a known plaintext except the last two blocks being $\text{len}(0^n) \parallel X$ for some known $X \in \{0, 1\}^n$, where $\text{len}(0^n) = 0^{120}10^7$ in the reference code. This can happen (e.g.) when a certain padding is already applied to the plaintext.

⁵ Rog04 defines the authenticity notion in the game that the adversary queries to the encryption oracle then outputs a query to the decryption oracle, but the response is not returned. The decryption oracle is not involved in the game and the success or failure of the forgery is determined outside the game. This definition itself is essentially the same as (1), and has no problem. However, because the adversary's final output does not tell whether the adversary wins or loses, we do not know how to to apply a hybrid argument of (52) using this definition.

OTHER OCB VERSIONS. Our attacks are only applicable to OCB2. For OCB1, the last block is encrypted by XE with a clearly separated mask. For OCB3, the last block is encrypted by XEX when it is n bits and otherwise by XE with a mask separated from those used by XEX.

OTHER DESIGNS BASED ON OCB. We have not found other AE algorithms based on OCB that could be affected by our attacks. OTR [Min14] is an inverse-free (for the absence of the blockcipher decryption in the scheme) parallelizable AE having a similar structure as OCB. As it only uses XE for the whole process, it is safe from our attacks. OPP [GJMN16] is a permutation-based AE based on OCB. It always uses XEX, or more precisely, a variant of XPX [Men16], because otherwise an offline permutation inverse query easily breaks the scheme. It is safe because of this consistent use of XPX.

STANDARDS AND LIBRARIES. ISO/IEC 19772:2009 specifies OCB2 and thus is affected. SJCL ⁶ is a well-known Javascript cryptography library and it implements OCB2, though we have not tested our attacks on it⁷.

8 Fixing OCB2

We suggest several ways to fix OCB2.

USE XEX FOR THE LAST BLOCK. The simplest way is to use XEX for the last block. We call it OCB2f (f for fix). Its pseudocode is obtained by just changing line 5 of OCB2.E_E and OCB2.D_E in Fig. 1 to

$$\begin{aligned} \text{Pad} &\leftarrow 2^m L \oplus E(2^m L \oplus \text{len}(M[m])), \text{ and} \\ \text{Pad} &\leftarrow 2^m L \oplus E(2^m L \oplus \text{len}(C[m])). \end{aligned}$$

OCB2f can be regarded as a mode of XEX^* , since the tweak spaces of XE and XEX in OCB2f are distinct. Specifically, we define $\Theta\text{CB2f}_{\tilde{E}}$ as the mode of \tilde{E}_K obtained by changing line 4 of $\Theta\text{CB2.E}_{\tilde{E}}$ in Fig. 4 to

$$\text{Pad} \leftarrow \tilde{E}^{*,1,N,m,0}(\text{len}(M[m])).$$

$\Theta\text{CB2f.D}_{\tilde{E}}$ is defined accordingly. Then $\Theta\text{CB2f}_{\tilde{E}}$ is equivalent to OCB2f_E if \tilde{E}_K is XEX_E^* .

The security bounds of OCB2f are the same as those claimed for OCB2:

⁶ <http://bitwiseshiftleft.github.io/sjcl/>

⁷ We need to natively use SJCL for attacks. The demo page of SJCL (<https://bitwiseshiftleft.github.io/sjcl/demo/>) cannot be used since (to our understanding) it only accepts ASCII strings for plaintext, while our attacks need a plaintext that cannot be encoded to a valid ASCII string.

Theorem 1. For $\text{OCB2f}_{\mathbb{P}}$, let \mathcal{A} be the privacy adversary and let \mathcal{A}_{\pm} be the authenticity adversary with single decryption query. Let σ_{priv} (σ_{auth}) be the total number of \mathbb{P} calls needed in the privacy (authenticity) game. Then we have

$$\begin{aligned} \text{Adv}_{\text{OCB2f}_{\mathbb{P}}}^{\text{priv}}(\mathcal{A}) &= \text{Adv}_{\text{OCB2f}_{\text{XEX}_{\mathbb{P}}^*}}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tsprp}}(\mathcal{B}) + \text{Adv}_{\Theta\text{CB2f}_{\mathbb{P}}}^{\text{priv}}(\mathcal{A}) \\ &\leq \frac{4.5\sigma_{\text{priv}}^2}{2^n} \end{aligned} \quad (54)$$

$$\begin{aligned} \text{Adv}_{\text{OCB2f}_{\mathbb{P}}}^{\text{auth}}(\mathcal{A}_{\pm}) &= \text{Adv}_{\text{OCB2f}_{\text{XEX}_{\mathbb{P}}^*}}^{\text{auth}}(\mathcal{A}_{\pm}) \leq \text{Adv}_{\text{XEX}_{\mathbb{P}}^*}^{\text{tsprp}}(\mathcal{B}_{\pm}) + \text{Adv}_{\Theta\text{CB2f}_{\mathbb{P}}}^{\text{auth}}(\mathcal{A}_{\pm}) \\ &\leq \frac{9.5\sigma_{\text{auth}}^2}{2^n} + \frac{2^{n-\tau}}{2^n - 1} \end{aligned} \quad (55)$$

hold for some tag-respecting \mathcal{B} using σ_{priv} queries and tag-respecting \mathcal{B}_{\pm} using σ_{auth} queries.

Intuitively, the security of OCB2f is proved because OCB2f is ΘCB2f using \tilde{E} instantiated by XEX^* , and the privacy and authenticity games do not force the adversary to violate the access rules (Definition 1). Combining (53) and the proofs of $\Theta\text{CB2}_{\mathbb{P}}$ from Rog04 with minor changes gives the desired results. In slightly more detail, the proof for privacy is trivially obtained in the same manner to that of Rog04. For the authenticity bound, since this is again almost the same as the proof of OCB2 , we provide a proof sketch that focuses on the important difference from the proof of OCB2 . For simplicity, we ignore AD (assuming it is empty for all queries) and $\tau = n$. A full proof will appear in the full version.

Proof. (Sketch) Suppose \mathcal{A}_{\pm} queries $(N_1, M_1), \dots, (N_q, M_q)$ to the encryption oracle and (N', C', T') to the decryption oracle. For $1 \leq i \leq q$, let (C_i, T_i) be the return value for the query (N_i, M_i) . Without loss of generality, we assume \mathcal{A}_{\pm} tries to forge after all encryption queries. We need to consider the following cases.

Case 1: $N' \notin \{N_1, \dots, N_q\}$

Case 2: $N' = N_i$ for some $i \in \{1, \dots, q\}$, $|C'|_n = |C_i|_n$

In these two cases, the proof is almost the same as $\Theta\text{CB2}_{\mathbb{P}}$ and the adversary can guess the valid tag, T^* , with probability at most $2^{n-\tau}/(2^n - 1)$.

Case 3: $N' = N_i$ for some $i \in \{1, \dots, q\}$, $|C'|_n < |C_i|_n$

Let $|C'|_n = c'$, $|C_i|_n = c$. In this case, unlike $\Theta\text{CB2}_{\mathbb{P}}$, $\tilde{E}^{*,1,N',c',0}$ is invoked to decrypt $C'[c']$, while $\tilde{E}^{*,1,N',c',0}$ (where the underlying blockcipher is \mathbb{P}) has also been used in an encryption query. Nevertheless, the adversary has to guess T^* without access to $\tilde{E}^{*,0,N',c',1}$ since $c' \neq c$. Since $\tilde{E}^{*,0,N',c',1}$ has never queried before, the probability of $T^* = T'$ is $1/2^\tau$.

Case 4: $N' = N_i$ for some $i \in \{1, \dots, q\}$, $|C'|_n > |C_i|_n$

In the same manner to the above case, the forging probability is $1/2^\tau$.

Use XEX^+ . Minematsu and Matsushima [MM09] proposed an extension of XEX^* called XEX^+ . This allows to use plain blockcipher calls in combination with XEX and XE . MM09 suggested how to use XEX^+ to instantiate a variant of OCB, where the last message block is encrypted by an unmasked blockcipher. This variant of OCB is not affected by our attacks and provably secure.

9 Conclusions

We have presented practical forgery attacks against OCB2, a high-profile, ISO-standard authenticated encryption scheme. This was possible due to the discrepancy between the proof of OCB2 and the actual construction, in particular about the interpretation of OCB2 as a mode of TBC which combines XEX and XE . While the latest OCB3 has a superior software performance from the previous ones, and is clearly recommended by the designers, we think OCB2 is still quite influential for its simple description and the sophisticated, modular design based on TBC. Our attacks show that, while the approach introduced by Rog04 is invaluable, we could not directly derive a secure AE from it without applying a fix.

We comment that, due to the errors in the proofs, provably-secure schemes can be broken, or schemes still remain secure but the proofs need to be fixed. Even if we limit our focus to AE, we have many examples, such as NSA’s Dual CTR [Rog04d], [DGW01], EAX-prime [MLMI13], GCM [IOM12], and some of the CAESAR submissions [Nan14], [BS16], [SMAP15] and more. We believe our work to emphasize the quality of security proofs and their active verifications.

Acknowledgements

The authors would like to thank Tetsu Iwata for useful comments.

References

- [ABL⁺14] Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. How to Securely Release Unverified Plaintext in Authenticated Encryption. In *ASIACRYPT (1)*, volume 8873 of *Lecture Notes in Computer Science*, pages 105–125. Springer, 2014.
- [ADL17] Tomer Ashur, Orr Dunkelman, and Atul Luykx. Boosting Authenticated Encryption Robustness with Minimal Modifications. In *CRYPTO (3)*, volume 10403 of *Lecture Notes in Computer Science*, pages 3–33. Springer, 2017.
- [AY13] Kazumaro Aoki and Kan Yasuda. The Security of the OCB Mode of Operation without the SPRP Assumption. In *ProvSec*, volume 8209 of *Lecture Notes in Computer Science*, pages 202–220. Springer, 2013.
- [BC09] John Black and Martin Cochran. MAC Reforgeability. In *FSE*, volume 5665 of *Lecture Notes in Computer Science*, pages 345–362. Springer, 2009.

- [BDJR97] Mihir Bellare, Anand Desai, E. Jokipii, and Phillip Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *FOCS '97*, pages 394–403. IEEE Computer Society, 1997.
- [BN17] Ritam Bhaumik and Mridul Nandi. Improved Security for OCB3. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in Computer Science*, pages 638–666. Springer, 2017.
- [BR02] John Black and Phillip Rogaway. A Block-Cipher Mode of Operation for Parallelizable Message Authentication. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 384–397. Springer, 2002.
- [BRW04] Mihir Bellare, Phillip Rogaway, and David A. Wagner. The EAX Mode of Operation. In *FSE*, volume 3017 of *Lecture Notes in Computer Science*, pages 389–407. Springer, 2004.
- [BS16] Raphael Bost and Olivier Sanders. Trick or Tweak: On the (In)security of OTR’s Tweaks. In *ASIACRYPT (1)*, volume 10031 of *Lecture Notes in Computer Science*, pages 333–353, 2016.
- [DGW01] Pompiliu Donescu, Virgil D. Gligor, and David Wagner. A Note on NSA’s Dual Counter Mode of Encryption, 2001. <http://www.cs.berkeley.edu/~daw/papers/dcm-prelim.ps/>.
- [Fer02] Niels Ferguson. Collision attacks on OCB. Comments to NIST, 2002. <https://csrc.nist.gov/CSRC/media/Projects/Block-Cipher-Techniques/documents/BCM/Comments/general-comments/papers/Ferguson.pdf/>.
- [FLLW17] Christian Forler, Eik List, Stefan Lucks, and Jakob Wenzel. Reforgeability of Authenticated Encryption Schemes. In *ACISP (2)*, volume 10343 of *Lecture Notes in Computer Science*, pages 19–37. Springer, 2017.
- [GJMN16] Robert Granger, Philipp Jovanovic, Bart Mennink, and Samuel Neves. Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In *EUROCRYPT (1)*, volume 9665 of *Lecture Notes in Computer Science*, pages 263–293. Springer, 2016.
- [IK03] Tetsu Iwata and Kaoru Kurosawa. OMAC: One-Key CBC MAC. In Thomas Johansson, editor, *FSE*, volume 2887 of *Lecture Notes in Computer Science*, pages 129–153. Springer, 2003.
- [IOM12] Tetsu Iwata, Keisuke Ohashi, and Kazuhiko Minematsu. Breaking and Repairing GCM Security Proofs. In *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 31–49. Springer, 2012.
- [ISO09] Information Technology - Security techniques - Authenticated encryption, ISO/IEC 19772:2009. International Standard ISO/IEC 19772, 2009.
- [KR11] Ted Krovetz and Phillip Rogaway. The Software Performance of Authenticated-Encryption Modes. In *FSE*, volume 6733 of *Lecture Notes in Computer Science*, pages 306–327. Springer, 2011.
- [LRW02] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2002.
- [Men16] Bart Mennink. XPX: Generalized Tweakable Even-Mansour with Improved Security Guarantees. In *CRYPTO (1)*, volume 9814 of *Lecture Notes in Computer Science*, pages 64–94. Springer, 2016.
- [Min14] Kazuhiko Minematsu. Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In *Eurocrypt*, 2014. to appear.

- [MLMI13] Kazuhiko Minematsu, Stefan Lucks, Hiraku Morita, and Tetsu Iwata. Attacks and Security Proofs of EAX-Prime. In *FSE*, volume 8424 of *Lecture Notes in Computer Science*, pages 327–347. Springer, 2013.
- [MM09] Kazuhiko Minematsu and Toshiyasu Matsushima. Generalization and Extension of XEX* Mode. *IEICE Transactions*, 92-A(2):517–524, 2009.
- [Nan14] Mridul Nandi. Forging Attacks on Two Authenticated Encryption Schemes COBRA and POET. In *ASIACRYPT (1)*, volume 8873 of *Lecture Notes in Computer Science*, pages 126–140. Springer, 2014.
- [RBBK01] Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: a block-cipher mode of operation for efficient authenticated encryption. In *ACM Conference on Computer and Communications Security*, pages 196–205. ACM, 2001.
- [RFC14] The OCB Authenticated-Encryption Algorithm. IRTF RFC 7253, 2014.
- [Rog02] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM Conference on Computer and Communications Security*, pages 98–107. ACM, 2002.
- [Rog04a] Phillip Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In *ASIACRYPT*, volume 3329 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2004.
- [Rog04b] Phillip Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. Full version of [Rog04a], 2004. available from <http://www.cs.ucdavis.edu/~rogaway/papers/>.
- [Rog04c] Phillip Rogaway. Nonce-Based Symmetric Encryption. In *FSE*, volume 3017 of *Lecture Notes in Computer Science*, pages 348–359. Springer, 2004.
- [Rog04d] Phillip Rogaway. On the Role Definitions in and Beyond Cryptography. In *ASIAN*, volume 3321 of *Lecture Notes in Computer Science*, pages 13–32. Springer, 2004.
- [SMAP15] Willem Schroé, Bart Mennink, Elena Andreeva, and Bart Preneel. Forgery and Subkey Recovery on CAESAR Candidate iFeed. In *SAC*, volume 9566 of *Lecture Notes in Computer Science*, pages 197–204. Springer, 2015.
- [SWZ12] Zhelei Sun, Peng Wang, and Liting Zhang. Collision Attacks on Variant of OCB Mode and Its Series. In *Inscrypt*, volume 7763 of *Lecture Notes in Computer Science*, pages 216–224. Springer, 2012.
- [VV18] Serge Vaudenay and Damian Vizár. Can Caesar Beat Galois? - Robustness of CAESAR Candidates Against Nonce Reusing and High Data Complexity Attacks. In *ACNS*, volume 10892 of *Lecture Notes in Computer Science*, pages 476–494. Springer, 2018.

A Code Example for Minimal Attack

1. Retrieve OCB2 reference code from <http://web.cs.ucdavis.edu/~rogaway/ocb/code-2.0.htm> and AES reference code (`rijndael-alg-fst.c`).
2. Change the main routine of `ocb.c` to the following snippet:

```
int
main(void)
{
    block nbits = {0};
```

```

block N = {0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15};
block K = {0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15};
byte Hatk[0]; /* any */
byte Matk[32] = {0};
byte Catk[32] = {0};
byte Hforge[0]; /* must be empty */
byte Cforge[16] = {0};
byte Mforge[16] = {0};
block T,Tatk,Tforge;
int res;
ocb_state *state;

/* Test for the minimal attack */
printf("Test for minimal attack \n");
state = ocb_init((byte *)"abcdefghijklmnop",
sizeof(T),sizeof(N),AES128);
memset(nbits,0,sizeof(block));
nbits[sizeof(block)-1] = 16 * 8; /* 128 bits */
memcpy(Matk,nbits,sizeof(block));
printf("Encryption query:\n");
pbuf(N,16, " Nonce");
pbuf(Matk,32, " Plaintext");
pbuf(Hatk,sizeof(Hatk), " AD");
ocb_provide_header(state,Hatk,sizeof(Hatk));
ocb_encrypt(state,N,Matk,sizeof(Matk),Catk,Tatk);
pbuf(Catk,32, " Ciphertext");
pbuf(Tatk,16, " Tag");
printf("Decryption query (forgery):\n");
memcpy(Cforge, Catk, 16);
xor_block(Cforge, Cforge, nbits);
pbuf(N,16, " Forged Nonce (the same as encryption)");
pbuf(Hforge, sizeof(Hforge), " Forged AD (empty)");
pbuf(Cforge,16, " Forged Ciphertext");
memcpy(Tforge, Matk+16, 16);
xor_block(Tforge, Tforge, Catk+16);
pbuf(Tforge, 16, " Forged Tag");
ocb_provide_header(state,Hforge,sizeof(Hforge));
res = ocb_decrypt(state,N,Cforge,sizeof(Cforge),Tforge,
Mforge);
ocb_zeroize(state);
printf("Tags match: %i.\n", res); /* 1 is "matched" */
pbuf(Mforge, sizeof(Mforge), " Forged Plaintext");
return 0;
}

```