



Generic Attack on Iterated Tweakable FX Constructions

Ferdinand Sibleyras

► To cite this version:

Ferdinand Sibleyras. Generic Attack on Iterated Tweakable FX Constructions. CT-RSA 2020 - The Cryptographers' Track at the RSA Conference 2020, Feb 2020, San Francisco, United States. pp.1–14, 10.1007/978-3-030-40186-3_1 . hal-02424953

HAL Id: hal-02424953

<https://inria.hal.science/hal-02424953>

Submitted on 29 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Generic Attack on Iterated Tweakable FX Constructions

Ferdinand Sibleyras

Inria, France

Keywords: Tweakable · Block Cipher · Provable Security · FX · Cryptanalysis · Optimality · XHX2

Abstract. Tweakable block ciphers are increasingly becoming a common primitive to build new resilient modes as well as a concept for multiple dedicated designs. While regular block ciphers define a family of permutations indexed by a secret key, tweakable ones define a family of permutations indexed by both a secret key and a public tweak. In this work we formalize and study a generic framework for building such a tweakable block cipher based on regular block ciphers, the iterated tweakable FX construction, which includes many such previous constructions of tweakable block ciphers. Then we describe a cryptanalysis from which we can derive a provable security upper-bound for all constructions following this tweakable iterated FX strategy. Concretely, the cryptanalysis of r rounds of our generic construction based on n -bit block ciphers with κ -bit keys requires $\mathcal{O}(2^{\frac{r}{r+1}(n+\kappa)})$ online and offline queries. For $r = 2$ rounds this interestingly matches the proof of the particular case of XHX2 by Lee and Lee (ASIACRYPT 2018) thus proving for the first time its tightness. In turn, the XHX and XHX2 proofs show that our generic cryptanalysis is information theoretically optimal for 1 and 2 rounds.

1 Introduction

Tweakable block ciphers have been the focus of many recent works in the field of symmetric cryptography as it provides a very interesting flexibility compared to regular block ciphers. Formally, a block cipher is defined as a family of permutations indexed by a secret key, thus an n -bit block cipher E indexed by a κ -bit key is an application $E : \{0, 1\}^\kappa \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. Whereas a tweakable block cipher is a family of permutations indexed by both a secret key and a public tweak, thus an n -bit tweakable block cipher $\tilde{E}$ indexed by a $\tilde{\kappa}$ -bit secret key and a τ -bit public tweak is an application $\tilde{E} : \{0, 1\}^{\tilde{\kappa}} \times \{0, 1\}^\tau \times \{0, 1\}^n \rightarrow \{0, 1\}^n$. They have been formalized by Liskov, Rivest and Wagner [LRW11].

On the other hand, regular block ciphers benefit from a longer history of research which gave birth to many designs and implementations notably including the DES [DES77] and the AES [AES01]. Therefore a natural question is: how can we build a tweakable block cipher out of regular block ciphers? In fact

this line of study inspired new modes of operations like OCB [RBBK01] and PMAC [BR02] that benefits from a relatively easy two-step proof: first we show that the main construction is secure when used along with a tweakable block cipher then we construct such tweakable block cipher with a regular block cipher to fully describe the mode. A first approach can be to append a tweak with the secret key such that the concatenation becomes the effective key to the regular block cipher. Given security under related key attacks this can work but at the cost of security: the size of the secret key will have to be reduced to make space for the tweak.

To go around this limitation Liskov et al. described two constructions LRW1 and LRW2 [LRW11]. In particular LRW2 is somehow remindful of the FX construction that adds an n -bit key before the input and another after the output of the underlying block cipher. The FX construction has been proposed by Kilian and Rogaway [KR96] in a different context: they investigated DESX, an easy solution to protect DES against an exhaustive key search. FX consists in adding one n -bit subkey before and another one after the block cipher. With such strategy they proved that the time complexity of the best generic cryptanalysis goes from $\mathcal{O}(2^\kappa)$ to $\mathcal{O}(2^{\kappa+n}/D)$ where D is the data or online query complexity. The FX construction has since been notably used in PRINCE [BCG⁺12] and PRIDE [ADK⁺14]. We can naturally iterate r rounds of the FX construction which requires to have r κ -bit subkeys along with $(r + 1)$ n -bit subkeys. Then the idea to build a tweakable block cipher is to blend the tweak and the master key together in a predefined key schedule to obtain all the required subkeys for the computation.

1.1 Notations

First we formally describe the r -round tweakable iterated FX construction (Figure 2) on which our results apply. Let $E_{1,2,\dots,r}(u, \cdot)$ be r block ciphers with κ -bit key u and n -bit input and output. Let k be the κ -bit master key of the tweakable block cipher construction. Let t be a tweak of arbitrary length. Let $\gamma_i(k, t)$ be the subkey for the i^{th} block cipher of length κ -bit for $1 \leq i \leq r$ and $\lambda_i(k, t)$ the n -bit subkeys to XOR in the state for $0 \leq i \leq r$. For example the $r = 2$ -round tweakable FX construction (Figure 1) $\tilde{E}_k(t, m)$ is described as:

$$\tilde{E}_k(t, m) = E_2(\gamma_2(k, t), E_1(\gamma_1(k, t), m \oplus \lambda_0(k, t)) \oplus \lambda_1(k, t)) \oplus \lambda_2(k, t)$$

We will focus on generic key recovery attacks. The goal of the cryptanalysis of $\tilde{E}_k(t, m)$ is to recover k by doing offline queries to $E_{1,2,\dots,r}(\cdot, \cdot)$ and online queries to $\tilde{E}_k(\cdot, \cdot)$. We don't count the number of calls to the γ and λ functions generating the subkeys as queries because we don't assume any security property for them. In fact it is common for the subkeys to assume some almost uniformity, almost universality or almost XOR-universality property with respect to the tweak (See Definition 1). This makes the analysis proper for most of the constructions we cite except for $\tilde{F}[2]$ by Mennink [Men15] which can be seen as a 1-round tweakable FX where the subkey functions reuse the block cipher itself.

Definition 1. Let $\delta > 0$ and a function $\lambda : \mathcal{K} \times \mathcal{T} \rightarrow \mathcal{Y}$ for non-empty sets $\mathcal{K}, \mathcal{T}, \mathcal{Y}$.

– $\lambda(k, t)$ is said to be δ -almost uniform if for any $t \in \mathcal{T}$ and any $y \in \mathcal{Y}$,

$$\Pr(k \leftarrow_{\$} \mathcal{K} : \lambda(k, t) = y) \leq \delta .$$

– $\lambda(k, t)$ is said to be δ -almost universal (AU) if for any distinct t and $t' \in \mathcal{T}$,

$$\Pr(k \leftarrow_{\$} \mathcal{K} : \lambda(k, t) = \lambda(k, t')) \leq \delta .$$

– $\lambda(k, t)$ is said to be δ -almost XOR-universal (AXU) if for any distinct t and $t' \in \mathcal{T}$ and any $y \in \mathcal{Y}$,

$$\Pr(k \leftarrow_{\$} \mathcal{K} : \lambda(k, t) \oplus \lambda(k, t') = y) \leq \delta .$$

While our results do not depend on the repartition of the tweak space, having arbitrary long tweaks is justified by the **XTX** transformation of Minematsu and Iwata [MI15]. Indeed **XTX** transforms a tweakable block cipher with a tweak of limited length to one with a tweak of arbitrary length without, in our case, affecting the general iterated tweakable FX structure as it simply affects the subkey functions.




Fig. 1. 2-Round Tweakable FX.

1.2 Previous Works

In the same paper where they formalize the concept of tweakable block ciphers, Liskov, Rivest and Wagner proposed two constructions often known as **LRW1** and **LRW2** [LRW11]. **LRW1** consists in adding the tweak between two calls of the block cipher while **LRW2** evaluates a keyed universal hash function on the tweak and adds it twice: before the input and after the output of the block cipher. These modes are described as $\tilde{E}_k(t, m) = E_k(t \oplus E_k(m))$ and $\tilde{E}_k(t, m) = E_k(m \oplus h(t)) \oplus h(t)$ respectively with the requirement that h be an almost XOR-universal function. They also provide security proofs roughly up to $2^{n/2}$ for both schemes. Matching attacks on **LRW1** and **LRW2** are trivial as they both allow for an easy distinguisher after the first collision at the birthday bound. Other

constructions of tweakable block cipher related to LRW2 include XE and XEX by Rogaway [Rog04] and used in the OCB mode of operation.

In the quest for optimal security Mennink proposed the constructions $\tilde{F}[1]$ and $\tilde{F}[2]$ [Men15]. The latter reaches a provable security of 2^n queries which is the optimal security in the standard model for regular block ciphers. Other works tried to build a tweakable block cipher based solely on public permutations in the style of Even-Mansour [EM93]. Such tweakable block ciphers includes TEM [CLS15] and XPX [Men16] that are also subject to a tight birthday bound security of $\mathcal{O}(2^{n/2})$. Then Jha, List, Minematsu, Mishra and Nandi described a framework called XHX [JLM⁺17] and proved its security up to $2^{(n+\kappa)/2}$. They also describe generalised XHX, GXHX. In particular this means that a provable security beyond 2^n is reachable but in the ideal cipher model where rekeying is possible. This framework uses a single-round FX framework where all 3 subkeys are derived from a universal hash function on the secret master key and an arbitrarily long tweak.

So far, with the exception of GXHX, the proofs of all schemes cited can be shown to be tight. However, things become more involved when trying to iterate those constructions. Landecker et al. [LST12] proposed to iterate two independent evaluations of LRW2 and proved a security up to $2^{2n/3}$ queries. An attack on cascaded LRW2 (or CLRW2) has been later proposed by Mennink [Men18] in query complexity $\mathcal{O}(2^{3n/4})$ not completely closing the gap. Then, recently, Lee and Lee proposed XHX2 [LL18] by iterating two independent rounds of XHX. They managed to prove a query security lower bound of $\min\{2^{\frac{2}{3}(n+\kappa)}, 2^{n+\kappa/2}\}$ and left the tightness of this bound as an open question which we will be able to answer positively in this work.

On the other hand, a generic cryptanalysis of the r -round iterated FX construction has already been made with the original attack by Gaži [Gaž13] in query complexity $\mathcal{O}(2^{\frac{r-1}{r}n+\kappa})$. Obviously this attack can be used against our tweakable version when we fix the tweak to a single value. As it is written, the attack starts by querying all the code books of the secret cipher that makes the maximum possible 2^n calls. However this natural limitation of regular block ciphers has no ground in the presence of tweaks. Much like one can have security proofs beyond 2^n calls, one could attack a tweakable cipher using more than 2^n tweak/plaintext/ciphertext triples.

1.3 Results

Our generic iterated tweakable FX framework is pertinent to all cited constructions as shown in Table 1. Using a single-round FX to blend in the tweak is the most common approach and may be considered as well understood. However there seem to be additional security to be gained in iterating those constructions. Some works [LST12] [LL18] tend to do and prove just that. The focus on 2 rounds is justified by the fact that we don't know of any constructions based on block ciphers using more than 2 rounds and the single-round ones mostly have already well understood matching attacks. However we believe it is also

Table 1. Some previously proposed schemes and description of how it fits in our iterated tweakable FX generic framework.
 Multiplications ($\times$) are over the finite field $\text{GF}(2^n)$.

Ref	Scheme	r	Subkey functions
[LRW11]	LRW2	1	$\lambda_0(k, t) = \lambda_1(k, t)$ a uniform and AXU function. $\gamma_1(k, t) = k$
[Men15]	$\tilde{F}[1]$	1	$\lambda_0(k, t) = \lambda_1(k, t) = t \times k$ $\gamma_1(k, t) = t \oplus k$
[Men15]	$\tilde{F}[2]$	1	$\lambda_0(k, t) = \lambda_1(k, t) = E_1(2 \times k, t)$ $\gamma_1(k, t) = t \oplus k$
[Men16]	XPX	1	$\kappa = 0$ so $E_1(\cdot, m) = P(m)$ $t = t_{11} \parallel t_{12} \parallel t_{21} \parallel t_{22}$ $\lambda_0(k, t) = t_{11}k \oplus t_{12}P(k)$ $\lambda_1(k, t) = t_{21}k \oplus t_{22}P(k)$
[JLM ⁺ 17]	XHX	1	$\gamma_1(k, t)$ a uniform and AU function. $\lambda_0(k, t) = \lambda_1(k, t)$ a uniform and AXU function.
[LRW11]	LRW1	2	$\lambda_0(k, t) = \lambda_2(k, t) = 0$ $\lambda_1(k, t) = t$ $\gamma_1(k, t) = \gamma_2(k, t) = k$
[LST12]	CLRW2	2	$\lambda_0(k, t)$ and $\lambda_2(k, t)$ two uniform and AXU functions. $\lambda_1(k, t) = \lambda_0(k, t) \oplus \lambda_2(k, t)$ $\gamma_1(k, t) = \gamma_2(k, t) = k$
[LL18]	XHX2	2	$\gamma_1(k, t)$ and $\gamma_2(k, t)$ two uniform and AU functions. $\lambda_0(k, t)$ and $\lambda_2(k, t)$ two uniform and AXU functions. $\lambda_1(k, t) = \lambda_0(k, t) \oplus \lambda_2(k, t)$

interesting to know what kind of security bounds we might hope to achieve by iterating even further.

So in this paper we ask ourselves what is the best security bound attainable when using the iterated FX paradigm for building tweakable block ciphers from regular block ciphers. To do this we improve on the attack described by Gaži [Gaž13] to apply it in the tweakable block cipher setting.

First we show an information theoretic attack for $r = 2$ rounds when $\kappa \leq 2n$ with offline and online query complexity of:

$$Q = \mathcal{O}(2^{\frac{2}{3}(n+\kappa)}) \cdot \sqrt[3]{\tilde{\kappa}/n}.$$

Note that $Q = \mathcal{O}(2^{\frac{2}{3}(n+\kappa)})$ under the reasonable assumption that the size of the master secret key is linear with respect to the state size, that is, $\tilde{\kappa} = \mathcal{O}(n)$.

The recent construction **XHX2** by Lee and Lee [LL18] is a particular case of our setting where $\lambda_1(k, t) = \lambda_0(k, t) \oplus \lambda_2(k, t)$. Their provable security bound is $2^{\frac{2}{3}(n+\kappa)}$ whenever $\kappa \leq 2n$ and therefore matches our attack. Thus our results prove the tightness of their bound and their bound proves the optimality of the attack.

We then extend the attack to multiple rounds of the same construction. This gives an attack on r rounds when $\kappa \leq rn$ with query complexity:

$$Q = \mathcal{O}(2^{\frac{r}{r+1}(n+\kappa)}) \cdot \sqrt[r+1]{\tilde{\kappa}/n}.$$

Again note that $Q = \mathcal{O}(2^{\frac{r}{r+1}(n+\kappa)})$ under the assumption that $\tilde{\kappa} = \mathcal{O}(n)$.

Table 2. Some previously proposed schemes with their known asymptotic bounds.

Ref	Scheme	r	Proof	Known Attack	Our Generic Attack
[LRW11]	LRW2	1	$2^{n/2}$	$2^{n/2}$	$2^{\frac{1}{2}(n+\kappa)}$
[Men15]	$\tilde{F}$ [1]	1	$2^{\frac{2}{3}n}$	2^n	2^n (as $\kappa = n$)
[Men16]	XPX	1	$2^{n/2}$	$2^{n/2}$	$2^{n/2}$ (as $\kappa = 0$)
[JLM ⁺ 17]	XHX	1	$2^{\frac{1}{2}(n+\kappa)}$	$2^{\frac{1}{2}(n+\kappa)}$	$2^{\frac{1}{2}(n+\kappa)}$
[JLM ⁺ 17]	GXXH	1	$2^{\frac{1}{2}(n+\kappa)}$	none	$2^{\frac{1}{2}(n+\kappa)}$
[Men15]	$\tilde{F}$ [2]	1	2^n	2^n	N.A.
[LRW11]	LRW1	2	$2^{n/2}$	$2^{n/2}$	$2^{\frac{2}{3}(n+\kappa)}$
[LST12]	CLRW2	2	$2^{2n/3}$	$2^{3n/4}$	$2^{\frac{2}{3}(n+\kappa)}$
[LL18]	XHX2	2	$2^{\frac{2}{3}(n+\kappa)}$	none	$2^{\frac{2}{3}(n+\kappa)}$

2 Cryptanalysis of 2-Round Tweakable FX

In this section we give an algorithm to extract the master key of a 2-round tweakable FX construction, Algorithm 1, then we show how it works by deriving the constants used and thus deriving the total query complexity.

2.1 The Algorithm

This cryptanalysis of Algorithm 1 is a key recovery attack and follows the idea of the original cryptanalysis by Gaži [Gaž13]: we want just enough data to construct contradictory paths for each wrong key. First we do all the required offline computations under all possible κ -bit key. Input values are the sets S_1 and S_2 which can be chosen randomly and the input/output pairs under the key j are stored in $\mathcal{L}_{j,1}$ and $\mathcal{L}_{j,2}$ for E_1 and E_2 respectively. Then we store all observable tweak/plaintext/ciphertext triples in $\mathcal{L}_0$. We don't need to choose the set S_0 of inputs to the tweakable block cipher as the attack works in the known plaintext setting. At last we can test all the κ -bit keys; potential master keys k only using the stored values by reconstructing the paths round by round.

Indeed sets $\mathcal{A}$ and $\mathcal{B}$ reconstruct the paths under the current key guess and the condition $\forall (t, m, b) \in \mathcal{B} : (t, m, b \oplus \gamma_5(k, t)) \in \mathcal{L}_0$ is checking whether there is a contradictory path (if not satisfied) or not (if satisfied). The additional condition $|\mathcal{B}| \geq \nu$ is simply here to ensure a good reduction.

For completeness we provide Algorithm 2 to show how to construct the sets $\mathcal{A}$ and $\mathcal{B}$. To construct $\mathcal{A}$ is to apply Algorithm 2 with inputs $S_0, \mathcal{L}_{\gamma_1(k,t),1}, \lambda_0(k, t)$. It is basically looking over all elements of the first set and checking if a shifted version of a value exists somewhere in the second set then, if found, it records the starting and ending values.

The constants ν and Q are derived in Section 2.2 and the algorithm already ensures that the total query complexity is of magnitude Q . Indeed once we construct the sets $\mathcal{L}_{j,i}$ and $\mathcal{L}_0$ we will have all the necessary queries to perform

Algorithm 1 Cryptanalysis of 2-round tweakable FX construction.

Input: $\tilde{\kappa}, n, \kappa \leq 2n, \tilde{E}, E_1, E_2, \gamma_1, \gamma_2, \lambda_0, \lambda_1, \lambda_2$

Output: k : the master key of $\tilde{E}$

$\nu \leftarrow \tilde{\kappa}/n$

$Q \leftarrow 2^{\frac{2}{3}(n+\kappa)} \cdot \sqrt[3]{\nu}$

▷ Constants derived in Section 2.2

Randomly sample $S_1 \subset \{0, 1\}^n$ with $|S_1| = Q/2^\kappa = 2^{\frac{2n-\kappa}{3}} \sqrt[3]{\nu}$.

Randomly sample $S_2 \subset \{0, 1\}^n$ with $|S_2| = Q/2^\kappa = 2^{\frac{2n-\kappa}{3}} \sqrt[3]{\nu}$.

for all $j \in \{0, 1\}^\kappa$ **do**

$\mathcal{L}_{j,1} \leftarrow \{(m, E_1(j, m)) : m \in S_1\}$

$\mathcal{L}_{j,2} \leftarrow \{(m, E_2(j, m)) : m \in S_2\}$

▷ Offline Queries Sets

end for

Let $S_0 \subset \{0, 1\}^* \times \{0, 1\}^n$ with $|S_0| = Q$ be an observable tweak/message set.

$\mathcal{L}_0 \leftarrow \{(t, m, \tilde{E}(t, m)) : (t, m) \in S_0\}$

▷ Online Queries Set

for all $k \in \{0, 1\}^{\tilde{\kappa}}$ **do**

$\mathcal{A} \leftarrow \{(t, m, a) : (t, m) \in S_0, (m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}\}$

$\mathcal{B} \leftarrow \{(t, m, b) : (t, m, a) \in \mathcal{A}, (a \oplus \lambda_1(k, t), b) \in \mathcal{L}_{\gamma_2(k, t), 2}\}$

▷ by Algorithm 2

if $|\mathcal{B}| \geq \nu$ **and** $\forall (t, m, b) \in \mathcal{B} : (t, m, b \oplus \lambda_2(k, t)) \in \mathcal{L}_0$ **then**

return k

end if

end for

return $\emptyset$

▷ No proper key in the set

the attack. Since $|\mathcal{L}_{j,i}| = |S_i| = Q/2^\kappa$ and there are 2^κ different possible subkeys then the total number of queries to E_1 and E_2 is Q . Then we also construct $\mathcal{L}_0$ so the number of online queries will also be $|\mathcal{L}_0| = |S_0| = Q$.

2.2 Deriving The Constants

The Query Complexity. To derive the constant Q used in Algorithm 1 we first focus on what happens when we guess the correct master key k . In that case we want to make sure that $|\mathcal{B}| \geq \nu$ happens with good probability as the other constraint is always true by construction of the scheme.

First let's look at the set $\mathcal{A}$:

$$\mathcal{A} \leftarrow \{(t, m, a) : (t, m) \in S_0, (m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}\}$$

By construction there are Q values $(t, m) \in S_0$ and, as S_1 is chosen randomly and independently, there is a $|S_1|/2^n$ probability that $(m \oplus \lambda_0(k, t)) \in S_1$ for each (t, m) observed and thus that there exists an a such that $(m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}$. Therefore in expectation we have $|\mathcal{A}| = Q^2/2^{n+\kappa}$.

We do the same reasoning for $\mathcal{B}$:

$$\mathcal{B} \leftarrow \{(t, m, b) : (t, m, a) \in \mathcal{A}, (a \oplus \lambda_1(k, t), b) \in \mathcal{L}_{\gamma_2(k, t), 2}\}$$

Algorithm 2 Set construction.

Input: $\mathcal{S}_1, \mathcal{S}_2, \ell$

Output: $\mathcal{S}_3 \leftarrow \{(e, s_3) : (e, s_1) \in \mathcal{S}_1, (s_1 \oplus \ell, s_3) \in \mathcal{S}_2\}$
 $\mathcal{S}_3 \leftarrow \emptyset$
for all $(e, s_1) \in \mathcal{S}_1$ **do**
 if $\exists s_3 : (s_1 \oplus \ell, s_3) \in \mathcal{S}_2$ **then**
 $\mathcal{S}_3 \leftarrow \mathcal{S}_3 \cup \{(e, s_3)\}$
 end if
end for
return $\mathcal{S}_3$

to find that in expectation $|\mathcal{B}| = Q^3/2^{2n+2\kappa}$.

With some regularity assumptions, if $|\mathcal{B}| = \nu$ in expectation then $|\mathcal{B}| \geq \nu$ with constant probability. Therefore we put:

$$Q^3/2^{2n+2\kappa} = \nu \implies Q = 2^{\frac{2}{3}(n+\kappa)} \cdot \sqrt[3]{\nu}$$

The Number of Paths. The constant Q was derived so that we don't have false negatives, that is, we succeed with good probability when we guess the good key k . Now we derive the constant ν so that we don't have any false positive that means the test fails with good probability for all the wrong guesses of k .

First notice that the fact that $|\mathcal{B}| = \nu$ in expectation is true for all guesses of k , good or wrong. If $|\mathcal{B}| < \nu$ then the test fails as it should. If $|\mathcal{B}| \geq \nu$ then we need to look at the second condition that is $\forall (t, m, b) \in \mathcal{B} : (t, m, b \oplus \lambda_3(k, t)) \in \mathcal{L}_0$. If the guess is wrong then for a given $(t, m, b) \in \mathcal{B}$ we have $(b \oplus \lambda_3(k, t)) = \tilde{E}(t, m)$ with a 2^{-n} probability. Since $|\mathcal{B}| \geq \nu$ then the second condition is satisfied with probability $(2^{-n})^\nu = 2^{-\nu \cdot n}$. The test must fail for all the wrong guesses and there are $2^{\tilde{\kappa}} - 1$ such wrong guesses so all the tests should fail at least with constant probability when:

$$2^{\tilde{\kappa}} \cdot 2^{-\nu \cdot n} \leq 1 \implies \tilde{\kappa} - \nu \cdot n \leq 0 \implies \nu \geq \tilde{\kappa}/n$$

thus we take $\nu = \tilde{\kappa}/n$.

2.3 Constraints

For all of this to work there are some constraints that need to be spelled out. First we require that:

$$\begin{aligned} 1 &\leq |S_i| \\ \iff 1 &\leq 2^{\frac{2}{3}n - \frac{1}{3}\kappa} \cdot \sqrt[3]{\nu} \\ \iff \kappa &\leq 2n + \log(\nu) \end{aligned}$$

which limits to possible size of κ to a multiple of the state size n . Very few block ciphers admit a key larger than $2n$ so this is not a strong limitation in practice.

We also need to have diverse tweakable subkeys. Indeed so far we did not require that the functions $\gamma_i(k, t)$ depends on t which means that the tweak can be put, or not, at any stage of the construction but we still require that the tweak changes something. Therefore we can deduce such requirement:

$$\begin{aligned} \forall k \in \{0, 1\}^{\tilde{\kappa}} \forall (t, m) \in S_0 \forall (t', m') \in S_0 : \\ t \neq t' \implies \exists i : \gamma_i(k, t) \neq \gamma_i(k, t') \text{ OR } \lambda_i(k, t) \neq \lambda_i(k, t') \end{aligned}$$

which means that for every pairs of two different observed tweaks at least one of the respective implied subkeys must be different. This condition mostly ensure that this is a reasonable tweakable block cipher construction. Indeed in the case where two tweaks imply the exact same subkeys then one can quickly realise that it gets the same permutation for two different tweaks which is a near zero probability event for a perfect tweakable block cipher and hence it's a distinguisher.

3 Cryptanalysis of r -Round Tweakable FX

Starting from the attack of Section 2 we show how to generalise it to attack $r \geq 1$ rounds of the same construction in $Q = \mathcal{O}(2^{\frac{r}{r+1}(n+\kappa)} \cdot {}^{r+1}\sqrt{\tilde{\kappa}/n})$ query complexity. The strategy is the same, we begin by doing all the necessary queries before reconstructing paths round by round to finally check whether there is a contradictory path or not. This is Algorithm 3.




Fig. 2. r -Round Tweakable FX.

3.1 Constants and Complexity

The Query Complexity. We derive the constant Q used in Algorithm 3 in the same way as we did for the 2-round version. First we focus on what happens when we guess the correct master key k . In that case we want to make sure that $|\mathcal{B}| \geq \nu$ happens with good probability as contradictory paths cannot exist under the correct key.

Let's look at the set $\mathcal{A}_1$:

$$\mathcal{A}_1 \leftarrow \{(t, m, a) : (t, m) \in S_0, (m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}\}$$

Algorithm 3 Cryptanalysis of r -round tweakable FX construction.

Input: $\tilde{\kappa}, n, \kappa \leq rn, \tilde{E}, E_1, E_2, \dots, E_r, \gamma_1, \gamma_2, \dots, \gamma_r, \lambda_0, \lambda_1, \lambda_2, \dots, \lambda_r$

Output: k : the master key of $\tilde{E}$

```

1:  $\nu \leftarrow \tilde{\kappa}/n$ 
2:  $Q \leftarrow 2^{\frac{r}{r+1}(n+\kappa)} \cdot {}^{r+1}\sqrt{\nu}$ 

3: for all  $i \in \{1, \dots, r\}$  do
4:   Randomly sample  $S_i \subset \{0, 1\}^n$  with  $|S_i| = Q/2^\kappa = 2^{\frac{rn-\kappa}{r+1}} {}^{r+1}\sqrt{\nu}$ .
5: end for
6: for all  $j \in \{0, 1\}^\kappa$  do
7:   for all  $i \in \{1, \dots, r\}$  do
8:      $\mathcal{L}_{j,i} \leftarrow \{(m, E_i(j, m)) : m \in S_i\}$  ▷ Offline Queries Sets
9:   end for
10: end for

11: Let  $S_0 \subset \{0, 1\}^* \times \{0, 1\}^n$  with  $|S_0| = Q$  be an observable tweak/message set.
12:  $\mathcal{L}_0 \leftarrow \{(t, m, \tilde{E}(t, m)) : (t, m) \in S_0\}$  ▷ Online Queries Set

13: for all  $k \in \{0, 1\}^{\tilde{\kappa}}$  do
14:    $\mathcal{A}_1 \leftarrow \{(t, m, a) : (t, m) \in S_0, (m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}\}$ 
15:   for all  $i \in \{2, \dots, r\}$  do
16:      $\mathcal{A}_i \leftarrow \{(t, m, a) : (t, m, \bar{a}) \in \mathcal{A}_{i-1}, (\bar{a} \oplus \lambda_{i-1}(k, t), a) \in \mathcal{L}_{\gamma_i(k, t), i}\}$ 
17:   end for ▷ by Algorithm 2
18:   if  $|\mathcal{A}_r| \geq \nu$  and  $\forall (t, m, a) \in \mathcal{A}_r : (t, m, a \oplus \lambda_r(k, t)) \in \mathcal{L}_0$  then
19:     return  $k$ 
20:   end if
21: end for
22: return  $\emptyset$  ▷ No proper key in the set

```

By construction there are Q values $(t, m) \in S_0$ and, as S_1 is chosen randomly and independently, there is a $|S_1|/2^n$ probability that $\exists a : (m \oplus \lambda_0(k, t), a) \in \mathcal{L}_{\gamma_1(k, t), 1}$ for all observed tweak/message pairs (t, m) . Therefore, in expectation, we have $|\mathcal{A}_1| = Q^2/2^{n+\kappa}$.

Then we can easily prove by induction that $|\mathcal{A}_i| = Q^{i+1}/2^{i(n+\kappa)}$ as it is true for $|\mathcal{A}_1|$ and $|\mathcal{A}_{i+1}| = |\mathcal{A}_i| \cdot |S_{i+1}|/2^n$. Thus we get $|\mathcal{A}_r| = Q^{r+1}/2^{r(n+\kappa)}$.

With some regularity assumptions, if in expectation $|\mathcal{A}_r| = \nu$ then $|\mathcal{A}_r| \geq \nu$ with constant probability. Therefore we put:

$$Q^{r+1}/2^{r(n+\kappa)} = \nu \implies Q = 2^{\frac{r}{r+1}(n+\kappa)} \cdot {}^{r+1}\sqrt{\nu}$$

The Number of Paths. The constant Q was derived so that we avoid false negative when we guess the good key k . Now we derive the constant ν to avoid false positives.

If $|\mathcal{A}_r| < \nu$ then the test fails as it should. If $|\mathcal{A}_r| \geq \nu$ then the second condition is satisfied with probability $(2^{-n})^\nu = 2^{-\nu \cdot n}$. The test must fail for all the $2^{\tilde{\kappa}} - 1$ wrong guesses so all the tests should fail at least with constant

probability when:

$$2^{\tilde{\kappa}} \cdot 2^{-\nu \cdot n} \leq 1 \implies \tilde{\kappa} - \nu \cdot n \leq 0 \implies \nu \geq \tilde{\kappa}/n$$

thus we take $\nu = \tilde{\kappa}/n$.

For all of this to work there are again some constraints. First we require that:

$$\begin{aligned} 1 &\leq |S_i| \\ \iff \kappa &\leq rn + \log(\nu) \end{aligned}$$

which limits to possible size of κ to a multiple of the state size n .

Then we have the condition that the tweak changes something:

$$\begin{aligned} \forall k \in \{0, 1\}^{\tilde{\kappa}} \forall (t, m) \in S_0 \forall (t', m') \in S_0 : \\ t \neq t' \implies \exists i : \gamma_i(k, t) \neq \gamma_i(k, t') \text{ OR } \lambda_i(k, t) \neq \lambda_i(k, t') \end{aligned}$$

Notice that this condition prevents the known matching attack on **XXH**. Indeed, as for **XXH** $r = 1$ and $\lambda_0 = \lambda_1$, a collision on the full subkeys is expected after trying $\mathcal{O}(2^{(n+\kappa)/2})$ different tweaks. Our attack has the same complexity and also work on the generalised setting **GXXH** that doesn't enforce $\lambda_0 = \lambda_1$. This shows that the security cannot improve even if a collision on the full subkeys is made hard by, for example, choosing many different subkey functions or by using a mode of operation that limits the amount of different observable tweaks.

3.2 Discussion

Using Tweakable Block Ciphers. If instead of regular block ciphers we use tweakable block ciphers then it is not trivial to adapt this attack. Indeed we use the fact that the master key and the tweak must be blended before computation and not separately plugged in a tweakable block cipher. Such construction of a tweakable block cipher based on another tweakable block cipher could be used to increase security and/or the size of the tweak in a way that the original FX construction builds a stronger block cipher from another block cipher. However on the cryptanalysis side what can always be done is to fix a single tweak and apply the original attack by Gaži [Gaž13] in query complexity $\mathcal{O}(2^{\frac{r-1}{r}n+\kappa})$ or $\mathcal{O}(2^{\frac{r}{r+1}(n+\kappa)})$ when $\kappa \leq \frac{n}{r}$.

Weaker Constructions. This attack is generic given any reasonable key schedule represented by the λ and γ functions. However they are particular cases where better attacks are possible. In particular the cascaded **LRW2** construction is a 2-round tweakable FX construction where the key in the block cipher does not vary with the tweaks (γ_1 and γ_2 don't depend on t). This construction permits an attack in $\mathcal{O}(2^{\frac{3n}{4}})$ by Mennink [Men18] using only two different tweaks which beats our generic attack as soon as $\kappa > \frac{n}{8}$.

Tweak-rekeying. In fact our generic attack being a key recovery attack it will require at least 2^κ calls to the underlying block cipher. As soon as $k \geq n$ this implies a complexity above 2^n . Mennink [Men17] showed that provable 2^n security is unattainable in the standard block cipher model used for the proofs of schemes without tweak-rekeying. Therefore our generic attack can only hope to be tight for schemes that use tweak-rekeying and thus that are proved in the ideal block cipher model.

Key recovery and distinguisher. The fact that the complexity of this cryptanalysis depends on the size of the master key, even if a little, makes it hardly comparable to distinguishers that are independent of the master key size. Instead of waiting for some bad event to occur we collect just enough information to completely determine the master key. In the case of **XXH** the known distinguisher has the same asymptotic complexity but the widely different approaches make them hard to combine: a bad event for the known distinguisher gives no information on the master key. However for **XXH2**, and generally for $r \geq 2$ rounds of the tweakable FX construction proved in the ideal cipher model, it may well be the case that a key recovery approach is more relevant than looking for a suitable bad event for a distinguisher.

Towards Simplicity. The attack on generic 2-round tweakable FX is also tight since Lee and Lee could prove with **XXH2** [LL18] that we can reach this level of security even when $\lambda_1(k, t) = \lambda_0(k, t) \oplus \lambda_2(k, t)$ with some conditions on those functions. Moreover the previously known matching attack on **XXH** [JLM⁺17] exploited the fact that $\lambda_0(k, t) = \lambda_1(k, t)$ but our generic attack shows that it cannot be made more secure without this simplification. Another way to say it is that enforcing $\lambda_0(k, t) = \lambda_1(k, t)$ does not affect the provable security bound.

Using this iterated tweakable FX paradigm, one can therefore wonder how much it is possible to simplify the subkey functions while maintaining an optimal provable security with respect to the generic security upper bound shown in this work.

Acknowledgement

The author would like to thank the 2018 Asian Symmetric Key Workshop and Gaëtan Leurent for useful discussions. This work was partially supported by the French DGA.

References

- ADK⁺14. Martin R. Albrecht, Benedikt Driessen, Elif Bilge Kavun, Gregor Leander, Christof Paar, and Tolga Yalçın. Block ciphers - focus on the linear layer (feat. PRIDE). In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 57–76. Springer, Heidelberg, August 2014.

- AES01. Advanced Encryption Standard (AES). National Institute of Standards and Technology (NIST), FIPS PUB 197, U.S. Department of Commerce, November 2001.
- BCG⁺12. Julia Borghoff, Anne Canteaut, Tim Güneysu, Elif Bilge Kavun, Miroslav Knežević, Lars R. Knudsen, Gregor Leander, Ventzislav Nikov, Christof Paar, Christian Rechberger, Peter Rombouts, Søren S. Thomsen, and Tolga Yalçın. PRINCE - A low-latency block cipher for pervasive computing applications - extended abstract. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 208–225. Springer, Heidelberg, December 2012.
- BR02. John Black and Phillip Rogaway. A block-cipher mode of operation for parallelizable message authentication. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 384–397. Springer, Heidelberg, April / May 2002.
- CLS15. Benoit Cogliati, Rodolphe Lampe, and Yannick Seurin. Tweaking Even-Mansour ciphers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 189–208. Springer, Heidelberg, August 2015.
- DES77. Data encryption standard. National Bureau of Standards, NBS FIPS PUB 46, U.S. Department of Commerce, January 1977.
- EM93. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *ASIACRYPT’91*, volume 739 of *LNCS*, pages 210–224. Springer, Heidelberg, November 1993.
- Gaž13. Peter Gaži. Plain versus randomized cascading-based key-length extension for block ciphers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 551–570. Springer, Heidelberg, August 2013.
- JLM⁺17. Ashwin Jha, Eik List, Kazuhiko Minematsu, Sweta Mishra, and Mridul Nandi. XHX - A framework for optimally secure tweakable block ciphers from classical block ciphers and universal hashing. Cryptology ePrint Archive, Report 2017/1075, 2017. <https://eprint.iacr.org/2017/1075>.
- KR96. Joe Kilian and Phillip Rogaway. How to protect DES against exhaustive key search. In Neal Koblitz, editor, *CRYPTO’96*, volume 1109 of *LNCS*, pages 252–267. Springer, Heidelberg, August 1996.
- LL18. ByeongHak Lee and Jooyoung Lee. Tweakable block ciphers secure beyond the birthday bound in the ideal cipher model. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 305–335. Springer, Heidelberg, December 2018.
- LRW11. Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. *Journal of Cryptology*, 24(3):588–613, July 2011.
- LST12. Will Landecker, Thomas Shrimpton, and R. Seth Terashima. Tweakable blockciphers with beyond birthday-bound security. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 14–30. Springer, Heidelberg, August 2012.
- Men15. Bart Mennink. Optimally secure tweakable blockciphers. In Gregor Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 428–448. Springer, Heidelberg, March 2015.
- Men16. Bart Mennink. XPX: Generalized tweakable Even-Mansour with improved security guarantees. In Matthew Robshaw and Jonathan Katz, editors,

- CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 64–94. Springer, Heidelberg, August 2016.
- Men17. Bart Mennink. Insuperability of the standard versus ideal model gap for tweakable blockcipher security. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 708–732. Springer, Heidelberg, August 2017.
- Men18. Bart Mennink. Towards tight security of cascaded LRW2. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 192–222. Springer, Heidelberg, November 2018.
- MI15. Kazuhiko Minematsu and Tetsu Iwata. Tweak-length extension for tweakable blockciphers. In Jens Groth, editor, *15th IMA International Conference on Cryptography and Coding*, volume 9496 of *LNCS*, pages 77–93. Springer, Heidelberg, December 2015.
- RBBK01. Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In Michael K. Reiter and Pierangela Samarati, editors, *ACM CCS 2001*, pages 196–205. ACM Press, November 2001.
- Rog04. Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In Pil Joong Lee, editor, *ASIACRYPT 2004*, volume 3329 of *LNCS*, pages 16–31. Springer, Heidelberg, December 2004.