Improved Security Analysis of XEX and LRW Modes

Kazuhiko Minematsu

NEC Corporation, 1753 Shimonumabe, Nakahara-Ku, Kawasaki 211-8666, Japan
k-minematsu@ah.jp.nec.com

Abstract. We study block cipher modes that turn a block cipher into a tweakable block cipher, which accepts an auxiliary variable called tweak in addition to the key and message. Liskov et al. first showed such a mode using two keys, where one is the block cipher's key and the other is used for some non-cryptographic function. Later, Rogaway proposed the XEX mode to reduce these two keys to one key. In this paper, we propose a generalization of the Liskov et al.'s scheme with a concrete security proof. Using this, we provide an improved security proof of the XEX and some improvements to the LRW-AES, which is a straightforward AES-based instantiation of Liskov et al.'s scheme proposed by the IEEE Security in Storage Workgroup.

1 Introduction

Tweakable block ciphers are block ciphers that accept a variable called tweak in addition to the key and message. They were formally defined by Liskov, Rivest, and Wagner [10]. In their definition, a tweak is used to provide variability: any two different tweaks give two instances of an ordinary (i.e., not tweakable) block cipher. Formally, tweakable block ciphers are defined as a function $\widetilde{E} : \mathcal{M} \times \mathcal{K} \times \mathcal{T} \to \mathcal{M}$, where $(\mathcal{M}, \mathcal{K}, \mathcal{T})$ denotes (message space, key space, tweak space). For any two tweak values, $T \neq T'$, the outputs of $\widetilde{E}_{K,T}$ should appear to be independent of outputs of $\widetilde{E}_{K,T'}$ even if T and T' are public but K is secret. Liskov et al. showed that a standard block cipher could be easily converted into a tweakable one by using a mode of operation similar to DESX [9]. They also pointed out that tweakable block ciphers are key components to build advanced modes such as authenticated encryption modes. Their proposal, which we call the LRW mode, is as follows. For plaintext M with tweak T, the ciphertext is $C = E_K(M \oplus \Delta(T)) \oplus \Delta(T)$, where Δ is a keyed function of T called the offset function. They proved that the LRW mode was provably secure if the key of Δ, denoted by K_Δ, was independent of K, and Δ was ϵ-almost XOR universal (ϵ-AXU) for sufficiently small ϵ (see Def. 2). The security considered here is the indistinguishability from the ideal tweakable block cipher using any combination of chosen-plaintext attack (CPA) and chosen-ciphertext attack (CCA) for chosen tweaks. A tweakable block cipher with this property is called a strong tweakable block cipher.

E. Biham and A.M. Youssef (Eds.): SAC 2006, LNCS 4356, pp. 96–113, 2007.

The LRW mode needs two independent keys. However, what would happen if K_Δ is *not* independent of the block cipher key, K? For example, is it safe to use E_K(constant) as (a part of) K_Δ? In this paper, we study this problem. Our main contribution is a general construction of strong tweakable block ciphers with a concrete security proof. Our scheme has basically the same structure as that of the LRW mode, but we allow Δ to invoke E_K and/or another function which is possibly independently keyed of K. Using our scheme, we provide some improvements to the previous modes. The first target is the XEX mode [19], an one-key tweakable block cipher similar to the LRW mode. Here, 'one-key' means that the key of the mode is the block cipher key, and only one block cipher key scheduling is needed. XEX mode has a parameter, and in the initial definition of XEX, it was claimed that it was strongly secure if its parameter provided "unique representations" [18]. However, providing unique representations is not a sufficient condition. XEX having a bad parameter is vulnerable to a very simple attack even if it provides unique representations (see [19] and Sect. 4.1 of this paper). The published version of XEX fixed this problem [19]. Our generalized construction clearly explains why this fix works well (which is only briefly mentioned in [19]) and provides a security proof of the fixed XEX, which improves the one shown in [19].

Our second target is the LRW-AES, which is a straightforward instantiation of LRW mode using AES [23]. It has been discussed by the IEEE security in storage working group (SISWG) as a standard mode for storage encryption. The offset function of LRW-AES uses multiplication in GF(2^{128}), where a tweak is multiplied by the 128-bit key independent of the block cipher's key. Using our scheme, we demonstrate how to reduce two keys of the LRW-AES to one key without increasing the computational cost or reducing the allowed tweak set. The underlying idea is similar (but not identical) to one applied to XEX. We also present an alternative mode of AES using the 4-round AES as the offset function. That is, the mode is essentially AES-based and no dedicated AXU function is needed. XEX mode of AES has this property too; however, it allows only incremental update of a tweak. In contrast, our proposal enables us to update a tweak arbitrarily at the cost of 4-round AES invocation. We provide an experimental implementation of our AES-based mode and demonstrate that ours is much more efficient than the reference LRW-AES implementation.

2 Preliminaries

2.1 Notation

Σ^n denotes $\{0,1\}^n$. If a random variable X is uniformly distributed over a set $\mathcal{X}$, we write $X \in_{\mathrm{u}} \mathcal{X}$. An n-bit block uniform random function (URF), denoted by R, is a random variable uniformly distributed over $\{f : \Sigma^n \to \Sigma^n\}$. Similarly, a random variable distributed over all n-bit permutations is an n-bit block uniform random permutation (URP) and is denoted by P. A tweakable n-bit URP with the tweak space $\mathcal{T}$ is defined by the set of $|\mathcal{T}|$ independent URPs (i.e., an independent n-bit URP is used for each tweak in $\mathcal{T}$) and is denoted

by $\widetilde{\mathsf{P}}$. If $F_K : \mathcal{X} \to \mathcal{Y}$ is a keyed function, then F_K is a random variable (not necessarily uniformly) distributed over $\{f : \mathcal{X} \to \mathcal{Y}\}$. If its key, K, is uniform over $\mathcal{K}$, we have $\Pr(F_K(x) = y) = \{k \in \mathcal{K} : f(k, x) = y\}/|\mathcal{K}|$ for some function $f : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$. If K is fixed to $k \in \mathcal{K}$, F_k denotes a function $f(k, *)$. If K is clear from the context, we will omit the subscript K and write $F : \mathcal{X} \to \mathcal{Y}$.

Elements of $\mathrm{GF}(2^n)$. We express the elements of field $\mathrm{GF}(2^n)$ by the n-bit coefficient vectors of the polynomials in the field. We alternatively represent n-bit coefficient vectors by integers $0, 1, \ldots, 2^n - 1$. For example, 5 corresponds to the coefficient vector $(00 \ldots 0101)$ (which corresponds to the polynomial $\mathtt{x}^2 + 1$) and 1 corresponds to $(00 \ldots 01)$, i.e., the identity element.

2.2 Security Notion

Definition 1. *Let F and G be two keyed n-bit block functions. Let us assume that the oracle has implemented H, which is identical to one of F or G. An adversary, A, guesses if H is F or G using CPA. The maximum CPA-advantage in distinguishing F from G is defined as*

$$\mathtt{Adv}^{\mathtt{cpa}}_{F,G}(q, \tau) \stackrel{\text{def}}{=} \max_{A:(q,\tau)\text{-CPA}} \left| \Pr(A^F = 1) - \Pr(A^G = 1) \right|, \tag{1}$$

where $A^F = 1$ denotes that A's guess is 1, which indicates one of F or G, and (q, τ)-CPA denotes a CPA that uses q queries with time complexity τ (see [1] for a detailed description of τ). If the attacks have unlimited computational power, we write $\mathtt{Adv}^{\mathtt{cpa}}_{F,G}(q)$.

Let E_K be an n-bit block cipher and let $\widetilde{E}_K$ be an n-bit tweakable block cipher with tweak space $\mathcal{T}$. For any $x \in \Sigma^n$, $t \in \mathcal{T}$, and $w \in \{0, 1\}$, we define:

$$E^{\pm}_K(x, w) \stackrel{\text{def}}{=} \begin{cases} E_K(x) & \text{if } w = 0 \\ E^{-1}_K(x) & \text{if } w = 1, \end{cases} \quad \widetilde{E}^{\pm}_K(x, t, w) \stackrel{\text{def}}{=} \begin{cases} \widetilde{E}_K(x, t) & \text{if } w = 0 \\ \widetilde{E}^{-1}_K(x, t) & \text{if } w = 1, \end{cases}$$

where E^{-1}_K denotes the inversion of E_K. The securities of E_K and $\widetilde{E}_K$ are measured by

$$\mathtt{Adv}^{\mathtt{sprp}}_{E_K}(q, \tau) \stackrel{\text{def}}{=} \mathtt{Adv}^{\mathtt{cca}}_{E_K,\mathsf{P}}(q, \tau) \stackrel{\text{def}}{=} \mathtt{Adv}^{\mathtt{cpa}}_{E^{\pm}_K,\mathsf{P}^{\pm}}(q, \tau), \quad \text{and} \tag{2}$$

$$\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\widetilde{E}_K}(q, \tau) \stackrel{\text{def}}{=} \mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\widetilde{E}_K,\widetilde{\mathsf{P}}}(q, \tau) \stackrel{\text{def}}{=} \mathtt{Adv}^{\mathtt{cpa}}_{\widetilde{E}^{\pm}_K,\widetilde{\mathsf{P}}^{\pm}}(q, \tau), \tag{3}$$

where P ($\widetilde{\mathsf{P}}$) is an n-bit URP (tweakable URP with tweak space $\mathcal{T}$). A keyed permutation that can not be efficiently distinguished from URP (i.e., CPA-advantage is negligibly small for any practical (q, τ)) is called a pseudorandom permutation (PRP) [3]. A PRP with a negligibly small Chosen ciphertext attack (CCA)-advantage (i.e., $\mathtt{Adv}^{\mathtt{sprp}}_{E_K}(q, \tau)$) is a strong PRP (SPRP). We focus on modes that turn an n-bit SPRP into an n-bit strong tweakable block cipher, which has negligibly small $\widetilde{\mathrm{CCA}}$-advantage, i.e., $\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\widetilde{E}_K}(q, \tau)$, for any practical (q, τ).

3 Previous Tweakable Block Cipher Modes

3.1 General DESX-Like Mode

All modes dealt within this paper including our proposals have the form defined as

$$\widetilde{E}_{K,K_\Delta}(T, M) = E_K(M \oplus \Delta(T)) \oplus \Delta(T), \tag{4}$$

where $T \in \mathcal{T}$ is a tweak and $M \in \Sigma^n$ is a plaintext. Here, E_K is n-bit block and $\Delta : \mathcal{T} \to \Sigma^n$ is a keyed function of tweak and called an *offset function*. Its key is denoted by K_Δ. Δ can invoke E_K and/or another function which is fixed or independently keyed of K. Thus K_Δ is not always independent of K.

3.2 LRW Mode

Liskov et al.'s scheme, which we call the LRW mode, uses the offset function Δ, where its key K_Δ is independent of the block cipher key, K. To prove its security, we need the notion of an ϵ-almost XOR universal hash function, which is as follows.

Definition 2. *Let $F_K : \mathcal{X} \to \Sigma^n$ be a keyed function with $K \in_{\mathrm{u}} \mathcal{K}$. If $\Pr(F_K(x) \oplus F_K(x') = c) \leq \epsilon$ for any $x \neq x'$ and $c \in \Sigma^n$, then F_K is an ϵ-almost XOR universal (ϵ-AXU) hash function.*

The following theorem proves the security of LRW mode. The proof is in Appendix B.

Theorem 1. *Let $\widetilde{E}_{K,K_\Delta}$ be the LRW mode using the offset function Δ, where its key is $K_\Delta \in_{\mathrm{u}} \mathcal{K}_\Delta$ and tweak $T \in \mathcal{T}$ (see Eq. (4)). If Δ is ϵ-AXU (for input T) and K_Δ is independent of the block cipher key K, then $\mathtt{Adv}^{\widetilde{\mathrm{sprp}}}_{\widetilde{E}_{K,K_\Delta}}(q, \tau) \leq \mathtt{Adv}^{\mathrm{sprp}}_{E_K}(q, \tau') + \epsilon q^2$, where $\tau' = \tau + O(q)$.*

This is better than the result of Liskov et al. (theorem 2 of [10]), as they showed $3\epsilon q^2$ instead of ϵq^2. A straightforward instantiation of the LRW is to define $\Delta(T) = K_\Delta \cdot T$, where $K_\Delta \in_{\mathrm{u}} \Sigma^n$ and $T \in \Sigma^n$, and $\cdot$ denotes multiplication in $\mathrm{GF}(2^n)$. This apparently has bias $\epsilon = 1/2^n$. The mode of AES with this offset function has been considered by the IEEE SISWG under the name LRW-AES.

3.3 XEX Mode

XEX mode was proposed by Rogaway [19]. It was designed to be a strong tweakable block cipher. According to the definition of XEX, a base is an element of $\Sigma^n \setminus \{0\}$, and a set $\mathbb{I}_1^d \stackrel{\mathrm{def}}{=} \mathbb{I}_1 \times \mathbb{I}_2 \times \cdots \times \mathbb{I}_d$ is called an index set, where $\mathbb{I}_i \subseteq \{0, 1, \ldots, 2^n - 1\}$ for all i. A pair of a list of bases $\alpha_1, \ldots, \alpha_d$ and an index set $\mathbb{I}_1^d$ is a parameter setting of XEX.

A XEX mode with a parameter setting $((\alpha_1, \ldots, \alpha_d), \mathbb{I}_1^d)$ has the tweak space $\mathbb{I}_1^d \times \Sigma^n$. Let $(i_1, \ldots, i_d, N) \in \mathbb{I}_1^d \times \Sigma^n$. The offset function of XEX is defined as:

$$\Delta(i_1, \ldots, i_d, N) = \alpha_1^{i_1} \cdot \alpha_2^{i_2} \cdot \cdots \cdot \alpha_d^{i_d} \cdot V, \text{ where } V = E_K(N). \tag{5}$$

Here, multiplications are done in GF(2^n). Since Δ uses E_K as the source of randomness, XEX mode is one-key and needs only one block cipher key scheduling. XEX mode is highly efficient: if we want to increment a tweak (i.e., increment one of i_j w/o changing other indexes), then it is done with one bitshift and one XOR operation. This technique is called the powering-up construction and has been adopted by other modes [5,6]. Consequently, XEX mode requires no special functions other than the block cipher. Although we can not change a tweak arbitrarily, we can still increment a tweak (with respect to one of i_j) with negligibly small cost.

4 Construction of Strong Tweakable Block Cipher

4.1 A Bug in the Initial XEX and an Attack Against OCB1

A parameter setting of the XEX is said to provide unique representations if it contains no collisions, i.e., $\prod_{j=1}^{d} \alpha_j^{i_j} \neq \prod_{j=1}^{d} \alpha_j^{i'_j}$ for any $(i_1, \ldots, i_d), (i'_1, \ldots, i'_d)$ such that $(i_1, \ldots, i_d) \neq (i'_1, \ldots, i'_d)$. The following example is a parameter setting providing unique representations shown by Rogaway [18].

Example 1. $\alpha_1 = 2, \alpha_2 = 3$ and $\mathbb{I}_1 = \{0, 1, \ldots, 2^{n/2}\}, \mathbb{I}_2 = \{0, 1\}$.

In the initial definition of XEX [18], it was claimed that XEX was a strong tweakable block cipher if its parameter setting provided unique representations. However, this claim turned out to be false, as pointed out by [19]. In general, XEX is broken if its parameter setting allows an index vector, $(i_1, \ldots, i_d)$, such that $\alpha_1^{i_1} \cdot \cdots \cdot \alpha_d^{i_d} = 1$. We call it a "reduced-to-1" index vector. For example, the parameter setting described in Ex. 1 allows the following attack [19].

1. Ask the oracle to decrypt $C_1 = 0$ with tweak $T_1 = (0, 0, N)$ for some N, and obtain a plaintext $M_1 = E_K^{-1}(E_K(N)) \oplus E_K(N) = N \oplus E_K(N)$. Compute $E_K(N) = M_1 \oplus N$.
2. Then, the encryption of $M_2 = 2 \cdot (M_1 \oplus N) \oplus N$ with tweak $T_2 = (1, 0, N)$, which is denoted by C_2, is predictable from $E_K(N)$: $C_2 = E_K(N) \oplus 2 \cdot E_K(N)$.

On The Security of OCB1. The above attack can be used as an attack against OCB1 [18,19], which is an improvement to the famous OCB mode proposed by Rogaway [17]. He proved that (a generalized form of) OCB1 could use any tweakable block cipher as its component, and that it was a secure AE mode if the underlying tweakable block cipher was strong, i.e., $\widetilde{\text{CCA}}$-secure. It would be natural to wonder if one can attack against OCB1 using the XEX with a bad parameter setting (i.e., one containing a "reduced-to-1" index vector). We show this holds true[1], if the *inverse* of XEX, denoted by XEX^{-1}, is used to instantiate OCB1. For instance, let us use the parameter setting of Ex. 1. Then, XEX^{-1} gives the ciphertext $C = E_K^{-1}(M \oplus \Delta(i_1, i_2, N)) \oplus \Delta(i_1, i_2, N)$ where

[1] The OCB1 defined in [18] and [19] are slightly different, however, our attack can be applied to both versions.

$\Delta(i_1, i_2, N) = 2^{i_1}3^{i_2}E_K(N)$. Although this implementation was not mentioned in [19], it was as efficient as the XEX-based one. Moreover, using XEX^{-1} would be preferable to using XEX in some situations. For example, it would be desirable to use XEX^{-1} if E_K is faster[2] than E_K^{-1} and fast operation of the receiver (rather than the sender) is required. Our attack is presented in Appendix C.

4.2 The Security of Fixed XEX

The attack presented in the previous section crucially depends on the existence of reduced-to-1 index vector. Thus it would be natural to think of the idea of removing reduced-to-1 index vector from the allowed tweak set. Here, we prove that this simple fix is theoretically fine.

Theorem 2. *Let* $\mathrm{XEX}[E_K]$ *be the XEX mode of* E_K *with a parameter setting providing unique representations and containing no "reduced-to-*1*" index vector. Then, we have* $\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{XEX}[E_K]}(q, \tau) \le \mathtt{Adv}^{\mathtt{sprp}}_{E_K}(2q, \tau') + \frac{4.5q^2}{2^n}$, where $\tau' = \tau + O(q)$.

For example, we can fix the parameter setting of Ex. 1 by removing $(i_1, i_2) = (0, 0)$. If $n = 128$, the fixed XEX is secure if $q \ll 2^{63}$. The same fix has already been proposed in [19]. However, our proof improves the one shown in [19], which proved $\frac{9.5q^2}{2^n}$ instead of $\frac{4.5q^2}{2^n}$. The proof of Theorem 2 will be provided in Sect. 4.3. One of our purposes is to provide a clear and comprehensive explanation why this fix works well.

4.3 The Proof of Theorem 2

Tools for the Proof. Since we will uses a methodology developed by Maurer [11], we briefly describe his notations. Consider event a_i defined for i input/output pairs (and possibly internal variables) of a keyed function, F. Here, we omit the description of key throughout. Let $\overline{a_i}$ be the negation of a_i. We assume a_i is monotone; i.e., a_i never occurs if $\overline{a_{i-1}}$ occurs. For instance, a_i is monotone if it indicates that all i outputs are distinct. An infinite sequence of monotone events $\mathcal{A} = a_0 a_1 \dots$ is called a *monotone event sequence* (MES). Here, a_0 denotes some tautological event. Note that $\mathcal{A} \wedge \mathcal{B} = (a_0 \wedge b_0)(a_1 \wedge b_1) \dots$ is a MES if $\mathcal{A} = a_0 a_1 \dots$ and $\mathcal{B} = b_0 b_1 \dots$ are both MESs. For any sequence of random variables, $X_1, X_2, \dots$, let X^i denote $(X_1, \dots, X_i)$. After this, $\mathrm{dist}(X^i)$ will denote an event where $X_1, X_2, \dots, X_i$ are distinct. If $\mathrm{dist}(X^i, Y^j)$ holds true, then we have no collision among $\{X_1, \dots, X_i, Y_1, \dots, Y_j\}$.

Let MESs $\mathcal{A}$ and $\mathcal{B}$ be defined for two keyed functions, $F : \mathcal{X} \to \mathcal{Y}$ and $G : \mathcal{X} \to \mathcal{Y}$, respectively. Let $X_i \in \mathcal{X}$ and $Y_i \in \mathcal{Y}$ be the i-th input and output. Let P^F be the probability space defined by F. For example, $P^F_{Y_i|X^iY^{i-1}}(y^i, x^i)$ means $\Pr[Y_i = y_i | X^i = x^i, Y^{i-1} = y^{i-1}]$ where $Y_j = F(X_j)$ for $j \ge 1$. If $P^F_{Y_i|X^iY^{i-1}}(y^i, x^i) = P^G_{Y_i|X^iY^{i-1}}(y^i, x^i)$ for all possible (y^i, x^i), then we write

[2] For instance, some AES software implementations, including the reference code [22], have this property.

$P^F_{Y_i|X^iY^{i-1}} = P^G_{Y_i|X^iY^{i-1}}$. Inequalities such as $P^F_{Y_i|X^iY^{i-1}} \leq P^G_{Y_i|X^iY^{i-1}}$ are similarly defined.

Definition 3. *We write* $F^{\mathcal{A}} \equiv G^{\mathcal{B}}$ *if* $P^F_{Y_ia_i|X^iY^{i-1}a_{i-1}} = P^G_{Y_ib_i|X^iY^{i-1}b_{i-1}}$ *holds for all* $i \geq 1$*, which means* $P^F_{Y_ia_i|X^iY^{i-1}a_{i-1}}(y^i, x^i) = P^G_{Y_ib_i|X^iY^{i-1}b_{i-1}}(y^i, x^i)$ *holds for all possible* (y^i, x^i) *such that both* $P^F_{a_{i-1}|X^{i-1}Y^{i-1}}(y^{i-1}, x^{i-1})$ *and* $P^G_{b_{i-1}|X^{i-1}Y^{i-1}}(y^{i-1}, x^{i-1})$ *are positive.*

Definition 4. *We write* $F|\mathcal{A} \equiv G|\mathcal{B}$ *if* $P^F_{Y_i|X^iY^{i-1}a_i}(y^i, x^i) = P^G_{Y_i|X^iY^{i-1}b_i}(y^i, x^i)$ *holds for all possible* (y^i, x^i) *and all* $i \geq 1$*.*

Note that if $F^{\mathcal{A}} \equiv G^{\mathcal{B}}$, then $F|\mathcal{A} \equiv G|\mathcal{B}$ (but not vice versa).

Definition 5. *For MES* $\mathcal{A}$ *defined for* F*,* $\nu(F, \overline{a_q})$ *denotes the maximal probability of* $\overline{a_q}$ *for any* (q, ∞)*-CPA that interacts with* F*.*

Note that, for any tweakable block cipher $\widetilde{E}_K$, $\nu(\widetilde{E}^{\pm}_K, \overline{a_q})$ is the maximal probability of $\overline{a_q}$ for any $\widetilde{\text{CCA}}$-attacker, i.e., CPA/CCA for chosen tweaks. For simplicity, it will be abbreviated to $\nu(\widetilde{E}_K, \overline{a_q})$. These equivalences are crucial to the information-theoretic security proof. For example, the following theorem holds true.

Theorem 3. *(Theorem* 1 (i) *of [11]) Let* $\mathcal{A}$ *and* $\mathcal{B}$ *be MESs defined for* F *and* G*. If* $F^{\mathcal{A}} \equiv G^{\mathcal{B}}$ *or* $F|\mathcal{A} \equiv G$*, then* $\mathtt{Adv}^{\mathtt{cpa}}_{F,G}(q) \leq \nu(F, \overline{a_q})$*.*

We will use some of Maurer's results including Theorem 3 to make simple and intuitive proofs[3]. For completeness, these results are cited in Appendix A.

General Scheme and Its Security Proof. We proceed as follows. First, we describe a general scheme (which has the form of Eq. (4)) for a tweakable block cipher. Then, we prove that it is a strong tweakable block cipher if its offset function satisfies certain conditions. As the fixed XEX satisfies these conditions, we immediately obtain Theorem 2 as a corollary.

For any two keyed n-bit block functions E_K and $G_{K'}$, let $\mathrm{TW}[E_K, G_{K'}]$ be an n-bit block tweakable block cipher with tweak space $\mathcal{T} = (\mathcal{L}, \Sigma^n)$ for some finite set $\mathcal{L}$. Here E_K must be invertible. Its offset function is defined as

$$\Delta(T) = (F_{K''}(L, G_{K'}(N))), \text{ where } T = (L, N) \in \mathcal{L} \times \Sigma^n. \tag{6}$$

Here, $F_{K''}$ is a keyed function : $\mathcal{L} \times \Sigma^n \to \Sigma^n$ with key $K'' \in_{\mathrm{u}} \mathcal{K}''$ (see Fig. 1). The key of the offset function is (K', K''). We assume that K and K' are not necessarily independent (e.g., $G_{K'} = E_K$ is possible). We also assume that K'' is independent of (K, K') or a *constant* k'' (i.e., $F_{K''}$ can be a fixed function $F_{k''}$). The ranges of keys can be different. What we want to do is to clarify the

[3] Maurer's methodology [11] can only be applied to information-theoretic settings. In most cases information-theoretic proofs can be easily converted into computational ones, but this is not always the case [12,16]. However, we do not encounter such difficulties in this paper. His methodology can also be applied to *random systems*, i.e., stateful probabilistic functions. However, none of our proposals require underlying functions to be stateful.

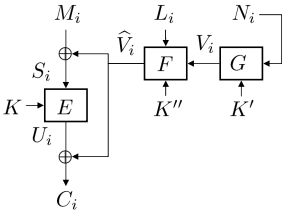

Fig. 1. General scheme for a tweakable block cipher

sufficient condition for $F_{K''}$ to make $\mathrm{TW}[E_K, E_K]$ provably secure. As a first step, we have

$$\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}]}(q) = \mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}],\widetilde{\mathsf{P}}}(q) \le \mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}],\mathrm{TW}[\mathsf{P},\mathsf{R}]}(q) + \mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{TW}[\mathsf{P},\mathsf{R}]}(q), \quad (7)$$

which follows from the triangle inequality. Here, P and R are the n-bit URP and URF, and $\widetilde{\mathsf{P}}$ is the n-bit tweakable URP with tweak space $\mathcal{T}$. Note that P and R in $\mathrm{TW}[\mathsf{P},\mathsf{R}]$ are independent, however, two Ps in $\mathrm{TW}[\mathsf{P},\mathsf{P}]$ denote the same function. We start by analyzing $\mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}],\mathrm{TW}[\mathsf{P},\mathsf{R}]}(q)$ in Eq. (7), which is the main technical part. We need some definitions before the analysis. For any $\mathrm{TW}[E_K, G_{K'}]$, let M_i (C_i) denote the i-th plaintext (ciphertext). In addition, let $T_i = (L_i, N_i)$ be the i-th tweak. We define internal variables of $\mathrm{TW}[E_K, G_{K'}]$ such as $V_i \stackrel{\mathrm{def}}{=} G_{K'}(N_i)$ and $\widehat{V}_i \stackrel{\mathrm{def}}{=} F_{K''}(L_i, V_i)$. Moreover, we have $S_i \stackrel{\mathrm{def}}{=} M_i \oplus \widehat{V}_i$ and $U_i \stackrel{\mathrm{def}}{=} C_i \oplus \widehat{V}_i$.

The following lemma tells us what probability we have to analyze.

Lemma 1. *Let a_q be $\mathrm{dist}(S^q, \mathrm{uni}(N^q))$, where $\mathrm{uni}(N^q)$ consists of all distinct elements among N^q. I.e., a_q means that all elements in $\{S_1, \ldots, S_q, N_1, \ldots, N_q\}$ are distinct except for the collisions between N_is. Similarly, let b_q denote $\mathrm{dist}(U^q, \mathrm{uni}(V^q))$. Here, if $\mathrm{uni}(N^q) = (N_{i_1}, \ldots, N_{i_\theta})$ for some $\{i_1, \ldots, i_\theta\} \subseteq \{1, \ldots, q\}$, then $\mathrm{uni}(V^q) = (V_{i_1}, \ldots, V_{i_\theta})$. Then, we have*

$$\mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}],\mathrm{TW}[\mathsf{P},\mathsf{R}]}(q) \le \nu(\mathrm{TW}[\mathsf{P},\mathsf{R}], \overline{a_q \wedge b_q}). \quad (8)$$

Proof. Let us consider the following probabilistic functions: $\Sigma^n \times \{0,1,2\} \to \Sigma^n$.

$$\mathrm{PP}(x,w) = \begin{cases} \mathsf{P}(x) & \text{if } w = 0 \text{ or } 2, \\ \mathsf{P}^{-1}(x) & \text{if } w = 1, \end{cases} \qquad \mathrm{PR}(x,w) = \begin{cases} \mathsf{P}(x) & \text{if } w = 0, \\ \mathsf{P}^{-1}(x) & \text{if } w = 1, \\ \mathsf{R}(x) & \text{if } w = 2. \end{cases}$$

Here, P and R are independent n-bit URP and URF. Observe that there exists a procedure, $\mathbb{F}$, such that $\mathrm{TW}[\mathsf{P},\mathsf{P}]$ ($\mathrm{TW}[\mathsf{P},\mathsf{R}]$) is equivalent to $\mathbb{F}[\mathrm{PP}]$ ($\mathbb{F}[\mathrm{PR}]$). Consider the game of distinguishing PP from PR using CPA (note that this game is quite easy to win). For PP and PR, let $(X_i, W_i) \in \Sigma^n \times \{0,1,2\}$ be the i-th query, and $Y_i \in \Sigma^n$ be the i-th output. For convenience, we allow adversaries to make colliding queries having $W_i = 2$ such as $(X_1, 2)$ and $(X_2, 2)$ where $X_1 = X_2$. Let $\mathcal{I} = \{i \in \{1, \ldots, q\} : W_i \in \{0, 2\}\}$. Let a'_q be the event that all

X_is with $i \in \mathcal{I}$ are distinct, except for the trivial collisions (i.e., $X_i = X_j$ such that $W_i = W_j = 2$ and $i \neq j$). Similarly, b'_q denotes the event that all Y_is with $i \in \mathcal{I}$ are distinct, except for the trivial collisions. Note that a'_q is equivalent to b'_q in PP, but not in PR. Then, for two MESs $\mathcal{A}' = a'_0 a'_1 \dots$ and $\mathcal{B}' = b'_0 b'_1 \dots$,

$$\mathrm{PP}|\mathcal{A}' \wedge \mathcal{B}' \equiv \mathrm{PP}|\mathcal{A}' \equiv \mathrm{PR}|\mathcal{A}' \wedge \mathcal{B}' \tag{9}$$

holds. Let Z^q be the q-th transcript (X^q, W^q, Y^q). Then, we obtain

$$P^{\mathrm{PR}}_{a'_q b'_q | Z^{q-1} X_q w_q a'_{q-1} b'_{q-1}} \leq P^{\mathrm{PP}}_{a'_q | Z^{q-1} X_q w_q a'_{q-1}} \tag{10}$$

since the r.h.s. of Eq. (10) is always 0 or 1 and if it is 0, then the l.h.s. is also 0 (recall that Eq. (10) means that the inequality holds for all possible arguments). From Eqs. (9) and (10) and Lemma 3, there is an MES defined for PP, $\mathcal{C}'$, such that

$$\mathrm{PP}^{\mathcal{A}' \wedge \mathcal{B}' \wedge \mathcal{C}'} \equiv \mathrm{PP}^{\mathcal{A}' \wedge \mathcal{C}'} \equiv \mathrm{PR}^{\mathcal{A}' \wedge \mathcal{B}'} \tag{11}$$

holds true. It is easy to see that $\mathcal{A}' \wedge \mathcal{B}'$ is equivalent to $\mathcal{A} \wedge \mathcal{B}$ where $\mathcal{A} = a_0 a_1 \dots$ and $\mathcal{B} = b_0 b_1 \dots$ are defined for $\mathrm{TW}[\mathsf{P}, \mathsf{P}]$ and $\mathrm{TW}[\mathsf{P}, \mathsf{R}]$. From this fact, and Eq. (11), and Lemma 4, we obtain

$$\mathrm{TW}[\mathsf{P}, \mathsf{P}]^{\mathcal{A} \wedge \mathcal{B} \wedge \mathcal{C}} \equiv \mathrm{TW}[\mathsf{P}, \mathsf{R}]^{\mathcal{A} \wedge \mathcal{B}}, \text{ for some MES } \mathcal{C} \text{ defined for } \mathrm{TW}[\mathsf{P}, \mathsf{P}]. \tag{12}$$

Combining Eq. (12) and Theorem 3 proves the lemma. □

Next, we have to evaluate the r.h.s. of Eq. (8). Our result is the following.

Lemma 2. *Let γ, ϵ, and ρ be given such that $F_{K''}$ satisfies the following three conditions when $V, V' \in_{\mathrm{u}} \Sigma^n$ and V is independent of V'.*

1. $\max_{l \in \Sigma^n, c \in \Sigma^n} \Pr(F_{K''}(l, V) = c) \leq \gamma$ *(here, probability is defined by $K'' \in_{\mathrm{u}} \mathcal{K}''$ and $V \in_{\mathrm{u}} \Sigma^n$).*
2. $\max_{l, l' \in \mathcal{L}, l \neq l', c \in \Sigma^n} \Pr(F_{K''}(l, V) \oplus F_{K''}(l', V) = c) \leq \epsilon$ *and* $\max_{l, l' \in \mathcal{L}, c \in \Sigma^n} \Pr(F_{K''}(l, V) \oplus F_{K''}(l', V') = c) \leq \epsilon$.
3. $\max_{l \in \mathcal{L}, c \in \Sigma^n} \Pr(F_{K''}(l, V) \oplus V = c) \leq \rho$.

Then we have

$$\nu(\mathrm{TW}[\mathsf{P}, \mathsf{R}], \overline{a_q \wedge b_q}) \leq \left(\gamma + \epsilon + \rho + \frac{1}{2^{n+1}}\right) q^2. \tag{13}$$

Proof. Let $\widetilde{\mathsf{P}}$ be the n-bit tweakable URP with tweak space $\mathcal{T} = \mathcal{L} \times \Sigma^n$. All variables and events defined for $\mathrm{TW}[\mathsf{P}, \mathsf{R}]$ are similarly defined for $\widetilde{\mathsf{P}}$ by using dummy functions. For example, $S_i = F_{K''}(L, \mathsf{R}(N)) \oplus M_i$ and $a_q = \mathrm{dist}(S^q, \mathrm{uni}(N^q))$. Note that $\mathrm{dist}(S^q)$ and $\mathrm{dist}(U^q)$ are equivalent in $\mathrm{TW}[\mathsf{P}, \mathsf{R}]$, but not in $\widetilde{\mathsf{P}}$. Let Z^q be the q-th transcript (M^q, C^q, T^q), and X_q be the q-th query (i.e., X_q is (M_q, T_q) or (C_q, T_q)), and Y_q be the q-th answer from the oracle, which is M_q or C_q. Let K_Δ be the key of Δ, which determines the instance of $(\mathsf{R}, F_{K''})$. From the assumption, K_Δ is uniformly distributed over $\mathcal{K}_\Delta \stackrel{\mathrm{def}}{=} \{f : \Sigma^n \to \Sigma^n\} \times \mathcal{K}''$ and independent of P.

Then, it is easy to verify that $P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{Y_q|Z^{q-1}X_qa_qb_qK_\Delta} = P^{\widetilde{\mathsf{P}}}_{Y_q|Z^{q-1}X_qa_qb_qK_\Delta}$ and $P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{K_\Delta|Z^{q-1}X_qa_qb_q} = P^{\widetilde{\mathsf{P}}}_{K_\Delta|Z^{q-1}X_qa_qb_q}$ hold. Therefore, we have

$$P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{Y_q|Z^{q-1}X_qa_qb_q} = \sum_{K_\Delta} P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{Y_q|Z^{q-1}X_qa_qb_qK_\Delta} \cdot P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{K_\Delta|Z^{q-1}X_qa_qb_q} \tag{14}$$

$$= \sum_{K_\Delta} P^{\widetilde{\mathsf{P}}}_{Y_q|Z^{q-1}X_qa_qb_qK_\Delta} \cdot P^{\widetilde{\mathsf{P}}}_{K_\Delta|Z^{q-1}X_qa_qb_q} = P^{\widetilde{\mathsf{P}}}_{Y_q|Z^{q-1}X_qa_qb_q}, \tag{15}$$

where summations are taken for all $\delta \in \mathcal{K}_\Delta$. This indicates the following conditional equivalence.

$$\mathrm{TW}[\mathsf{P},\mathsf{R}]|\mathcal{A} \wedge \mathcal{B} \equiv \widetilde{\mathsf{P}}|\mathcal{A} \wedge \mathcal{B}. \tag{16}$$

Then, we determine if

$$P^{\widetilde{\mathsf{P}}}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}K_\Delta} \leq P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}K_\Delta} \tag{17}$$

holds. We first analyze the r.h.s. of Eq. (17). Let us assume the variables in the condition are fixed such as $(Z^{q-1}, X_q, K_\Delta) = (z^{q-1}, x^q, \delta)$ and the q-th query is a chosen-plaintext query. Then, all variables except U_q are uniquely determined. Therefore, whether $a_q^+ \stackrel{\mathrm{def}}{=} a_q \wedge \mathrm{dist}(U^{q-1}, \mathrm{uni}(V^q))$ holds or not is a function of (Z^{q-1}, X_q, K_Δ). If a_q^+ holds, then U_q is uniform over $\Omega \stackrel{\mathrm{def}}{=} \Sigma^n \setminus \{U_1 \ldots, U_{q-1}\}$, and $a_q \wedge b_q$ occurs if $U_q \in \Omega \setminus \{V_1, \ldots, V_{q-1}\}$. Note that $\{U_1 \ldots, U_{q-1}\} \cap \{V_1, \ldots, V_q\} = \emptyset$ if a_q^+ holds. From these observations, we have

$$P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}K_\Delta} = \begin{cases} 0 & \text{if } a_q^+ \text{ does not hold,} \\ 1 - \frac{\theta}{2^n-(q-1)} & \text{otherwise ,} \end{cases} \tag{18}$$

where θ denotes the number of unique elements among $\{V_1, \ldots, V_q\}$. How about the l.h.s. of Eq. (17)? The occurrence of a_q^+ is a function of (Z^{q-1}, X_q, K_Δ) as well as the r.h.s. However, the distribution of U_q is different. Let Ψ be a set of indexes defined by $\Psi = \{i \in \{1, \ldots, q-1\} : T_q = T_i\}$ and let $|\Psi|$ be ψ. If a_q^+ holds, U_q is uniform over $\Omega' \stackrel{\mathrm{def}}{=} \Sigma^n \setminus \{U_i : i \in \Psi\}$ and $a_q \wedge b_q$ occurs if $U_q \in \Omega' \setminus \{\{V_1, \ldots, V_q\} \cup \{U_i : i \in \Psi^c\}\}$. Therefore, we have

$$P^{\widetilde{\mathsf{P}}}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}K_\Delta} = \begin{cases} 0 & \text{if } a_q^+ \text{ does not hold,} \\ 1 - \frac{\theta+q-\psi-1}{2^n-\psi} & \text{otherwise .} \end{cases} \tag{19}$$

Note that $0 \leq \psi \leq q-1$ and $1 \leq \theta \leq q$. Thus, when $q \leq 2^n - \theta + 1$ we obtain

$$\frac{\theta+q-\psi-1}{2^n-\psi} - \frac{\theta}{2^n-(q-1)} = \frac{(q-\psi-1)\cdot(2^n-(q-1)-\theta)}{(2^n-\psi)\cdot(2^n-(q-1))} \geq 0. \tag{20}$$

Since $\theta \leq q$, Eq. (20) holds unless $q > 2^{n-1} + 0.5$. The same analysis holds when the q-th query is a chosen-ciphertext query. Therefore, Eq. (17) holds if $q \leq 2^{n-1}$.

It is almost trivial to see that $P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{K_\Delta|Z^{q-1}X_qa_{q-1}b_{q-1}} = P^{\widetilde{\mathsf{P}}}_{K_\Delta|Z^{q-1}X_qa_{q-1}b_{q-1}}$. By combining this and Eqs. (17), we have

$$P^{\widetilde{\mathsf{P}}}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}} \le P^{\mathrm{TW}[\mathsf{P},\mathsf{R}]}_{a_qb_q|Z^{q-1}X_qa_{q-1}b_{q-1}}, \quad \text{if } q \le 2^{n-1}. \tag{21}$$

From this inequality, and Eq. (16), and Lemma 3, $\mathrm{TW}[\mathsf{P},\mathsf{R}]^{\mathcal{A}\wedge\mathcal{B}\wedge\mathcal{C}} \equiv \widetilde{\mathsf{P}}^{\mathcal{A}\wedge\mathcal{B}}$ holds for some MES $\mathcal{C} = c_0c_1\ldots$, if $q \le 2^{n-1}$. Therefore, using Lemma 5, we obtain

$$\nu(\mathrm{TW}[\mathsf{P},\mathsf{R}], \overline{a_q \wedge b_q}) \le \nu(\mathrm{TW}[\mathsf{P},\mathsf{R}], \overline{a_q \wedge b_q \wedge c_q}) = \nu(\widetilde{\mathsf{P}}, \overline{a_q \wedge b_q}), \text{ if } q \le 2^{n-1}.$$

Note that any adversary's strategy against $\widetilde{\mathsf{P}}$ must be independent of R and $F_{K''}$, as they do not affect the input or output of $\widetilde{\mathsf{P}}$. Therefore, evaluating $\nu(\widetilde{\mathsf{P}}, \overline{a_q \wedge b_q})$ is quite easy: we only have to consider non-adaptive strategies. Let $P^{\widetilde{\mathsf{P}}}$ denote the probability space defined by $\widetilde{\mathsf{P}}$ and some fixed q inputs. Then, the second condition of Lemma 2 implies that $P^{\widetilde{\mathsf{P}}}(S_i = S_j) \le \epsilon$ and $P^{\widetilde{\mathsf{P}}}(U_i = U_j) \le \epsilon$ if $i \ne j$. Moreover, we have $P^{\widetilde{\mathsf{P}}}(S_i = N_j) = P^{\widetilde{\mathsf{P}}}(\widehat{V}_i = N_j \oplus M_i) \le \gamma$ for any i, j, and $P^{\widetilde{\mathsf{P}}}(U_i = V_j) = P^{\widetilde{\mathsf{P}}}(\widehat{V}_i = C_i \oplus V_j) \le \rho$ for any i, j (more precisely, it is at most ρ if $N_i = N_j$ and $1/2^n$ otherwise). Therefore, if $q \le 2^{n-1}$, we have

$$\begin{aligned}\nu(\widetilde{\mathsf{P}}, \overline{a_q \wedge b_q}) &\le P^{\widetilde{\mathsf{P}}}(\overline{\mathrm{dist}(S^q)}) + P^{\widetilde{\mathsf{P}}}(\overline{\mathrm{dist}(U^q)}) + P^{\widetilde{\mathsf{P}}}(S_i = N_j \text{ for some } i, j \le q) \\ &\quad + P^{\widetilde{\mathsf{P}}}(U_i = V_j \text{ for some } i, j \le q) + P^{\widetilde{\mathsf{P}}}(V_i = V_j \text{ for some } i, j \le q,\ i \ne j) \\ &\le 2\binom{q}{2}\epsilon + q^2\gamma + q^2\rho + \binom{q}{2}2^{-n} \le q^2(\epsilon + \gamma + \rho + 2^{-n-1}). \end{aligned} \tag{22}$$

This upper bound reaches 1 if $q \sim 2^{n/2}$, thus the condition $q \le 2^{n-1}$ is redundant. This concludes the proof. □

Note that the second condition of Lemma 2 implies that the offset function is ϵ-AXU for input $T = (L, N) \in \mathcal{L} \times \Sigma^n$. By combining Eq. (7) and Lemmas 1, 2 and Theorem 1, the security of $\mathrm{TW}[\mathsf{P},\mathsf{P}]$ is proved in the following theorem.

Theorem 4. *If the assumption of Lemma 2 holds true for $F_{K''}$, we have*

$$\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}]}(q) \le \left(2\epsilon + \gamma + \rho + \frac{1}{2^{n+1}}\right) q^2. \tag{23}$$

The proof of Theorem 2. From Theorem 4, we can easily see that the following offset function enables a simple one-key tweakable block cipher.

Corollary 1. *Let* $\mathrm{TW}[\mathsf{P},\mathsf{P}]$ *use the offset function defined as* $\Delta(T) = L \cdot E_K(N)$, *where* $T = (L, N) \in (\Sigma^n \setminus \{0, 1\}) \times \Sigma^n$. *Then, we have* $\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{TW}[\mathsf{P},\mathsf{P}]}(q) \le \frac{4.5q^2}{2^n}$. *Moreover, for any block cipher* E_K,

$$\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\mathrm{TW}[E_K,E_K]}(q, \tau) \le \mathtt{Adv}^{\mathtt{sprp}}_{E_K}(2q, \tau') + \frac{4.5q^2}{2^n}, \text{ where } \tau' = \tau + O(q).$$

Proof. Note that $L \cdot V$, $L \cdot V \oplus L' \cdot V$, and $L \cdot V \oplus V$ are permutations of V for any $L, L' \in \Sigma^n \setminus \{0, 1\}$ such that $L \neq L'$. This indicates $\epsilon = \gamma = \rho = 1/2^n$ and thus Theorem 4 proves the first claim. The second claim follows from the first and the standard conversion from the information-theoretic setting to the computational setting. □

Recall that an output of XEX's offset function is $\prod_{j=1}^{d} \alpha_j^{i_j} \cdot E_K(N)$, where a tweak is $(i_1, \ldots, i_d, N)$. In the fixed XEX, $\prod_{j=1}^{d} \alpha_j^{i_j} \neq \prod_{j=1}^{d} \alpha_j^{i'_j}$ whenever $(i_1, \ldots, i_d) \neq (i'_1, \ldots, i'_d)$, and $\prod_{j=1}^{d} \alpha_j^{i_j}$ never be 0 or 1 (in $\mathrm{GF}(2^n)$). Therefore, Theorem 2 is immediately obtained from Corollary 1.

Applications of Theorem 4. Theorem 4 provides not only the improved proof of XEX, but also useful tools for the design of strong tweakable block cipher. For example, consider the LRW mode based on a dedicated AXU hash function such as MMH or NMH (see e.g., [2]). Then, Theorem 4 tells us what properties are needed (in addition to the AXU property) if we want to substitute (a part of) the key of LRW's offset function with an encryption of the block cipher. This is achieved by our generalized construction. In particular, for the offset function of the form $\Delta(L, N) = g(L \oplus E_K(N))$ where g is a fixed n-bit permutation, the conditions of Lemma 2 become simpler: since $g(l \oplus V)$ and $g(l \oplus V) \oplus g(l' \oplus V')$ are permutations of V, γ and the second ϵ in the second condition are naturally $1/2^n$. The remaining conditions can be interpreted such that g is differentially ϵ-uniform [15] and is a $(2^n\rho - 1)$-almost orthomorphism [21] (equivalently, a permutation with maximum self-differential probability ρ [13], where self-differential means the differential between the input and output). An example of such a permutation is the inversion on $\mathrm{GF}(2^n)$, $\mathrm{inv}(*)$, where $\mathrm{inv}(x) = x^{-1}$ if $x \neq 0$, and $\mathrm{inv}(0) = 0$. If g is the inversion on $\mathrm{GF}(2^n)$, $\epsilon = 4/2^n$ holds from [15], and a simple analysis proves that $\rho = 3/2^n$. Consequently, the mode with the offset function $\Delta(L, N) = \mathrm{inv}(L \oplus E_K(N))$ is provably secure and has the bound $(2\epsilon + \gamma + \rho + 0.5)\frac{q^2}{2^n} = \frac{12.5q^2}{2^n}$. This demonstrates that strong tweakable ciphers with arbitrary tweak update are possible from permutations with good differential and self-differential property. We will use this idea in the next section.

5 Improving LRW-AES

Theorem 4 also gives some improvements to the LRW-AES described in Sect. 3.2. Here, we propose two improvements.

LRW-AES-Sqr: One-Key LRW-AES having 2^n tweak values. As mentioned, LRW-AES is the mode for AES that provides a strong tweakable block cipher using $\Delta(T) = K_\Delta \cdot T$, where $T \in \Sigma^n$ and $K_\Delta \in_{\mathrm{u}} \Sigma^n$ is independent of the key of the AES. Although the original LRW-AES needs two keys, Corollary 1 provides some ways to reduce these two keys to the one AES key. The simplest fix is the same as one used for the XEX: let $K_\Delta = E_K(0)$ and $T \in \Sigma^n \setminus \{0, 1\}$. However, the resulting mode is not strictly compatible with LRW-AES because

of the reduced tweak set. However, we still have several options to achieve one-key LRW-AES having 2^n tweak values. An efficient one among these options is to use squaring, which is as follows. We first generate $V = E_K(0)$ in the preprocessing. For tweak $T \in \Sigma^n$, the offset function is defined as:

$$\Delta(T) = \begin{cases} V^2 & \text{if } T = 0, \\ a \cdot V^2 & \text{if } T = 1, \\ T \cdot V & \text{otherwise.} \end{cases} \tag{24}$$

Here, a is a fixed element of $\Sigma^n \setminus \{0, 1\}$. This requires only one AES encryption in the preprocessing, and the cost for updating a tweak (i.e., the cost for computing $\Delta(T)$) is almost the same as that of the original LRW-AES, namely one GF multiplication. To be precise, the computation of $a \cdot V^2$ requires two multiplications; however the cost for multiplication by constant a can be negligibly small with the powering-up construction. The security of this scheme, which we call LRW-AES-Sqr, is proved as follows.

Theorem 5. $\widetilde{\mathtt{Adv}^{\mathtt{sprp}}_{\text{LRW-AES-Sqr}}}(q, \tau) \leq \mathtt{Adv}^{\mathtt{sprp}}_{\text{AES}_K}(q + 1, \tau') + \frac{7.5q^2}{2^{128}}$, where $\tau' = \tau + O(q)$.

Proof. We apply Lemma 2 to the offset function in Eq. (24). Since squaring in a field with characteristic two is a permutation, both V^2 and $a \cdot V^2$ are permutations of V. Also, $T \cdot V$ with $T \notin \{0, 1\}$ is a permutation. Thus we have $\gamma = 1/2^n$. Every sum of two offset values (i.e., $T \cdot V \oplus T' \cdot V$, $V^2 \oplus a \cdot V^2$, $V^2 \oplus T \cdot V$, and $a \cdot V^2 \oplus T \cdot V$ for any $T, T' \in \Sigma^n \setminus \{0, 1\}$ with $T \neq T'$) is a quadratic or linear function of V, but can not be reduced to a constant since $a \notin \{0, 1\}$. As a function with degree d has at most d solutions, every sum has bias of at most $2/2^n$, which means $\epsilon = 2/2^n$. Moreover, both $V^2 \oplus V$ and $a \cdot V^2 \oplus V$ have bias $2/2^n$, and $T \cdot V \oplus V$ with $T \notin \{0, 1\}$ has bias $1/2^n$. Therefore we have $\rho = 2/2^n$. Note that AES is invoked $q + 1$ times in LRW-AES-Sqr. Combining these facts and Theorem 4 proves the theorem. □

LRW-AES-4r: LRW-AES WITHOUT MULTIPLICATION. Both LRW-AES and LRW-AES-Sqr require GF multiplication in order to be able to update a tweak arbitrarily. Here, we provide an interesting alternative to the multiplication: the reduced-round of AES. This idea is basically the same as the recent proposal of AES-based message authentication codes [13]. More precisely, what we use is the 4-round AES, denoted by $\text{AES}^{(4)}_{K_{\text{sub}}}$, where $K_{\text{sub}} \in_{\text{u}} (\Sigma^{128})^3$ consists of the round keys for the last three rounds. The first round key is set to 0. We first generate $V = \text{AES}_K(0)$ and $K_{\text{sub}} \in_{\text{u}} (\Sigma^{128})^3$. For tweak $T \in \Sigma^{128}$, we use the offset function such as $\Delta(T) = \text{AES}^{(4)}_{K_{\text{sub}}}(T \oplus V)$. This scheme, which we call LRW-AES-4r is essentially AES-based while the cost for updating a tweak is less than an AES encryption. XEX mode also has this property (if we fix N to some constant), but a tweak can be updated only in an incremental order.

SECURITY OF LRW-AES-4r. The differential and linear properties of the AES and its reduced-round version have been extensively studied. Particularly,

Table 1. Mean speed of LRW-AES and our improvements for random 2^{20} messages and tweaks on a PC (Pentium III (Coppermine), 1 GHz, 16KB L1 cache). Alg 1 and 2 denote the multiplication algorithms specified in [23]. Preprocessing includes key schedulings for both AES and its inverse, and precomputation for multiplication, and one AES encryption: $V = \mathrm{AES}_K(0)$.

Mode	Preproc (cycles)	Enc (cycle/byte)	Dec (cycle/byte)
LRW-AES (alg 1)	1248	234	241
LRW-AES (alg 2)	289506	155	161
LRW-AES-Sqr (alg 1)	1696	235	241
LRW-AES-Sqr (alg 2)	289966	155	161
LRW-AES-4r	1653	39	45

Keliher proved that the maximum expected differential probability of $\mathrm{AES}^{(4)}_{K_{\mathrm{sub}}}$ was at most 2^{-113} [8], if $K_{\mathrm{sub}} \in_{\mathrm{u}} (\Sigma^{128})^3$. This means that $\mathrm{AES}^{(4)}_{K_{\mathrm{sub}}}(T \oplus V)$ is 2^{-113}-AXU, when (K_{sub}, V) is the key and T is the input.

The security of LRW-AES-4r is proved as follows.

Theorem 6. $\widetilde{\mathtt{Adv}^{\mathtt{sprp}}_{\text{LRW-AES-4r}}}(q,\tau) \le \mathtt{Adv}^{\mathtt{sprp}}_{\mathrm{AES}_K}(q+1,\tau') + \frac{(2^{16}+2.5)q^2}{2^{128}}$, where $\tau' = \tau + O(q)$.

Proof. We have $\epsilon < 2^{-113} = 2^{15}/2^{128}$ from [8]. Moreover, we have $\gamma = 1/2^{128}$. Note that the output of $\mathrm{AES}^{(4)}_{K_{\mathrm{sub}}}$ is completely random and independent of the input, as each round key is XORed to the intermediate message and uniformly distributed. This indicates $\rho = 1/2^{128}$. □

We have to mention that LRW-AES-4r is not an ideal substitute for the LRW-AES. The security of the LRW-AES-4r is moderately degraded compared with the original LRW-AES. That is, LRW-AES-4r has $112/2 = 56$-bit security (i.e., q must be much smaller than 2^{56}), while the original LRW-AES has 63-bit security. This means that the lifetime of key should be slightly shortened. In addition, the key of the LRW-AES-4r is longer (512 bits) than that of the LRW-AES (256 bits), though both require only one AES key scheduling. We implemented our proposals and the original LRW-AES in software. Our implementation was based on the reference AES code [22]. We used two naive algorithms for multiplication in $\mathrm{GF}(2^{128})$ that were specified in the document of LRW-AES [23]. The performance of LRW-AES-4r is quite remarkable, as Table 1 shows.

Acknowledgments

We would like to thank Peng Wang for pointing out the reference on XEX. We also thank Etsuko Tsujihara for the implementation and anonymous reviewers for very useful comments.

References

1. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A Concrete Security Treatment of Symmetric Encryption. In: Proceedings of the 38th Annual Symposium on Foundations of Computer Science, FOCS '97, pp. 394–403 (1997)
2. Black, J.: Message Authentication Code. PhD dissertation (2000)
3. Goldreich, O.: Modern Cryptography, Probabilistic Proofs and Pseudorandomness. Springer, Heidelberg
4. Halevi, S., Rogaway, P.: A Tweakable Enciphering Mode. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 482–499. Springer, Heidelberg (2003)
5. Halevi, S., Rogaway, P.: A Parallelizable Enciphering Mode. In: Okamoto, T. (ed.) CT-RSA 2004. LNCS, vol. 2964, pp. 292–304. Springer, Heidelberg (2004)
6. Iwata, T., Kurosawa, K.: On the Universal Hash Functions in Luby-Rackoff Cipher. In: Lee, P.J., Lim, C.H. (eds.) ICISC 2002. LNCS, vol. 2587, pp. 226–236. Springer, Heidelberg (2003)
7. Iwata, T., Kurosawa, K.: OMAC: One-Key CBC MAC. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 129–153. Springer, Heidelberg (2003)
8. Keliher, L., Sui, J.: Exact Maximum Expected Differential and Linear Probability for 2-Round Advanced Encryption Standard (AES). IACR ePrint Archive (2005)/321
9. Kilian, J., Rogaway, P.: How to Protect DES Against Exhaustive Key Search. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 252–267. Springer, Heidelberg (1996)
10. Liskov, M., Rivest, R., Wagner, D.: Tweakable Block Ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer, Heidelberg (2002)
11. Maurer, U.: Indistinguishability of Random Systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002)
12. Maurer, U., Pietrzak, K.: Composition of Random Systems: When Two Weak Make One Strong. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 410–427. Springer, Heidelberg (2004)
13. Minematsu, K., Tsunoo, Y.: Provably Secure MACs From Differentially-uniform Permutations and AES-based Implementations. In: Robshaw, M. (ed.) FSE 2006. LNCS, vol. 4047, Springer, Heidelberg (2006)
14. Naor, M., Reingold, O.: On the Construction of Pseudorandom Permutations: Luby-Rackoff Revisited. Journal of Cryptology 12(1), 29–66 (1999)
15. Nyberg, K.: Differentially Uniform Mappings for Cryptography. In: Helleseth, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 55–64. Springer, Heidelberg (1994)
16. Pietrzak, K.: Composition Does Not Imply Adaptive Security. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 55–65. Springer, Heidelberg (2005)
17. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: ACM Conference on Computer and Communications Security ACM CCS'01, pp. 196–205 (2001)
18. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC (the early version of [19]), `http://www.cs.ucdavis.edu/~rogaway/papers/offsets.pdf`
19. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 16–31. Springer, Heidelberg (2004)
20. Wegman, M., Carter, L.: New Hash Functions and Their Use in Authentication and Set Equality. Journal of Computer and System Sciences 22, 265–279 (1981)

21. Vaudenay, S.: On the Lai-Massey Scheme. In: Lam, K.-Y., Okamoto, E., Xing, C. (eds.) ASIACRYPT 1999. LNCS, vol. 1716, pp. 9–19. Springer, Heidelberg (1999)
22. http://homes.esat.kuleuven.be/~rijmen/rijndael/rijndael-fst-3.0.zip
23. Draft Proposal for Tweakable Narrow-block Encryption (2004), http://www.siswg.org/docs/LRW-AES-10-19-2004.pdf

A Theorems and Lemmas Proved by Maurer [11]

Let us now describe some of Maurer's results [11] other than Theorem 3. They were used in our analysis.

Lemma 3. *(Lemma* 1 (iv) *of [11]) Let MESs $\mathcal{A}$ and $\mathcal{B}$ be defined for F and G. Moreover, let X_i and Y_i denote the i-th input and output of F (or G), respectively. Assume $F|\mathcal{A} \equiv G|\mathcal{B}$. If $P^F_{a_i|X^iY^{i-1}a_{i-1}} \leq P^G_{b_i|X^iY^{i-1}b_{i-1}}$ for $i \geq 1$, which means $P^F_{a_i|X^iY^{i-1}a_{i-1}}(x^i, y^{i-1}) \leq P^G_{b_i|X^iY^{i-1}b_{i-1}}(x^i, y^{i-1})$ holds for all (x^i, y^{i-1}) such that $P^F_{a_{i-1}|X^{i-1}Y^{i-1}}(x^{i-1}, y^{i-1})$ and $P^G_{b_{i-1}|X^{i-1}Y^{i-1}}(x^{i-1}, y^{i-1})$ are positive, then there exists an MES $\mathcal{C}$ defined for G such that $F^{\mathcal{A}} \equiv G^{\mathcal{B}\wedge\mathcal{C}}$.*

Lemma 4. *(Lemma* 4 (ii) *of [11]) Let F and G be two compatible keyed functions, and $\mathbb{F}$ be the function of F and G (i.e., $\mathbb{F}[F]$ is a function that internally invokes F, possibly several times, to process its inputs). Here, $\mathbb{F}$ can be probabilistic, and if so, we assume $\mathbb{F}$ is independent of F or G. If $F^{\mathcal{A}} \equiv G^{\mathcal{B}}$ holds for some MESs $\mathcal{A}$ and $\mathcal{B}$, we have $\mathbb{F}[F]^{\mathcal{A}'} \equiv \mathbb{F}[G]^{\mathcal{B}'}$. Here, MES $\mathcal{A}' = a'_0a'_1\ldots$ is defined such that a'_i denotes $\mathcal{A}$-event is satisfied for the time period i. For example, if $\mathbb{F}[F]$ always invoke F c times for any input, then $a'_i = a_{ci}$. $\mathcal{B}'$ is defined in the same way.*

Lemma 5. *(Lemma* 6 (ii) *of [11]) If $F^{\mathcal{A}} \equiv G^{\mathcal{B}}$, then $\nu(F, \overline{a_q}) = \nu(G, \overline{b_q})$.*

Lemma 6. *(Lemma* 6 (iii) *of [11]) $\nu(F, \overline{a_q \wedge b_q}) \leq \nu(F, \overline{a_q}) + \nu(F, \overline{b_q})$.*

B Proof of Theorem 1

The structure of the proof is almost the same as the proofs of Lemmas 1 and 2. Let $\widetilde{E}$ denote the LRW mode using the offset function Δ, and $\widetilde{\mathsf{P}}$ denote the URP compatible (i.e., the block size and tweak space are the same as those of $\widetilde{E}$) with $\widetilde{E}$. Let M_i and C_i be the i-th plaintext and ciphertext, and let T_i be the i-th tweak. Let S_i be the i-th input to E_K, i.e., $S_i = \Delta(T_i) \oplus M_i$. Similarly, we define $U_i = \Delta(T_i) \oplus C_i$. Note that these variables can be defined for both $\widetilde{E}$ and $\widetilde{\mathsf{P}}$. We use two MESs $\mathcal{A} = a_0a_1\ldots$ and $\mathcal{B} = b_0b_1\ldots$ where $a_i \stackrel{\text{def}}{=} \mathrm{dist}(S^i)$ and $b_i \stackrel{\text{def}}{=} \mathrm{dist}(U^i)$. An analysis similar to that used in the proof of Lemma 2 provides that the equivalences $\widetilde{E}|\mathcal{A} \equiv \widetilde{\mathsf{P}}|\mathcal{A}\wedge\mathcal{B}$ and $\widetilde{E}^{\mathcal{A}\wedge\mathcal{C}} \equiv \widetilde{\mathsf{P}}^{\mathcal{A}\wedge\mathcal{B}}$ hold for some MES $\mathcal{C} = c_0c_1\ldots$ defined for $\widetilde{E}$. Thus we have

$$\mathtt{Adv}^{\widetilde{\mathtt{sprp}}}_{\widetilde{E}}(q) = \mathtt{Adv}^{\widetilde{\mathtt{cca}}}_{\widetilde{E},\widetilde{\mathsf{P}}}(q) \leq \nu(\widetilde{\mathsf{P}}, \overline{a_q \wedge b_q}) \leq \nu(\widetilde{\mathsf{P}}, \overline{a_q}) + \nu(\widetilde{\mathsf{P}}, \overline{b_q}), \tag{25}$$

where the first inequality follows from $\widetilde{E}^{\mathcal{A}\wedge\mathcal{C}} \equiv \widetilde{\mathsf{P}}^{\mathcal{A}\wedge\mathcal{B}}$ and Theorem 3, and the last follows from Lemma 6. It is almost trivial to see that any adaptive strategy against $\widetilde{\mathsf{P}}$ to invoke $\overline{a_q}$ or $\overline{b_q}$ is no better than the best non-adaptive strategy. Therefore, we have

$$\nu(\widetilde{\mathsf{P}}, \overline{a_q}) \leq \max_{m^q, t^q} P^{\widetilde{\mathsf{P}}}(S_i = S_j \text{ for some } 1 \leq i < j \leq q | M^q = m^q, T^q = t^q) \leq \epsilon q^2/2,$$

where $P^{\widetilde{\mathsf{P}}}$ denotes the probability space defined by $\widetilde{\mathsf{P}}$ and the maximum is taken for all q plaintexts and tweaks satisfying $(m_i, t_i) \neq (m_j, t_j)$ for any $i \neq j$. The second inequality follows from the assumption on Δ. Similarly, we obtain $\nu(\widetilde{\mathsf{P}}, \overline{b_q}) \leq \epsilon q^2/2$, and thus, $\nu(\widetilde{\mathsf{P}}, \overline{a_q}) + \nu(\widetilde{\mathsf{P}}, \overline{b_q}) \leq \epsilon q^2$. This concludes the proof.

C An Attack Against OCB1 Using Flawed XEX^{-1}

OCB1 [18] is an authenticated encryption mode for any finite-length message. A ciphertext consists of a nonce, and an encryption of a message, and an authentication tag, which we simply call a tag. The OCB1 defined in [18] and [19] are slightly different, but our attack is applicable to both. For simplicity, we only describe a version of OCB1 defined in [18] for a message of length cn bits for some positive integer c. We also assume that the tag is n-bit. Let the message M be $(M[0], \ldots, M[c-1])$, where each $M[i] \in \Sigma^n$. Let $\widetilde{E}_K$ be an n-bit block strong tweakable block cipher having the tweak space $\{0, 1, \ldots, 2^{n/2}\} \times \{0, 1\} \times \Sigma^n$. To encrypt M with nonce $N \in \Sigma^n$, we first let $C[i] = \widetilde{E}_K(M[i], (i, 0, N))$, where the second argument of $\widetilde{E}_K$ is a tweak, for $i = 0, \ldots, c-2$. The last block, $M[c-1]$, is encrypted such as $C[c-1] = M[c-1] \oplus \widetilde{E}_K(v, (c-1, 0, N))$, where v denotes the bit length of the last block, which is assumed to be n, in some deterministic encoding. Then, we compute the sum of all message blocks, namely $\mathsf{sum} = M[0] \oplus M[1] \oplus \cdots \oplus M[c-1]$. The tag is $\mathsf{tag} = \widetilde{E}_K(\mathsf{sum}, (c-1, 1, N))$, and the ciphertext C is $(N, C[0], \ldots, C[c-1], \mathsf{tag})$. To decrypt it, we compute $\widehat{M}[i] = \widetilde{E}_K^{-1}(C[i], (i, 0, N))$ for $1 \leq i \leq c-2$. For C_{c-1}, we have $\widehat{M}[c-1] = C_{c-1} \oplus \widetilde{E}_K(v, (c-1, 0, N))$ and $\widehat{\mathsf{sum}} = \widehat{M}[0] \oplus \widehat{M}[1] \oplus \cdots \oplus \widehat{M}[c-1]$. Then, we check if $\widetilde{E}_K(\widehat{\mathsf{sum}}, (c-1, 1, N))$ and tag are the same. If they are the same, we say the ciphertext is authenticated, and otherwise it is faked.

XEX^{-1} gives ciphertext $C = E_K^{-1}(M \oplus \Delta(i_1, i_2, N)) \oplus \Delta(i_1, i_2, N)$ where $\Delta(i_1, i_2, N)$ equals $2^{i_1}3^{i_2}E_K(N)$ for all $(i_1, i_2) \in \{0, 1, \ldots, 2^{n/2}\} \times \{0, 1\}$. Recall that this provides unique representations but does not exclude a reduced-to-1 index vector. Our attack is against the tag-generation part and is based on the information of two ciphertexts. We assume the nonce is set to N at the beginning of the attack.

1. Ask the oracle (who implemented the XEX^{-1}-based OCB1) to encrypt a $2n$-bit message, $M_1 = (M_1[0], M_1[1]) = (0, m)$, for some $m \in \Sigma^n$ and receive the ciphertext $C_1 = (N, C_1[0], C_1[1], \mathsf{tag}_1)$, where $C_1[0] = \widetilde{E}_K(M_1[0], (0, 0, N)) = E_K(N) \oplus N$ and $C_1[1] = m \oplus E_K^{-1}(v \oplus 2 \cdot E_K(N)) \oplus 2 \cdot E_K(N)$. The tag is $\mathsf{tag}_1 = E_K^{-1}(m \oplus 2 \cdot 3 \cdot E_K(N)) \oplus 2 \cdot 3 \cdot E_K(N)$.

2. Ask the oracle to encrypt $M_2 = (M_2[0], M_2[1]) = (0, m')$ for some $m' \in \Sigma^n, m' \neq m$ and receive the ciphertext $C_2 = (N', C_2[0], C_2[1], \mathsf{tag}_2)$, where $C_2[0] = \widetilde{E}_K(M_2[0], (0, 0, N')) = E_K(N') \oplus N'$ and $N' \neq N$. We do not use $C_2[1]$ and tag_2.
3. Then, issue the faked ciphertext $C' = (N, C'[0], C'[1], \mathsf{tag}')$. Here, $C'[0] = C_1[0] \oplus N \oplus N'$, and $C'[1] = C_1[0] \oplus C_1[1] \oplus C_2[0] \oplus N \oplus N'$, and $\mathsf{tag}' = \mathsf{tag}_1$.

The above faked ciphertext will be always accepted as authentic by the oracle, since the decrypted message will be:

$$\widehat{M}'[0] = E_K(C'[0] \oplus E_K(N)) \oplus E_K(N) = E_K(N) \oplus E_K(N') \tag{26}$$

$$\widehat{M}'[1] = C'[1] \oplus E_K^{-1}(v \oplus 2 \cdot E_K(N)) \oplus 2 \cdot E_K(N) = E_K(N) \oplus E_K(N') \oplus m. \tag{27}$$

These equations indicate that $\widehat{\mathsf{sum}}' = \widehat{M}'[0] \oplus \widehat{M}'[1] = m$, and therefore, we have $\mathsf{tag}' = \mathsf{tag}_1$.