

AES-OTR v1

Designer/Submitter : Kazuhiko Minematsu
NEC Corporation, Japan
Contact address : k-minematsu@ah.jp.nec.com

2014.03.14

1 Specification

OTR, which stands for Offset Two-Round, is a blockcipher mode of operation to realize an authenticated encryption with associated data (AEAD), proposed by Minematsu [20, 21]. Our CAESAR submission, AES-OTR, is an instantiation of OTR based on AES blockcipher with an additional version for associated data processing.

1.1 Parameters

AES-OTR has the following four parameters.

- Key length : 16 bytes (128 bits), 24 bytes (192 bits), 32 bytes (256 bits).
- Nonce length : 1 byte (8 bits) to 15 bytes (120 bits).
- Tag length : 4 bytes (32 bits) to 16 bytes (128 bits).
- Associated Data Processing (ADP) : Parallel (p) or Serial (s).

1.2 Recommended Parameter Set

- Primary recommended parameter set `aes128otrpv1` : 16-byte key for AES-128, 12-byte nonce, 16-byte tag, and parallel ADP.
- Secondary recommended parameter set `aes128otrsv1` : 16-byte key for AES-128, 12-byte nonce, 16-byte tag, and serial ADP.
- Third recommended parameter set `aes256otrpv1` : 32-byte key for AES-256, 12-byte nonce, 16-byte tag, and parallel ADP.
- Fourth recommended parameter set `aes256otrsv1` : 32-byte key for AES-256, 12-byte nonce, 16-byte tag, and serial ADP.

1.3 Algorithm

■ **Basic Notations.** Let $\{0, 1\}^*$ and $\{0, 1\}^{8^*}$ be the sets of all finite-length binary strings and all finite-length byte strings, both including the empty string ε . Here $\{0, 1\}^{8^*} \subset \{0, 1\}^*$. For binary string $X \in \{0, 1\}^*$, its bit length is denoted by $|X|$. We define $|X|_a \stackrel{\text{def}}{=} \max\{|X|/a, 1\}$. For $X = \varepsilon$ we have $|X|_a = 1$ for any $a \geq 1$ and $|X| = 0$. For any $X, Y \in \{0, 1\}^*$ we write $X\|Y$ or simply XY to denote the concatenation of X and Y . A sequence of a zeros is denoted by 0^a . For $X \in \{0, 1\}^*$, we write $(X[1], \dots, X[x]) \stackrel{\leftarrow}{=} X$ to denote the n -bit block partitioning of X . That is, $X[1]\|X[2]\|\dots\|X[x] = X$ where $x = |X|_n$, and $|X[i]| = n$ for $i < x$ and $|X[x]| \leq n$. If $X = \varepsilon$ we have $X[1] \stackrel{\leftarrow}{=} X$ with $X[1] = \varepsilon$, for any $n \geq 1$. The sequence of first c bits of $X \in \{0, 1\}^*$ is denoted by $\text{msb}_c(X)$. We have $\text{msb}_0(X) = \varepsilon$ for any X . For $X \in \{0, 1\}^*$ with $0 \leq |X| < n$, \underline{X} denotes the 10^* padding written as $X\|10^{n-|X|-1}$. When $|X| = n$, we have $\underline{X} = X$, and when $X = \varepsilon$ we have $\underline{X} = 10^{n-1}$.

Unless otherwise stated, we assume a variable n to mean 128. We also assume E_K to mean AES encryption function with key K , for some fixed $|K| \in \{128, 192, 256\}$. For $X, Y \in \{0, 1\}^n$, $Y = E_K(X)$ means that Y is the ciphertext obtained by AES encryption for plaintext X with key K . We may simply write E to denote E_K .

When we write $\Pr[X : Y]$, it means the probability of event Y defined over the experiment X . We may omit X if it is obvious.

If adversary \mathcal{A} accesses an oracle O and returns a value x as a final output, we write $\mathcal{A}^O \Rightarrow x$ to denote the corresponding event.

The specification of our scheme assumes big endian.

■ **Input and Output.** The encryption algorithm of AES-OTR using blockcipher (here AES) E , tag bit length τ , with parallel ADP is written as $\text{OTR-}\mathcal{E}_{E, \tau, p}$. When ADP is serial we similarly write $\text{OTR-}\mathcal{E}_{E, \tau, s}$.

to denote the encryption algorithm.

Both algorithms take the following byte sequences.

- Key $K \in \mathcal{K}_{ae}$,
- Nonce (a public message number) $N \in \mathcal{N}_{ae}$,
- Associated Data (AD) $A \in \mathcal{A}_{ae}$, and
- Plaintext $M \in \mathcal{M}_{ae}$.

The output is a byte sequence

- Ciphertext $(C, T) \in \mathcal{M}_{ae} \times \mathcal{T}_{ae}$,

where a ciphertext is a concatenation of unauthenticated ciphertext C fulfilling $|C| = |M|$ and a tag $T \in \mathcal{T}_{ae}$. There is no secret message number, hence the length of secret message number is assumed to be zero. Here $\mathcal{N}_{ae} = \{0, 1\}^{|N|}$, $\mathcal{K}_{ae} = \{0, 1\}^{|K|}$ with fixed $|N|$ and $|K|$ satisfying $|N| \in \{8, 16, \dots, 120\}$ and $|K| \in \{128, 192, 256\}$, and $\mathcal{M}_{ae} = \mathcal{A}_{ae} = \{0, 1\}^{8*}$, but we restrict these lengths as $|M|_8 \leq 2^{64}$ and $|A|_8 \leq 2^{64}$ (so 2^{60} 16-byte blocks). Both M and A can be empty (possibly simultaneously, though artificial). Tag space \mathcal{T}_{ae} is $\{0, 1\}^\tau$ with fixed $\tau \in \{32, 40, \dots, 128\}$.

The decryption algorithm of AES-OTR using blockcipher E , and tag bit length τ , with parallel ADP is written as $\text{OTR-}\mathcal{D}_{E,\tau,p}$. Similarly we write $\text{OTR-}\mathcal{D}_{E,\tau,s}$ to denote the decryption algorithm with serial ADP. Both algorithms take the following byte sequences.

- Key $K \in \mathcal{K}_{ae}$,
- Nonce (a public message number) $N \in \mathcal{N}_{ae}$,
- Associated Data (AD) $A \in \mathcal{A}_{ae}$, and
- Ciphertext $(C, T) \in \mathcal{M}_{ae} \times \mathcal{T}_{ae}$,

where a ciphertext is a concatenation of unauthenticated ciphertext C and a tag $T \in \mathcal{T}_{ae}$. If the result is determined as authentic, the output is a byte sequence

- Plaintext $M \in \mathcal{M}_{ae}$ satisfying $|M| = |C|$,

otherwise a rejection symbol \perp .

We remark that parallel ADP version is identical to the algorithm of [20,21], while serial ADP version has a different AD processing.

■ **Algorithms for Parallel ADP.** The algorithms of $\text{OTR-}\mathcal{E}_{E,\tau,p}$ and $\text{OTR-}\mathcal{D}_{E,\tau,p}$ are described in Fig. 1, and $\text{OTR-}\mathcal{E}_{E,\tau,p}$ is illustrated in Fig 2. In Fig. 1 the functions $\text{OTR-}\mathcal{E}_{E,\tau,p}$ and $\text{OTR-}\mathcal{D}_{E,\tau,p}$ are further decomposed into the encryption and decryption cores, EF_E , DF_E , and the authentication core, AF_E . Below we provide a brief explanation on Fig. 1. For $\text{OTR-}\mathcal{E}_{E,\tau,p}$, we first partition a plaintext M into n -bit blocks, i.e. $(M[1], \dots, M[m]) \stackrel{\leftarrow}{=} M$ (Fig. 1, Line 3 of EF_E). For presentation purpose we also assume a partition into $2n$ -bit blocks, which we call chunks, denoted as $(\mathbf{M}[1], \dots, \mathbf{M}[\ell]) \stackrel{\leftarrow}{=} M$. For each $i < \ell$, the i -th chunk $\mathbf{M}[i] = (M[2i-1], M[2i])$ is encrypted by a two-round Feistel permutation with masks as

$$C[2i-1] = E_K(2^{i-1}L \oplus M[2i-1]) \oplus M[2i], \quad (1)$$

$$C[2i] = E_K(2^{i-1}L \oplus \delta \oplus C[2i-1]) \oplus M[2i-1], \quad (2)$$

where $\delta = E_K(\underline{N})$ and $L = 4\delta$, to form the i -th ciphertext chunk, $\mathbf{C}[i] = (C[2i-1], C[2i])$ (Fig. 1, Lines 5 and 6 of EF_E). For the last chunk (which may be partial), a variation of the above procedure is applied depending on the length of the last chunk. When it is more than n bits we apply a variant of two-round Feistel (Fig. 1, Lines 10 to 13 of EF_E), and otherwise a variant of CTR mode is applied (Fig. 1, Line 17 of EF_E). This yields the unauthenticated ciphertext, $\mathbf{C} = (\mathbf{C}[1], \dots, \mathbf{C}[\ell])$. For computing the tag, the check sum is computed by taking the sum (XOR) of right halves of the plaintext chunks, $\mathbf{M}[1], \dots, \mathbf{M}[\ell-1]$, i.e., even plaintext blocks (Fig. 1, Line 7 of EF_E). The check sum also involves a value from the final chunk $\mathbf{M}[\ell]$ which depends on the chunk length (Fig. 1, Lines 14 and 18 of EF_E). Then the check sum is given to the AES encryption with a dedicated mask, depending on the length of

last chunk. The result is the untruncated plaintext tag, TE (Fig. 1, Lines 19 to 20 of EF_E).

In parallel to the above procedure, an associated data A is given to AF_E . It uses a parallelizable PRF, which is a variant of PMAC [23]. Then we obtain the untruncated associated data tag, TA . This is when A is non-empty. When $A = \varepsilon$, we have $TA = 0^n$.

Finally, we take a sum of TE and TA and apply the truncation to obtain the tag T . By concatenating C and T we obtain the ciphertext (C, T) .

For $OTR-\mathcal{D}_{E,\tau,p}$, we first perform the inverse operation of (1) and (2) and the inverse operation for the last chunk, to obtain the (unverified) plaintext $M = (\mathbf{M}[1], \dots, \mathbf{M}[\ell])$, and compute the check sum to obtain TE . By computing TA and taking a sum with TE , we obtain the unverified tag \hat{T} . If $T = \hat{T}$ we decide ciphertext (C, T) is authentic and the final output is M . Otherwise, the output is \perp .

■ **Algorithms for Serial ADP.** For serial ADP, the procedures are very similar to the case of parallel ADP, except the handling of associated data. We have Fig. 3 to show the algorithms of $OTR-\mathcal{E}_{E,\tau,s}$ and $OTR-\mathcal{D}_{E,\tau,s}$, and $OTR-\mathcal{E}_{E,\tau,s}$ is also illustrated in Fig 4. The algorithms of Fig. 3 are further decomposed into encryption core $EF-S_E$, decryption core $DF-S_E$, and authentication core $AF-S_E$. For $EF-S_E$ and $DF-S_E$ the pseudocodes are the same as EF_E and DF_E except line 2. The difference in line 2 is the addition of TA to $E_K(N)$ to yield δ . Here TA is generated by $AF-S_E$ taking A and $AF-S_E$ is a variant of CMAC [4], also known as OMAC [15].

2 Security Goals

We consider the standard security notions for nonce-based AEs [24], see Section 3. Nonce (public message number) must be unique for all encryption queries, however the same nonce can be used for an encryption and a decryption query, or for two decryption queries. We basically do not claim any security when nonce is reused for encryption queries, though the serial ADP version slightly relaxes the condition (See the last of this section).

The security goal of our cipher for confidentiality (or privacy) of plaintext is shown in Table 1, and that for integrity (or authenticity) of plaintext, associated data, and public message number is shown in Table 2. We believe that, the security goal of our cipher is mostly explained by the security bounds presented in Section 3. These tables come from a natural interpretation of these bounds, where the variables in the tables denote the required workload of an adversary to break the cipher, in logarithm base 2. If one of the variables reaches the required amount it means a break. For Table 1, Data and Time denote σ_{priv} and t' in the privacy bounds (Theorems 3.1 and Theorem 3.3). For Table 2, Data, Verify, and Time denote σ_{auth} , q_v , and t' in the authenticity bounds (Theorem 3.2 and Theorem 3.4). The inclusion of Verify in Table 2 is redundant since σ_{auth} includes q_v and τ is 128 for the parameter sets here, however, we explicitly present it to make clear that $q_v \ll \tau$ is a security condition, which is meaningful if τ is small. Strictly speaking, the Data figures in the tables should be slightly smaller with respect to the bounds, say Data 64 in Table 1 of `aes128otrpv1` and `aes256otrpv1` is $64 - \log_2(c)$ with $c = \sqrt{6}$ from Theorem 3.1.

As an additional security goal, we claim that the length of nonce can be changed without key renewal, provided all nonces are unique for all encryptions, for both parallel and serial ADP versions. This setting does not change the security bounds of Section 3. In addition, the security of serial ADP holds as far as a pair of AD and nonce (A, N) is unique for all encryption queries, for privacy and authenticity notions.

Table 1 Security goal for confidentiality (privacy).

	aes128otrpv1		aes128otrsv1		aes256otrpv1		aes256otrsv1	
	Data	Time	Data	Time	Data	Time	Data	Time
Confidentiality (Privacy)	64	128	64	128	64	256	64	256

Table 2 Security goal for integrity (authenticity).

	aes128otrpv1			aes128otrsv1			aes256otrpv1			aes256otrsv1		
	Data	Verify	Time	Data	Verify	Time	Data	Verify	Time	Data	Verify	Time
Integrity (Authenticity)	64	128 (τ)	128	64	128 (τ)	128	64	128 (τ)	256	64	128 (τ)	256

<p>Algorithm OTR-$\mathcal{E}_{E,\tau,p}(N, A, M)$</p> <ol style="list-style-type: none"> 1. $(C, TE) \leftarrow \text{EF}_E(N, M)$ 2. if $A \neq \varepsilon$ then $TA \leftarrow \text{AF}_E(A)$ 3. else $TA \leftarrow 0^n$ 4. $T \leftarrow \text{msb}_\tau(TE \oplus TA)$ 5. return (C, T) 	<p>Algorithm OTR-$\mathcal{D}_{E,\tau,p}(N, A, C, T)$</p> <ol style="list-style-type: none"> 1. $(M, TE) \leftarrow \text{DF}_E(N, C)$ 2. if $A \neq \varepsilon$ then $TA \leftarrow \text{AF}_E(A)$ 3. else $TA \leftarrow 0^n$ 4. $\widehat{T} \leftarrow \text{msb}_\tau(TE \oplus TA)$ 5. if $\widehat{T} = T$ return M 6. else return \perp
<p>Algorithm $\text{EF}_E(N, M)$</p> <ol style="list-style-type: none"> 1. $\Sigma \leftarrow 0^n$ 2. $\delta \leftarrow E(\underline{N}), L \leftarrow 4\delta$ 3. $(M[1], \dots, M[m]) \stackrel{\leftarrow}{\leftarrow} M$ 4. for $i = 1$ to $\lceil m/2 \rceil - 1$ do 5. $C[2i - 1] \leftarrow E(L \oplus M[2i - 1]) \oplus M[2i]$ 6. $C[2i] \leftarrow E(L \oplus \delta \oplus C[2i - 1]) \oplus M[2i - 1]$ 7. $\Sigma \leftarrow \Sigma \oplus M[2i]$ 8. $L \leftarrow 2L$ 9. if m is even 10. $L^* \leftarrow L \oplus \delta$ 11. $Z \leftarrow E(L \oplus M[m - 1])$ 12. $C[m] \leftarrow \text{msb}_{ M[m] }(Z) \oplus M[m]$ 13. $C[m - 1] \leftarrow E(L^* \oplus \underline{C[m]}) \oplus M[m - 1]$ 14. $\Sigma \leftarrow \Sigma \oplus Z \oplus \underline{C[m]}$ 15. if m is odd 16. $L^* \leftarrow L$ 17. $C[m] \leftarrow \text{msb}_{ M[m] }(E(L^*)) \oplus M[m]$ 18. $\Sigma \leftarrow \Sigma \oplus \underline{M[m]}$ 19. if $M[m] \neq n$ then $TE \leftarrow E(3L^* \oplus \Sigma)$ 20. else $TE \leftarrow E(3L^* \oplus \delta \oplus \Sigma)$ 21. $C \leftarrow (C[1], \dots, C[m])$ 22. return (C, TE) 	<p>Algorithm $\text{DF}_E(N, C)$</p> <ol style="list-style-type: none"> 1. $\Sigma \leftarrow 0^n$ 2. $\delta \leftarrow E(\underline{N}), L \leftarrow 4\delta$ 3. $(C[1], \dots, C[m]) \stackrel{\leftarrow}{\leftarrow} C$ 4. for $i = 1$ to $\lceil m/2 \rceil - 1$ do 5. $M[2i - 1] \leftarrow E(L \oplus \delta \oplus C[2i - 1]) \oplus C[2i]$ 6. $M[2i] \leftarrow E(L \oplus M[2i - 1]) \oplus C[2i - 1]$ 7. $\Sigma \leftarrow \Sigma \oplus M[2i]$ 8. $L \leftarrow 2L$ 9. if m is even 10. $L^* \leftarrow L \oplus \delta$ 11. $M[m - 1] \leftarrow E(L^* \oplus \underline{C[m]}) \oplus C[m - 1]$ 12. $Z \leftarrow E(L \oplus M[m - 1])$ 13. $M[m] \leftarrow \text{msb}_{ C[m] }(Z) \oplus C[m]$ 14. $\Sigma \leftarrow \Sigma \oplus Z \oplus \underline{C[m]}$ 15. if m is odd 16. $L^* \leftarrow L$ 17. $M[m] \leftarrow \text{msb}_{ C[m] }(E(L^*)) \oplus C[m]$ 18. $\Sigma \leftarrow \Sigma \oplus \underline{M[m]}$ 19. if $C[m] \neq n$ then $TE \leftarrow E(3L^* \oplus \Sigma)$ 20. else $TE \leftarrow E(3L^* \oplus \delta \oplus \Sigma)$ 21. $M \leftarrow (M[1], \dots, M[m])$ 22. return (M, TE)
<p>Algorithm $\text{AF}_E(A)$</p> <ol style="list-style-type: none"> 1. $\Xi \leftarrow 0^n$ 2. $\gamma \leftarrow E(0^n), Q \leftarrow 4\gamma$ 3. $(A[1], \dots, A[a]) \stackrel{\leftarrow}{\leftarrow} A$ 4. for $i = 1$ to $a - 1$ do 5. $\Xi \leftarrow \Xi \oplus E(Q \oplus A[i])$ 6. $Q \leftarrow 2Q$ 7. $\Xi \leftarrow \Xi \oplus \underline{A[a]}$ 8. if $A[a] \neq n$ then $TA \leftarrow E(Q \oplus \gamma \oplus \Xi)$ 9. else $TA \leftarrow E(Q \oplus 2\gamma \oplus \Xi)$ 10. return TA 	

Fig. 1 Algorithms of AES-OTR with parallel ADP. Tag bit size is $0 < \tau \leq n$, and \underline{X} denotes the 10^* padding of X (See Section 1.3).

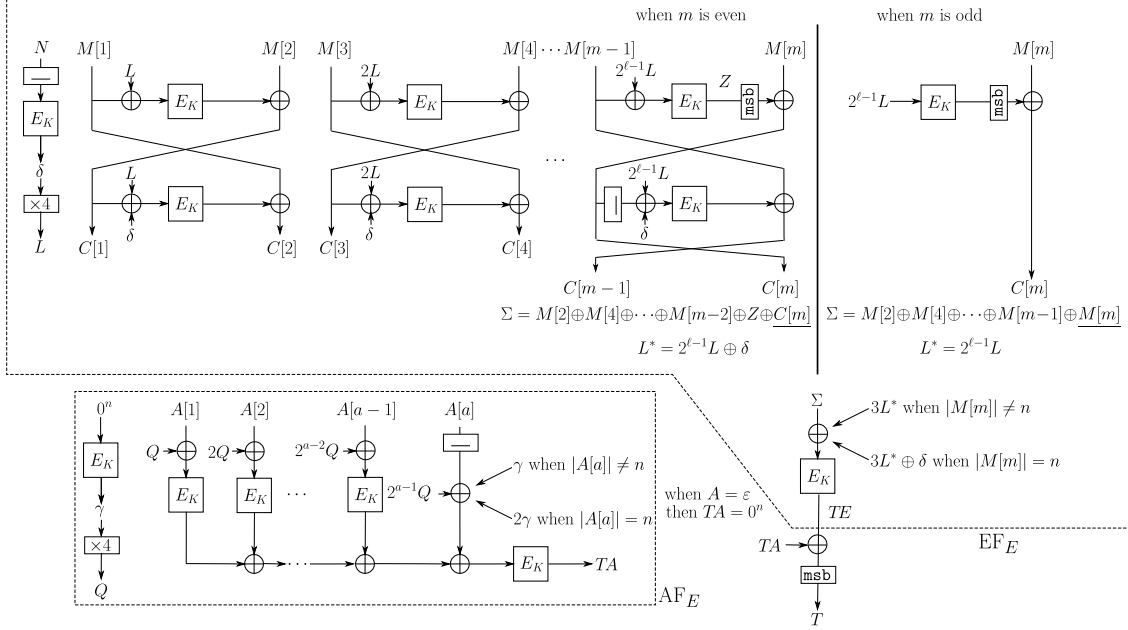


Fig. 2 Encryption of AES-OTR with parallel ADP. A box with underline and \underline{X} denote the 10^* padding of input X .

3 Security Analysis

■ **Provable Security Paradigm.** AES-OTR has a provable security based on the assumption that AES is a pseudorandom function. Below we provide a brief explanation on the security model we consider, which is common to many nonce-based AEAD blockcipher modes.

Let $\text{AE}[\tau]$ be an AEAD having τ -bit tag, where the encryption and decryption algorithms are $\text{AE-}\mathcal{E}_\tau$ and $\text{AE-}\mathcal{D}_\tau$. Then, a PRIV-adversary \mathcal{A} against $\text{AE}[\tau]$ accesses $\text{AE-}\mathcal{E}_\tau$, where the i -th query consists of nonce N_i , associated data A_i , and plaintext M_i . The final output of \mathcal{A} is a binary variable indicating \mathcal{A} 's guess. We define \mathcal{A} 's parameter list to be (q, σ_A, σ_M) , where q denotes the number of queries, and $\sigma_A \stackrel{\text{def}}{=} \sum_{i=1}^q |A_i|_n$ and $\sigma_M \stackrel{\text{def}}{=} \sum_{i=1}^q |M_i|_n$. We assume \mathcal{A} is nonce-respecting, i.e., all N_i s are distinct. We also define random-bit oracle, \mathcal{S} , which takes (N, A, M) and returns $(C, T) \stackrel{\mathcal{S}}{\leftarrow} \{0, 1\}^{|M|} \times \{0, 1\}^\tau$. The privacy notion for \mathcal{A} is defined as

$$\text{Adv}_{\text{AE}[\tau]}^{\text{priv}}(\mathcal{A}) \stackrel{\text{def}}{=} \Pr[K \stackrel{\mathcal{S}}{\leftarrow} \mathcal{K} : \mathcal{A}^{\text{AE-}\mathcal{E}_\tau} \Rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{S}} \Rightarrow 1]. \quad (3)$$

An AUTH-adversary \mathcal{A} against $\text{AE}[\tau]$ accesses $\text{AE-}\mathcal{E}_\tau$ and $\text{AE-}\mathcal{D}_\tau$, using q encryption queries and q_v decryption queries. Let $(N_1, A_1, M_1), \dots, (N_q, A_q, M_q)$ and $(N'_1, A'_1, C'_1, T'_1), \dots, (N'_{q_v}, A'_{q_v}, C'_{q_v}, T'_{q_v})$ be all the encryption and decryption queries made by \mathcal{A} . We define \mathcal{A} 's parameter list to be $(q, q_v, \sigma_A, \sigma_M, \sigma_{A'}, \sigma_{C'})$, where $\sigma_{A'} \stackrel{\text{def}}{=} \sum_{i=1}^{q_v} |A'_i|_n$ and $\sigma_{C'} \stackrel{\text{def}}{=} \sum_{i=1}^{q_v} |C'_i|_n$, in addition to σ_A and σ_M . The authenticity notion for the AUTH-adversary \mathcal{A} is defined as

$$\text{Adv}_{\text{AE}[\tau]}^{\text{auth}}(\mathcal{A}) \stackrel{\text{def}}{=} \Pr[K \stackrel{\mathcal{S}}{\leftarrow} \mathcal{K} : \mathcal{A}^{\text{AE-}\mathcal{E}_\tau, \text{AE-}\mathcal{D}_\tau} \text{ forges }], \quad (4)$$

where \mathcal{A} forges if $\text{AE-}\mathcal{D}_\tau$ returns a bit string (other than \perp) for a decryption query (N'_i, A'_i, C'_i, T'_i) for some $1 \leq i \leq q_v$ such that $(N'_i, A'_i, C'_i, T'_i) \neq (N_j, A_j, C_j, T_j)$ for all $1 \leq j \leq q$. We assume AUTH-adversary \mathcal{A} is always nonce-respecting with respect to encryption queries; using the same N for encryption and

<p>Algorithm OTR-$\mathcal{E}_{E,\tau,s}(N, A, M)$</p> <ol style="list-style-type: none"> 1. if $A \neq \varepsilon$ then $TA \leftarrow \text{AF-S}_E(A)$ 2. else $TA \leftarrow 0^n$ 3. $(C, TE) \leftarrow \text{EF-S}_E(N, M, TA)$ 4. $T \leftarrow \text{msb}_\tau(TE)$ 5. return (C, T) 	<p>Algorithm OTR-$\mathcal{D}_{E,\tau,s}(N, A, C, T)$</p> <ol style="list-style-type: none"> 1. if $A \neq \varepsilon$ then $TA \leftarrow \text{AF-S}_E(A)$ 2. else $TA \leftarrow 0^n$ 3. $(M, TE) \leftarrow \text{DF-S}_E(N, C, TA)$ 4. $\hat{T} \leftarrow \text{msb}_\tau(TE)$ 5. if $\hat{T} = T$ return M 6. else return \perp
<p>Algorithm EF-S$_E(N, M, TA)$</p> <ol style="list-style-type: none"> 1. $\Sigma \leftarrow 0^n$ 2. $\delta \leftarrow E(\underline{N}) \oplus TA, L \leftarrow 4\delta$ 3. $(M[1], \dots, M[m]) \stackrel{r}{\leftarrow} M$ 4. for $i = 1$ to $\lceil m/2 \rceil - 1$ do 5. $C[2i-1] \leftarrow E(L \oplus M[2i-1]) \oplus M[2i]$ 6. $C[2i] \leftarrow E(L \oplus \delta \oplus C[2i-1]) \oplus M[2i-1]$ 7. $\Sigma \leftarrow \Sigma \oplus M[2i]$ 8. $L \leftarrow 2L$ 9. if m is even 10. $L^* \leftarrow L \oplus \delta$ 11. $Z \leftarrow E(L \oplus M[m-1])$ 12. $C[m] \leftarrow \text{msb}_{ M[m] }(Z) \oplus M[m]$ 13. $C[m-1] \leftarrow E(L^* \oplus \underline{C[m]}) \oplus M[m-1]$ 14. $\Sigma \leftarrow \Sigma \oplus Z \oplus \underline{C[m]}$ 15. if m is odd 16. $L^* \leftarrow L$ 17. $C[m] \leftarrow \text{msb}_{ M[m] }(E(L^*)) \oplus M[m]$ 18. $\Sigma \leftarrow \Sigma \oplus \underline{M[m]}$ 19. if $M[m] \neq n$ then $TE \leftarrow E(3L^* \oplus \Sigma)$ 20. else $TE \leftarrow E(3L^* \oplus \delta \oplus \Sigma)$ 21. $C \leftarrow (C[1], \dots, C[m])$ 22. return (C, TE) 	<p>Algorithm DF-S$_E(N, C, TA)$</p> <ol style="list-style-type: none"> 1. $\Sigma \leftarrow 0^n$ 2. $\delta \leftarrow E(\underline{N}) \oplus TA, L \leftarrow 4\delta$ 3. $(C[1], \dots, C[m]) \stackrel{r}{\leftarrow} C$ 4. for $i = 1$ to $\lceil m/2 \rceil - 1$ do 5. $M[2i-1] \leftarrow E(L \oplus \delta \oplus C[2i-1]) \oplus C[2i]$ 6. $M[2i] \leftarrow E(L \oplus M[2i-1]) \oplus C[2i-1]$ 7. $\Sigma \leftarrow \Sigma \oplus M[2i]$ 8. $L \leftarrow 2L$ 9. if m is even 10. $L^* \leftarrow L \oplus \delta$ 11. $M[m-1] \leftarrow E(L^* \oplus \underline{C[m]}) \oplus C[m-1]$ 12. $Z \leftarrow E(L \oplus M[m-1])$ 13. $M[m] \leftarrow \text{msb}_{ C[m] }(Z) \oplus C[m]$ 14. $\Sigma \leftarrow \Sigma \oplus Z \oplus \underline{C[m]}$ 15. if m is odd 16. $L^* \leftarrow L$ 17. $M[m] \leftarrow \text{msb}_{ C[m] }(E(L^*)) \oplus C[m]$ 18. $\Sigma \leftarrow \Sigma \oplus \underline{M[m]}$ 19. if $C[m] \neq n$ then $TE \leftarrow E(3L^* \oplus \Sigma)$ 20. else $TE \leftarrow E(3L^* \oplus \delta \oplus \Sigma)$ 21. $M \leftarrow (M[1], \dots, M[m])$ 22. return (M, TE)
<p>Algorithm AF-S$_E(A)$</p> <ol style="list-style-type: none"> 1. $\Xi \leftarrow 0^n$ 2. $\gamma \leftarrow E(0^n)$ 3. $(A[1], \dots, A[a]) \stackrel{r}{\leftarrow} A$ 4. for $i = 1$ to $a-1$ do 5. $\Xi \leftarrow E(A[i] \oplus \Xi)$ 6. $\Xi \leftarrow \Xi \oplus \underline{A[a]}$ 7. if $A[a] \neq n$ then $TA \leftarrow E(2\gamma \oplus \Xi)$ 8. else $TA \leftarrow E(4\gamma \oplus \Xi)$ 9. return TA 	

Fig. 3 Algorithms of AES-OTR with serial ADP. Tag bit size is $0 < \tau \leq n$, and \underline{X} denotes the 10^* padding of X (See Section 1.3)

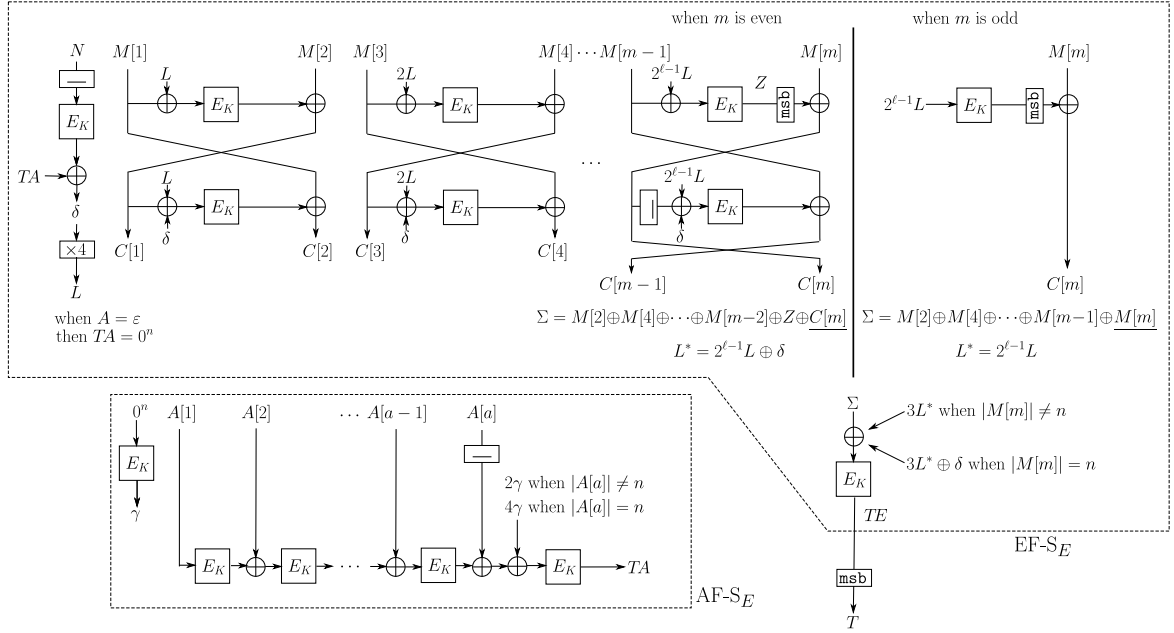


Fig. 4 Encryption of AES-OTR with serial ADP. A box with underline and \underline{X} denote the 10* padding of input X .

decryption queries is allowed, and the same N can be repeated within decryption queries, i.e. N_i is different from N_j for any $j \neq i$ but N'_i may be equal to N_j or $N'_{i'}$ for some j and $i' \neq i$.

The above definitions omit a parameter to specify the adversary's computation power (i.e., for unlimited computational power). If we need to specify the adversary's computation power, we will additionally define the time complexity, t , as the sum of running time of attack and the code size, under some fixed computation model. For details, see e.g. [8].

■ **PRP and PRF.** Let $\text{Func}(a, b)$ be the set of all functions: $\{0, 1\}^a \rightarrow \{0, 1\}^b$, and let $\text{Perm}(a)$ be the set of all permutation over $\{0, 1\}^a$. A uniform random function (URF) and a uniform random permutation (URP) is defined as a random function distributed uniformly random over $\text{Func}(a, b)$ and $\text{Perm}(a)$.

For n -bit blockcipher E with key $K \in \mathcal{K}_{ae}$, Let $\text{Adv}_{E_K}^{\text{prp}}(q, t)$ be the distinguishing advantage between E_K and an n -bit URP, P , using q chosen-plaintext queries and time complexity t . More formally, we have

$$\text{Adv}_{E_K}^{\text{prp}}(q, t) \stackrel{\text{def}}{=} \max_{\mathcal{A}} \{\Pr[K \xleftarrow{\$} \mathcal{K}_{ae} : \mathcal{A}^{E_K} \rightarrow 1] - \Pr[P \xleftarrow{\$} \text{Perm}(n) : \mathcal{A}^P \rightarrow 1]\} \quad (5)$$

where the maximum is taken for all adversaries using q chosen-plaintext queries with time complexity t . By convention we say a blockcipher E_K is a pseudorandom permutation (PRP) if $\text{Adv}_{E_K}^{\text{prp}}(q, t)$ is sufficiently small [18]*¹. Similarly we define

$$\text{Adv}_{E_K}^{\text{prf}}(q, t) \stackrel{\text{def}}{=} \max_{\mathcal{A}} \{\Pr[K \xleftarrow{\$} \mathcal{K}_{ae} : \mathcal{A}^{E_K} \rightarrow 1] - \Pr[R \xleftarrow{\$} \text{Func}(n, n) : \mathcal{A}^R \rightarrow 1]\} \quad (6)$$

and (conventionally) say E_K is a pseudorandom function (PRF) when $\text{Adv}_{E_K}^{\text{prf}}(q, t)$ is sufficiently small. Thanks to the well-known PRP-PRF switching lemma (e.g. see [11]), the both statements are exchangeable as long as $q \ll 2^{n/2}$.

*¹ This is actually a convention because the formal definition of PRP is for a family of keyed permutation, and we here did not specify a formal definition of "sufficiently small".

■ **Security Bounds.** First we provide the security bounds of OTR with parallel ADP.

Theorem 3.1 Fix $\tau \in \{1, \dots, n\}$. For any PRIV-adversary \mathcal{A} with parameter (q, σ_A, σ_M) and time complexity t ,

$$\text{Adv}_{\text{OTR}[E_K, \tau, p]}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_{E_K}^{\text{prp}}(2\sigma_{\text{priv}}, t') + \frac{6\sigma_{\text{priv}}^2}{2^n}$$

holds for $\sigma_{\text{priv}} = q + \sigma_A + \sigma_M$ and t' is about t plus the complexity of $2\sigma_{\text{priv}}$ AES encryption.

Theorem 3.2 Fix $\tau \in \{1, \dots, n\}$. For any AUTH-adversary \mathcal{A} with parameter $(q, q_v, \sigma_A, \sigma_M, \sigma_{A'}, \sigma_{C'})$ and time complexity t ,

$$\text{Adv}_{\text{OTR}[E_K, \tau, p]}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_{E_K}^{\text{prp}}(2\sigma_{\text{auth}}, t') + \frac{6\sigma_{\text{auth}}^2}{2^n} + \frac{q_v}{2^\tau}$$

holds for $\sigma_{\text{auth}} = q + q_v + \sigma_A + \sigma_M + \sigma_{A'} + \sigma_{C'}$ and t' is about t plus the complexity of $2\sigma_{\text{auth}}$ AES encryption.

Next we provide the security bounds of OTR with serial ADP for privacy and authenticity notions.

Theorem 3.3 Fix $\tau \in \{1, \dots, n\}$. For any CPA-adversary \mathcal{A} with parameter (q, σ_A, σ_M) and time complexity t ,

$$\text{Adv}_{\text{OTR}[E_K, \tau, s]}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_{E_K}^{\text{prp}}(2\sigma_{\text{priv}}, t') + \frac{5.5\sigma_{\text{priv}}^2}{2^n}$$

holds for $\sigma_{\text{priv}} = q + \sigma_A + \sigma_M$ and t' is about t plus the complexity of $2\sigma_{\text{priv}}$ AES encryption.

Theorem 3.4 Fix $\tau \in \{1, \dots, n\}$. For any CCA-adversary \mathcal{A} with parameter $(q, q_v, \sigma_A, \sigma_M, \sigma_{A'}, \sigma_{C'})$ and time complexity t ,

$$\text{Adv}_{\text{OTR}[E_K, \tau, s]}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_{E_K}^{\text{prp}}(2\sigma_{\text{auth}}, t') + \frac{7.5\sigma_{\text{auth}}^2}{2^n} + \frac{q_v}{2^\tau}$$

holds for $\sigma_{\text{auth}} = q + q_v + \sigma_A + \sigma_M + \sigma_{A'} + \sigma_{C'}$ and t' is about t plus the complexity of $2\sigma_{\text{auth}}$ AES encryption.

A part of proofs for parallel ADP version will appear at [21]. Full proofs of all theorems will appear in [20].

4 Features

AES-OTR has the following features.

- The key is one AES key, K .
- Inverse-free, that is, encryption and decryption can be done by AES encryption function, E_K .
- Rate-1 processing for both encryption and decryption. More precisely, for a -block AD and m -block plaintext, the number of AES calls is $a + m + 2$ (with one preprocessing call).
- On-line, one-pass, and parallel encryption and decryption, under two-block partition. For serial ADP version, associated data processing is serial but plaintext can be processed in parallel.
- Provable security up to about $2^{n/2}$ input blocks, based on the assumption that E_K is a pseudorandom function (PRF). From the PRP-PRF switching lemma, this holds also for the assumption that E_K is a pseudorandom permutation (PRP).

■ **Advantages over Previous AEAD Modes.** A comparison of OTR with other modes is shown in Table 3. Being rate-1 means the computation cost is halved from common rate-2 schemes. Compared with AES-GCM, AES-OTR does not require a full GF multiplier, while the computation cost of AES-OTR is

comparable to the encryption mode without the blockcipher decryption, such as AES-CTR. Roughly, the computation cost for encryption is comparable to that of AES-OCB. For decryption, their performances can be different depending on the platform, see below.

In return for these attractive features, one potential drawback of OTR is that it needs two-block partition, which requires more state memories required than that of OCB, though OCB needs the additional implementation of blockcipher decryption. The parallelizability of our scheme is up to the half of the message blocks, while OCB has full parallelizability, up to the number of message blocks. On-line processing capability is restrictive as it needs buffering of consecutive two input blocks.

For memory consumption, all inverse-free modes including OTR have a similar profile, as long as the blockcipher encryption is the dominant factor. An exception is GCM since field multiplication needs large memories for fast operation. The design of OTR avoids a large memory consumption.

Table 3 A comparison of AEAD modes. Calls denotes the number of primitive calls for m -block message and a -block header and one-block nonce, without constants.

Mode	Calls	On-line	Parallel	Primitive
CCM [3]	$a + 2m$	no	no	E
GCM [5]	m [E] and $a + m$ [Mul]	yes	yes	E, Mul^\dagger
EAX [9]	$a + 2m$	yes	no	E
OCB [17, 23, 25]	$a + m$	yes	yes	E, E^{-1}
CCFB [19]	$a + cm$ for some $1 < c^\ddagger$	yes	no	E
OTR	$a + m$	yes [¶]	yes [¶]	E

[†] $\text{GF}(2^n)$ multiplication

[‡] Security degrades as c approaches 1

[¶] two-block partition, serial AD for serial ADP version.

■ **Benefits of inverse-freeness.** The features listed above are already covered by OCB mode, except that OCB needs both blockcipher encryption and decryption (more precisely the OCB decryption needs both functions). As mentioned by Iwata and Yasuda [16], inverse-freeness can contribute to efficiency and security. First, for efficiency, the initialization AES-OTR needs AES key schedule only for the forward direction. In the standard, so-called T-table implementation, we can remove some tables required for the inverse, resulting in smaller ROM and RAM consumption on software. When side-channel protection is needed, inverse-freeness can help reduce the critical functions to protect. Moreover, AES decryption can be slower than AES encryption on some platforms (e.g., see AES implementation for Atmel AVR by Osvik et al. [22]). This property is a consequence of the initial design choice [12], focusing on inverse-free modes. The uneven performance figures of blockcipher enc/dec functions is undesirable in practice, when the mode uses both functions.

This phenomena, however, is not always true. Most notably, recent Intel’s CPUs have AES instructions called AESNI, which enables the identical performance for AES encryption and decryption, and AMD CPUs have an equivalent ones, too. Therefore, when our proposal uses AESNI, the performance would be roughly similar to that of AES-OCB with AESNI, for both encryption and decryption, though the complexity outside the blockcipher may degrade the result. We have other SW platforms where hardware AES is available but decryption is slower (e.g., [14]), or only encryption function is available (e.g., [13]). Basically, the value of our proposal is *not* to provide the fastest operation on modern CPUs, instead, to increase the availability of rate-1 performance for various platforms, using single algorithm based on the minimal cryptographic primitive and assumption.

For hardware the algorithm is parallelizable and also pipelinable. AES-OTR’s inverse-freeness significantly contributes in hardware size reduction for various forms of hardware implementation from standard round-based implementation to extremely high speed implementations using pipeline or multiple AES cores.

Next, inverse-freeness can contribute to security. Usually the security of a mode using both enc/dec functions of a blockcipher, denoted by E and E^{-1} , needs (E, E^{-1}) to be a strong pseudorandom permutation (Strong PRP or SPRP). This holds true for the original security proofs of all versions of

OCB [17, 23, 25], though a recent work of Aoki and Yasuda [7] showed a relaxation on the security condition for the blockcipher used by OCB. In contrast, when the mode uses only E , the security assumption is relaxed to PRP or PRF.

■**Security.** AES-OTR has a provable security reduction to the pseudorandomness of AES, which is a quite popular assumption used by many cryptographic designs. The proved security bounds show that one can not break the scheme in the senses of privacy and authenticity, up to around the birthday bound, i.e. 2^{64} blocks. Here, however, we have to warn that the security is proved for the standard nonce-respecting adversary, i.e. the encryption never processes duplicate nonces. We do not claim any security guarantee for adversaries beyond this condition, except a slight relaxation for serial ADP, shown in Section 2, which says that the security is preserved as long as a pair (A, N) is unique, i.e., it can work as nonce.

■**Justifications for Recommended Parameter Sets.** The nonce length of 12 bytes is a reasonable choice for simplicity, and is widely employed by many protocols. The tag length of 16 bytes offers high-security against even intensive forgery attempts. The AES key length of 128 bits provides adequate security against key exhaustive searches, and 256 key bits are sometimes required for extra security. For parallel ADP we need more memory than serial ADP, hence the latter is preferable for serial computing environment. In addition the computation of serial ADP is simpler than parallel ADP. See Section 5.

In addition, we specify the minimum tag length to 4 bytes, though the recommended parameter sets have 16-byte tag. Such short tag is usually dangerous, however, still useful to suppress the communication overhead with an appropriate limitation on the amount of data processed for one key. In fact we can find many examples of 4-byte tags in the realm of low-power wireless sensor networks, e.g., 6LoWPAN and Zigbee, both using AES-CCM with tag length between 4 to 16 bytes [1, 2].

5 Design Rationale

Our goal is to provide a fast and compact AEAD suitable for various platforms, using a minimum cryptographic primitive, i.e., encryption function of a blockcipher. We choose AES as the underlying blockcipher, because the security of AES has been extensively studied. The cipher achieves rate-1 parallelizable processing without using AES inverse. The computation cost is comparable to encryption modes using AES encryption function (e.g. CTR). Since we can not achieve privacy notion of AE with fewer AES calls than AES-CTR, this implies that the computation cost of AES-OTR is asymptotically minimum as a mode of AES encryption function.

■**Masks.** For masking applied to AES inputs inside AES-OTR, we employed constant GF multiplications, called GF doublings. This is used by many schemes. For hardware, doubling is simple. For software, doubling is basically not simple as it requires bit-shift of whole 128-bit block. However, recent studies reported that the optimized doubling software implementation can be fast [6]. Considering this we employ on-the-fly doubling as a practical masking option. We remark that using the same mask for the round functions in a two-round Feistel (say omitting xoring δ in Line 6 of EF_E in Fig. 1), does not work. This is because the two-round Feistel then becomes an involution, which allows the adversary to control the checksum value in the decryption, hence breaks authenticity. We also remark that the all masks for EF_E depend on N , hence do not allow precomputation before N is given, which is different from the latest OCB3 [17]. The reason of our choice is that we want our scheme not to waste $E_K(0^n)$ when AD is always empty, and not to require a large RAM for precomputation, considering applicability to constrained devices. Our goal is not to ultimately focus on one platform, hence we consider keeping the amount of precomputation small is good to achieve a balanced performance under multiple platforms.

■**Associated Data Processing.** We have two versions depending on associated data processing, ADP. For parallel ADP version, the processing of AD is based on (a variant of) PMAC, and the computation can be done in parallel to the processing of plaintext/ciphertext. This is to maximize the parallel computation capability. We can efficiently handle static AD, and employ counter-nonce caching due to Wu (mentioned by Bernstein [10]), if needed. For serial ADP version, the processing of AD is based on CMAC, hence is

inherently serial. At the cost of losing parallelizability for AD, this version reduces the computation cost of ADP in serial environment. Also, the point where CMAC output (TA) is XORed is chosen to reduce the required state memory, while keeping the capability of static AD and counter-nonce caching. These design choices make AES-OTR with serial ADP version suitable to, e.g., embedded software.

The designer has not hidden any weaknesses in this cipher.

6 Intellectual Property

We have a pending patent relating to AES-OTR, JP application No. 2013-161446. In case that our submission is included into the final portfolio, we are willing to make AES-OTR widely available with considering the possibility of licensing it at royalty free in implementing AES-OTR. Any suggestions and feedbacks are welcome on this issue.

If any of this information changes, the submitter will promptly (and within at most one month) announce these changes on the `crypto-competitions` mailing list.

7 Consent

The submitter hereby consents to all decisions of the CAESAR selection committee regarding the selection or non-selection of this submission as a second-round candidate, a third-round candidate, a finalist, a member of the final portfolio, or any other designation provided by the committee. The submitter understands that the committee will not comment on the algorithms, except that for each selected algorithm the committee will simply cite the previously published analyses that led to the selection of the algorithm. The submitter understands that the selection of some algorithms is not a negative comment regarding other algorithms, and that an excellent algorithm might fail to be selected simply because not enough analysis was available at the time of the committee decision. The submitter acknowledges that the committee decisions reflect the collective expert judgments of the committee members and are not subject to appeal. The submitter understands that if they disagree with published analyses then they are expected to promptly and publicly respond to those analyses, not to wait for subsequent committee decisions. The submitter understands that this statement is required as a condition of consideration of this submission by the CAESAR selection committee.

References

- [1] IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals, <https://tools.ietf.org/html/rfc4919/>
- [2] ZigBee Alliance, <http://www.zigbee.org/>
- [3] 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.
- [4] Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication. NIST Special Publication 800-38B (2005), national Institute of Standards and Technology.
- [5] Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D (2007), national Institute of Standards and Technology.
- [6] Aoki, K., Iwata, T., Yasuda, K.: How Fast Can a Two-Pass Mode Go? A Parallel Deterministic Authenticated Encryption Mode for AES-NI. DIAC 2012: Directions in Authenticated Ciphers (2012), available from <http://hyperelliptic.org/DIAC/>
- [7] Aoki, K., Yasuda, K.: The Security of the OCB Mode of Operation without the SPRP Assumption. In: Susilo, W., Reyhanitabar, R. (eds.) ProvSec. Lecture Notes in Computer Science, vol. 8209, pp. 202–220. Springer (2013)
- [8] Bellare, M., Kilian, J., Rogaway, P.: The Security of the Cipher Block Chaining Message Authentication Code. J. Comput. Syst. Sci. 61(3), 362–399 (2000)
- [9] Bellare, M., Rogaway, P., Wagner, D.: The EAX Mode of Operation. In: Roy and Meier [26], pp.

- [10] Bernstein, D.J., Schwabe, P.: New AES Software Speed Records. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT. Lecture Notes in Computer Science, vol. 5365, pp. 322–336. Springer (2008)
- [11] Black, J., Rogaway, P.: CBC MACs for Arbitrary-Length Messages: The Three-Key Constructions. In: Bellare, M. (ed.) CRYPTO. Lecture Notes in Computer Science, vol. 1880, pp. 197–215. Springer (2000)
- [12] Daemen, J., Rijmen, V.: AES Proposal: Rijndael (1999)
- [13] Didla, S., Ault, A., Bagchi, S.: Optimizing AES for embedded devices and wireless sensor networks. In: de Leon, M.P. (ed.) TRIDENTCOM. p. 4. ICST (2008)
- [14] Gouvêa, C.P.L., López, J.: High Speed Implementation of Authenticated Encryption for the MSP430X Microcontroller. In: Hevia, A., Neven, G. (eds.) LATINCRYPT. Lecture Notes in Computer Science, vol. 7533, pp. 288–304. Springer (2012)
- [15] Iwata, T., Kurosawa, K.: OMAC: One-Key CBC MAC. In: Johansson, T. (ed.) FSE. Lecture Notes in Computer Science, vol. 2887, pp. 129–153. Springer (2003)
- [16] Iwata, T., Yasuda, K.: BTM: A Single-Key, Inverse-Cipher-Free Mode for Deterministic Authenticated Encryption. In: Jr., M.J.J., Rijmen, V., Safavi-Naini, R. (eds.) Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 5867, pp. 313–330. Springer (2009)
- [17] Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) FSE. Lecture Notes in Computer Science, vol. 6733, pp. 306–327. Springer (2011)
- [18] Luby, M., Rackoff, C.: How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM J. Comput.* 17(2), 373–386 (1988)
- [19] Lucks, S.: Two-Pass Authenticated Encryption Faster Than Generic Composition. In: Gilbert, H., Handschuh, H. (eds.) FSE. Lecture Notes in Computer Science, vol. 3557, pp. 284–298. Springer (2005)
- [20] Minematsu, K.: Parallelizable Authenticated Encryption from Functions. *IACR Cryptology ePrint Archive* 2013, 628 (2013)
- [21] Minematsu, K.: Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In: *Eurocrypt* (2014), to appear
- [22] Osvik, D.A., Bos, J.W., Stefan, D., Canright, D.: Fast Software AES Encryption. In: Hong, S., Iwata, T. (eds.) FSE. Lecture Notes in Computer Science, vol. 6147, pp. 75–93. Springer (2010)
- [23] Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT. Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer (2004)
- [24] Rogaway, P.: Nonce-Based Symmetric Encryption. In: Roy and Meier [26], pp. 348–359
- [25] Rogaway, P., Bellare, M., Black, J.: OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Trans. Inf. Syst. Secur.* 6(3), 365–403 (2003)
- [26] Roy, B.K., Meier, W. (eds.): Fast Software Encryption, 11th International Workshop, FSE 2004, Delhi, India, February 5-7, 2004, Revised Papers, Lecture Notes in Computer Science, vol. 3017. Springer (2004)