Trick or Tweak: On the (In)security of OTR's Tweaks

Raphael Bost[1,2(✉)] and Olivier Sanders[3]

[1] Direction Générale de l'Armement - Maîtrise de l'Information, Bruz, France
raphael_bost@alumni.brown.edu
[2] Université de Rennes 1, Rennes, France
[3] Orange Labs, Cesson-Sévigné, France

Abstract. Tweakable blockcipher (TBC) is a powerful tool to design authenticated encryption schemes as illustrated by Minematsu's Offset Two Rounds (OTR) construction. It considers an additional input, called tweak, to a standard blockcipher which adds some variability to this primitive. More specifically, each tweak is expected to define a different, independent pseudo-random permutation.

In this work we focus on OTR's way to instantiate a TBC and show that it does not achieve independence for a large amount of parameters. We indeed describe collisions between the input masks derived from the tweaks and explain how they result in practical attacks against this scheme, breaking privacy, authenticity, or both, using a single encryption query, with advantage at least 1/4.

We stress however that our results do not invalidate the OTR construction as a whole but simply prove that the TBC's input masks should be designed differently.

1 Introduction

Communications over an insecure channel usually rise the issue of confidentiality and authenticity of data exchanged through this channel. Although efficient solutions are known for each of these properties individually, their combination to ensure both is not obvious [BN00, Kra01] and has, in practice, resulted in security breaches (*e.g.* [Kra01, AP13]). Also, the combination of different constructions, potentially relying on different primitives, may reveal quite costly.

Designing an *authenticated encryption* (AE) scheme, which efficiently achieves both authenticity and confidentiality, has thus become a major topic in cryptography, with many past contributions [Dwo04, Dwo07, MV04, BRW04, Rog04, KR11]. Since the beginning of the CAESAR competition [CAE14], a large number of new constructions have been proposed, from blockcipher modes of operation [IMGM15, Min14, AFF+15, DN14, HKR15] to ad-hoc designs [Nik14], or sponge-based constructions [BDP+14, ABB+14]. Among the former, OTR [Min14] follows an approach based on tweakable blockciphers (TBC), a powerful primitive introduced by Liskov, Rivest and Wagner [LRW02].

J.H. Cheon and T. Takagi (Eds.): ASIACRYPT 2016, Part I, LNCS 10031, pp. 333–353, 2016.
DOI: 10.1007/978-3-662-53887-6_12

1.1 Tweakable Blockcipher

Compared to a regular blockcipher, a TBC $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ takes an additional input $T \in \mathcal{T}$, called a tweak, which adds some variability. As illustrated in [LRW02], a TBC enables simpler designs and security proofs for AE schemes, and can be instantiated from a blockcipher. To achieve efficiency, the design of the input masks must take into account the fact that the TBC is generally not used alone but rather in a mode of operation. In particular, the cost of changing the tweak must be much smaller than the cost of changing the key.

The now common constructions to build a TBC out of a block cipher are the Xor-Encrypt (XE) and Xor-Encrypt-Xor (XEX) constructions of [Rog04]. The principle of XE is to derive an input mask Δ from the tweak and xor it with the message before calling E_K (XEX also xors this mask to the output). The efficiency comes from designing the input mask Δ in such a way that Δ_{i+1} (used to encrypt the i-th message block) can be easily derived from Δ_i. For example, in OCB2 [Rog04], Δ_{i+1} is obtained from Δ_i by multiplying the latter by some elements of $\mathbb{F}_{2^n}$ (namely X or $(X+1)$, where X generates $\mathbb{F}_{2^n}^*$).

OTR's masks slightly differs from OCB2's one by using, among others, $\Delta_{i,0} = X^{i+1}\delta$ for the $2i-1$-th block and $\Delta_{i,1} = (X^{i+1}+1)\delta$ for the $2i$-th block (where δ is the encryption of the nonce). This approach is very well suited to the Feistel-based construction of OTR.

1.2 Our Contribution

However, we show in this paper that this solution is, at best, unsafe and even totally insecure in many cases. Indeed, the security of XE relies on the hardness of constructing collisions among the input masks Δ_i.

This can easily be proven for OCB2 due to the form of $\Delta = X^i(X+1)^j E_K(N)$. A collision in the offsets means that $X^i(1+X)^j = X^{i'}(1+X)^{j'}$ for some integers i, i', j and j', and so that $(1+X)^{j-j'} = X^{i'-i}$. This equation, along with the discrete logarithm of $X+1$ in base X, allows to define bounds on i and j excluding any collision. Unfortunately, this is no longer true for OTR due to the special form of its offsets. For example, if we just consider the input masks $\Delta_{i,0} = X^{i+1}\delta$ and $\Delta_{i,1} = (X^{i+1}+1)\delta$, it is impossible to formally exclude collisions: there are no algebraic reason why X^i should differ from $X^j + 1$ for any $i, j \leq B$, for some bound B.

The simple fact that no formal proof can be provided should itself call for another design of the masks, nevertheless one might still wonder if these collisions are likely.

In this work, we investigate this issue and show that, for a large family of blocksize $n \leq 10000$ (OTR is defined for any blockcipher size $n \in \mathbb{N}^*$), standard choices of parameters lead to trivial collisions. Moreover, we show that the block sizes outside this family are not necessarily secure and need a specific, costly study to exclude collision for reasonable B. We focus on the most popular choices, namely $n = 64$ and $n = 128$, and present a collision for the former case when

$\mathbb{F}_{2^{64}}$ is generated, as usual, using the primitive pentanomial $P = X^{64} + X^4 + X^3 + X + 1$. We get similar results for $n = 128$ when $\mathbb{F}_{2^{128}}$ is generated by some specific primitive pentanomials. However, the latter do not include the usually used one, namely $P = X^{128} + X^7 + X^2 + X + 1$. We therefore study more thoroughly this case and propose a bound $B = 2^{45}$ excluding collisions. We do not claim that this bound is optimal but we provide evidence that collisions are likely to occur between 2^{45} and 2^{64}.

In a second part, we describe concrete attacks against privacy and authenticity resulting from these collisions. They show that the latter do not simply invalidate the security proof but also completely break the security of the construction.

Finally, we describe some ways of constructing the input masks which prevent collisions. We therefore emphasize that our work does not question the intrinsic security of OTR seen as a TBC mode of operation, but simply shows that the instantiation of the TBC in [Min14] should be fixed. In particular, due to our attack, Minematsu modified the masks generation in the last version of the CAESAR submission, AES-OTRv3 [Min16].

2 Preliminaries

2.1 Basic Notations

For sake of clarity, we will use the same notations as the ones of [Min14]. The set of all finite-length binary strings, including the empty string ϵ, is denoted by $\{0,1\}^*$. $\forall S \in \{0,1\}^*$, $|S|$ denotes the length of S and $|S|_a = \mathtt{max}\{\lceil(|S|/a)\rceil, 1\}$. The concatenation of two binary strings S and T is written ST. $\forall S \in \{0,1\}^*, (S[1], \ldots, S[m]) \xleftarrow{n} S$ denotes the n-bit block partitioning of S, *i.e.* $S = S[1] \ldots S[m]$, where $|S[i]| = n$ for $i < m$ and $|S[m]| \leq n$ (we thus have $m = |S|_n$). The sequence of a zeros is denoted by 0^a. For all $n \in \mathbb{N}$ and S such that $|S| \leq n$, $\underline{S}_n$ denotes the padding $S10^{n-|S|-1}$ if $|S| < n$ and S otherwise. In the following, we will omit the subscript n if it is made obvious by the context. For a finite set $\mathcal{S}$, we write $S \xleftarrow{\$} \mathcal{S}$ if S is uniformly chosen from $\mathcal{S}$.

2.2 Blockciphers and Tweakable Blockciphers

We review the standard definitions of blockciphers and tweakable blockciphers from [LRW02,Rog04]. A blockcipher is a function $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ where $n \in \mathbb{N}$, $\mathcal{K} \neq \emptyset$ is a finite set and $E(K,.) = E_K(.)$ is a permutation for each $K \in \mathcal{K}$. The PRF and PRP advantages of E against adversary $\mathcal{A}$ are defined as:

$$\mathtt{Adv}^{\mathtt{prf}}_E(\mathcal{A}) = \mathbb{P}[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{E_K(.)} \Rightarrow 1] - \mathbb{P}[\rho \xleftarrow{\$} \mathrm{Func}(n) : \mathcal{A}^{\rho(.)} \Rightarrow 1]$$

$$\mathtt{Adv}^{\mathtt{prp}}_E(\mathcal{A}) = \mathbb{P}[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{E_K(.)} \Rightarrow 1] - \mathbb{P}[\pi \xleftarrow{\$} \mathrm{Perm}(n) : \mathcal{A}^{\pi(.)} \Rightarrow 1]$$

where $\mathrm{Func}(n)$ (resp. $\mathrm{Perm}(n)$) is the set of all the functions (resp. permutations) $\{0,1\}^n \to \{0,1\}^n$.

A tweakable blockcipher is a blockcipher with an additional public input. It is formalized as a function $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ where $n \in \mathbb{N}$, $\mathcal{K}, \mathcal{T} \neq \emptyset$ are finite sets and $\widetilde{E}(K,T,.) = \widetilde{E}_K(T,.) = \widetilde{E}_K^T(.)$ is a permutation for each $K \in \mathcal{K}$ and $T \in \mathcal{T}$. The tweakable PRF and tweakable PRP advantages of $\widetilde{E}$ against adversary $\mathcal{A}$ is defined as:

$$\mathtt{Adv}^{\widetilde{\mathtt{prf}}}_{\widetilde{E}}(\mathcal{A}) = \mathbb{P}[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\widetilde{E}_K(.,.)} \Rightarrow 1] - \mathbb{P}[\widetilde{\rho} \xleftarrow{\$} \mathrm{Func}(\mathcal{T},n) : \mathcal{A}^{\widetilde{\rho}(.,.)} \Rightarrow 1]$$
$$\mathtt{Adv}^{\widetilde{\mathtt{prp}}}_{\widetilde{E}}(\mathcal{A}) = \mathbb{P}[K \xleftarrow{\$} \mathcal{K} : \mathcal{A}^{\widetilde{E}_K(.,.)} \Rightarrow 1] - \mathbb{P}[\widetilde{\pi} \xleftarrow{\$} \mathrm{Perm}(\mathcal{T},n) : \mathcal{A}^{\widetilde{\pi}(.,.)} \Rightarrow 1]$$

where $\mathrm{Func}(\mathcal{T},n)$ (resp. $\mathrm{Perm}(\mathcal{T},n)$) is the set of all mappings from $\mathcal{T}$ to functions (resp permutations) $\{0,1\}^n \to \{0,1\}^n$.

2.3 Authenticated Encryption

Definition. An authenticated encryption AE[τ] having a τ-bit tag consists of an encryption algorithm AE-$\mathcal{E}_\tau$ and a decryption algorithm AE-$\mathcal{D}_\tau$. The former takes as input a key $K \in \mathcal{K}_{ae}$, a nonce $N \in \mathcal{N}_{ae}$ and an associated data $A \in \mathcal{A}_{ae}$ along with a message $M \in \mathcal{M}_{ae}$ and outputs a ciphertext $C \in \mathcal{M}_{ae}$ as well as a tag $T_E \in \{0,1\}^\tau$. On input (K,N,A,C,T_E), the latter outputs a plaintext M such that $|M| = |C|$ or an error symbol $\perp$. The sets $\mathcal{K}_{ae}$, $\mathcal{N}_{ae}$, $\mathcal{A}_{ae}$ and $\mathcal{M}_{ae}$ are assumed to be non-empty and finite.

Security Model. The security properties expected from an authenticated encryption scheme are privacy and authenticity. The former informally requires that no adversary, even given access to encryption queries, is able to distinguish AE[τ] from an oracle \$ returning a random pair $(C,T_E) \xleftarrow{\$} \{0,1\}^{|M|} \times \{0,1\}^\tau$ on input (N,A,M). This is formally defined by the following advantage:

$$\mathtt{Adv}^{\mathtt{priv}}_{\mathrm{AE}[\tau]}(\mathcal{A}) = \mathtt{Pr}[K \xleftarrow{\$} \mathcal{K}_{ae} : \mathcal{A}^{\mathrm{AE}-\mathcal{E}_\tau} \to 1] - \mathtt{Pr}[\mathcal{A}^{\$} \to 1].$$

We say an adversary $\mathcal{A}$ is *nonce-respecting* if it cannot submit two queries (N_i,A_i,M_i) and (N_j,A_j,M_j) with $N_i = N_j$ for $i \neq j$. In this paper, we will always consider nonce-respecting adversaries. It is claimed in [Min14] that $\mathtt{Adv}^{\mathtt{priv}}_{\mathrm{OTR}[\tau]}(\mathcal{A}) \leq \frac{6(q+\sigma_A+\sigma_M)^2}{2^n}$ where q is the number of encryption queries and $(\sigma_A,\sigma_M) = (\sum_i^q |A_i|, \sum_i^q |M_i|)$.

Authenticity informally requires that no adversary, even with access to encryption and decryption queries, is able to produce a valid tuple (N,A,C,T_E), *i.e.* one such that AE-$\mathcal{D}_\tau(N,A,C,T_E) \neq \perp$. Obviously, (N,A,C,T_E) must not have been previously returned by the encryption oracle. The authenticity notion is defined by the advantage:

$$\mathtt{Adv}^{\mathtt{auth}}_{\mathrm{AE}[\tau]}(\mathcal{A}) = \mathtt{Pr}[K \xleftarrow{\$} \mathcal{K}_{ae} : \mathcal{A}^{\mathrm{AE}-\mathcal{E}_\tau,\mathrm{AE}-\mathcal{D}_\tau} \text{ forges}]$$

where $\mathcal{A}$ forges if one of the decryption query $(N'_i,A'_i,C'_i,T'_{E,i})$ does not return $\perp$. Notice that N'_i may be equal to N_j or $N'_{i'}$ for all i,i' and j. It is claimed in [Min14]

that $\mathtt{Adv}^{\mathtt{auth}}_{\mathrm{OTR}[\tau]}(\mathcal{A}) \leq \frac{6(q+q'+\sigma_A+\sigma_M+\sigma_{A'}+\sigma_{C'})^2}{2^n}$ where q (resp. q') is the number of encryption (resp. decryption) queries, $(\sigma_A, \sigma_M) = (\sum_i^q |A_i|, \sum_i^q |M_i|)$ and $(\sigma_{A'}, \sigma_{C'}) = (\sum_i^q |A'_i|, \sum_i^q |C'_i|)$.

2.4 Galois Field

For all non negative integers n, we denote by $\mathbb{F}_{2^n}$ the field with 2^n elements and by $\mathbb{F}^*_{2^n}$ its multiplicative group. To represent this field one [IK03, Rog04, Min14] usually selects the lexicographically first polynomial P among the primitive polynomials of degree n with coefficients in $\mathbb{F}_2$ having a minimum number of non-zero coefficients, and use $\mathbb{F}_2[X]/P(X)$ as a representation of $\mathbb{F}_{2^n}$. [Ser98] provides such polynomials for $n \leq 10000$. An element $a \in \mathbb{F}_{2^n}$ can then be written as a formal polynomial $b_1X^{n-1} + \ldots + b_{n-1}X + b_n$ of degree $n-1$ or equivalently as a n-bit string $b_1 \ldots b_n$. In the following, we will use both notations interchangeably.

For any $a = b_1X^{n-1} + \ldots + b_n$ and $c = b'_1X^{n-1} + \ldots + b'_n$ in $\mathbb{F}_{2^n}$, the product $a \cdot c$ is $(\sum_{i=1}^n b_iX^{n-i})(\sum_{j=1}^n b'_jX^{n-j})$ mod $P(X)$. In particular, it is worthy to note that $a \cdot X$ can be computed very efficiently with a shift and a conditional xor, hence the interest of a low-weight polynomial P. For example, for $n = 119$, one would select $P(X) = X^{119} + X^8 + 1$ [Ser98], so $a \cdot X = (a << 1) \oplus 0^{110}b_1 0^7 b_1$.

The table in [Ser98] shows that, up to $n = 10000$, primitive trinomials exist for slightly over one half of the values of n. In this case, the field $\mathbb{F}_{2^n}$ is usually generated by $X^n + X^j + 1$ for some $j \in [1, n-1]$. Otherwise, the table shows that, for $n \leq 10000$, one can at least find an irreducible pentanomial. For example, for $n = 128$, one can use $P(X) = X^{128} + X^7 + X^2 + X + 1$.

3 Description of OTR

Before describing our attack, we recall the AE scheme of [Min14], OTR$[E, \tau]$, parametrized by a keyed permutation $E_K : \{0,1\}^n \to \{0,1\}^n$, and a tag length $\tau \leq n$. Its encryption algorithm OTR-$\mathcal{E}_{E,\tau}$ consists of an encryption core EF_E and an authentication core AF_E which processes the additional authenticated data. Since our attack applies on EF_E, we omit the description of AF_E in Fig. 1 and assume that the string A (authenticated data) is empty.

EF_E can be seen as a variation of the tweakable blockcipher based authenticated encryption mode OCB [Rog04]. In OTR, tweakable blockciphers are instantiated using a two-rounds Feistel permutation where internal round functions are PRFs with tweak-dependent input masks. Algorithm 1 gives a formal description of the authenticated encryption algorithm $\mathbb{EF}[\widetilde{\rho}, \tau]$ that uses a tweakable random function $\widetilde{\rho}$. As defined in [Min14], the tweak space of $\widetilde{\rho}$ is $\mathcal{T} = (\{0,1\}^n \times \mathbb{N} \times \{0,1\}) \cup (\{*\} \times \{0,1\}^n \times \mathbb{N} \times \{0,1\} \times \{0,1\})$.[1]

An important theorem in the security proof of OTR is that, if $\widetilde{\rho}$ is a tweakable random function, then $\mathbb{EF}[\widetilde{\rho}, \tau]$ is a secure authenticated encryption scheme.

[1] We slightly changed the notations from [Min14] to give a more formal construction of the tweakable PRF.

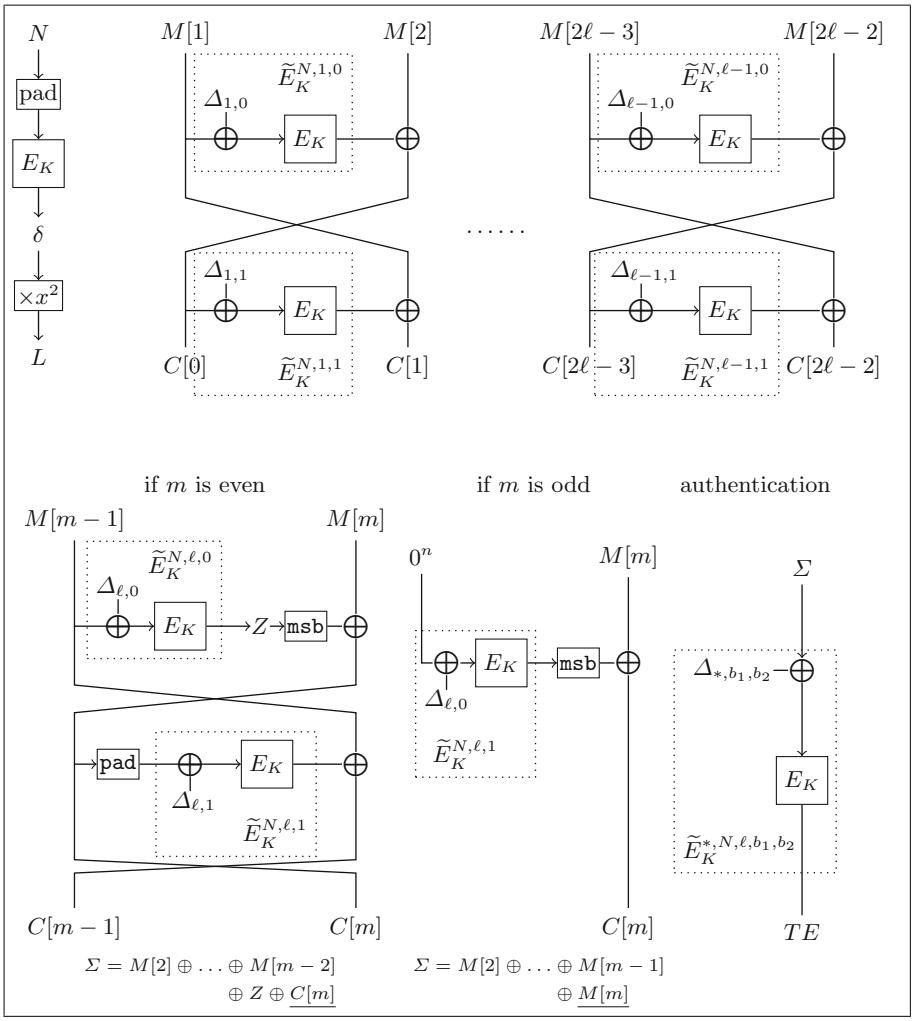

Fig. 1. Encryption core EF_E of OTR for a message $M = M[1] \ldots M[m]$ and a blocksize n. The integer ℓ is defined as $\lceil \frac{m}{2} \rceil$. $\Delta_{i,b} = (X^{i+1} + b)\delta$, for $i = 1, \ldots, \ell$ and $b \in \{0, 1\}$. $\Delta_{*,b_1,b_2} = [(X+1)X^{\ell+1} + X \cdot b_1 + b_1 + b_2]\delta$ with $b_1 = 0$ if m is odd and 1 otherwise while $b_2 = 0$ if $|M[m]| < n$ and 1 otherwise. The dotted boxes represent the tweakable random functions of the $\mathbb{OTR}$ construction.

Theorem 1 (Theorem 3 of [Min14]). *Fix $\tau \in \{1, \ldots, n\}$. For any adversary $\mathcal{A}$, and tweakable random function $\widetilde{\rho}$*

$$\mathtt{Adv}^{\mathtt{priv}}_{\mathbb{EF}[\widetilde{\rho},\tau]}(\mathcal{A}) = 0.$$

Moreover, for any adversary $\mathcal{A}$ making q encryption queries and q_v decryption queries,

$$\mathtt{Adv}^{\mathtt{auth}}_{\mathbb{EF}[\widetilde{\rho},\tau]}(\mathcal{A}) \leq \frac{2q_v}{2^n} + \frac{q_v}{2^\tau}.$$

We refer to the original paper for the full proof of this theorem. Minematsu also instantiates $\widetilde{\rho}$ using the XE approach [Rog04]:

$$\widetilde{E}^{N,i,a}_K(P) = E_K(P + \Delta_{i,a}) \text{ with } \Delta_{i,a} = X^{i-1}L + a \cdot \delta$$
$$\widetilde{E}^{*,N,i,b_1,b_2}_K(P) = E_K(P + \Delta_{*,i,b_1,b_2}) \text{ with } \Delta_{*,i,b_1,b_2} = (X+1)(X^{i-1}L + b_1 \cdot \delta) + b_2 \cdot \delta$$

where $\delta = E_K(N)$ and $L = X^2\delta$. Once developed, the final expression of the Δ values is

$$\Delta_{i,a} = (X^{i+1} + a)\delta$$
$$\Delta_{*,i,b_1,b_2} = (X^{i+2} + X^{i+1} + b_1X + b_1 + b_2)\delta.$$

Algorithm 1. Description of $\mathbb{EF}[\widetilde{\rho}, \tau]$.

1: $\Sigma \leftarrow 0^n$
2: $(M[1], \ldots, M[m]) \xleftarrow{n} M$
3: $\ell \leftarrow \lceil m/2 \rceil$
4: **for** $i = 1$ **to** $\ell - 1$ **do**
5: $\quad C[2i-1] \leftarrow \widetilde{\rho}^{N,i,0}(M[2i-1]) \oplus M[2i]$
6: $\quad C[2i] \leftarrow \widetilde{\rho}^{N,i,1}(C[2i-1]) \oplus M[2i-1]$
7: $\quad \Sigma \leftarrow \Sigma \oplus M[2i]$
8: **end for**
9: **if** m is even **then**
10: $\quad Z \leftarrow \widetilde{\rho}^{N,\ell,0}(M[m-1])$
11: $\quad C[m] \leftarrow \mathtt{msb}_{|M[m]|}(Z) \oplus M[m]$
12: $\quad C[m-1] \leftarrow \widetilde{\rho}^{N,\ell,1}(\underline{C[m]}) \oplus M[m-1]$
13: $\quad \Sigma \leftarrow \Sigma \oplus Z \oplus \underline{C[m]}$
14: $\quad$ **if** $|M[m]| \neq n$ **then** $TE \leftarrow \widetilde{\rho}^{*,N,\ell,1,0}(\Sigma)$
15: $\quad$ **else** $TE \leftarrow \widetilde{\rho}^{*,N,\ell,1,1}(\Sigma)$
16: **else** ▷ *m is odd*
17: $\quad C[m] \leftarrow \mathtt{msb}_{|M[m]|}(\widetilde{\rho}^{N,\ell,0}(0^n)) \oplus M[m]$
18: $\quad \Sigma \leftarrow \Sigma \oplus \underline{M[m]}$
19: $\quad$ **if** $|\underline{M[m]}| \neq n$ **then** $TE \leftarrow \widetilde{\rho}^{*,N,\ell,0,0}(\Sigma)$
20: $\quad$ **else** $TE \leftarrow \widetilde{\rho}^{*,N,\ell,0,1}(\Sigma)$
21: **end if**
22: $C \leftarrow (C[1], \ldots, C[m])$
23: **return** (C, TE)

To finish the proof of security, [Min14] uses the Lemma 1, claiming the CPA security of the tweakable PRF $\widetilde{E}$, provided that E is a perfect blockcipher (a random permutation):

Lemma 2 (Lemma 1 of [Min14]). *For any adversary $\mathcal{A}$ making q queries,*

$$\mathtt{Adv}^{\widetilde{\mathtt{prf}}}_{\widetilde{E}}(\mathcal{A}) \leq \frac{5q^2}{2^n}.$$

The proof of Lemma 1 relies on the fact that the masks Δ are assumed to be "differentially uniform" for any two distinct inputs. However, we show below that this is not the case for a large choice of parameters n, and that it actually completely breaks the security of OTR.

4 Collision in Masks Polynomials

4.1 Flaw in OTR's Proof

In [Min14], all possible masks Δ are regrouped in a set

$$\mathcal{S}_1(\delta) = \big\{X^{i+1}\delta, (X^{i+1}+1)\delta, (X^{i+2}+X^{i+1})\delta, (X^{i+2}+X^{i+1}+X)\delta, \\ (X^{i+2}+X^{i+1}+1)\delta, (X^{i+2}+X^{i+1}+X+1)\delta\big\}_{i=1}$$

(no upper bound on i is given but we can suppose that it is bounded by the maximum number of blocks one can query for an encryption, and that is it at most $2^{n/2}$) and it is claimed that for any $\Delta, \Delta' \in \mathcal{S}_1(\delta_1) \cup \mathcal{S}_1(\delta_2)$ such that Δ and Δ' are generated from two different expressions, and $d \in \{0,1\}^n$,

$$\Pr_{\delta_1,\delta_2 \xleftarrow{\$} \{0,1\}^n}[\Delta + \Delta' = d] \leq \frac{1}{2^n}$$

where the probability is taken over the random choices of δ_1 and δ_2. This is true if $\Delta \in \mathcal{S}_1(\delta_1)$ and $\Delta' \in \mathcal{S}_1(\delta_2)$, but not if both Δ and Δ' are generated from the same δ.

Namely, suppose that there are two integers i and $j \geq 2$ such that

$$X^i = X^j + 1 \tag{1}$$
$$\text{or } X^i = X^{j+1} + X^j + r(X) \tag{2}$$
$$\text{or } X^{i+1} + X^i = X^{j+1} + X^j + r(X) \tag{3}$$

with $r(X) \in \{0, 1, X, X+1\}$. Then we directly have a collision inside $\mathcal{S}_1(\delta)$ for any δ. This problem is not highlighted in the proof and we will show that we can actually find (and use) such pairs of integers.

In the following, we will use the terms 'type-1', 'type-2', and 'type-3' for collisions satisfying, respectively, Eqs. (1), (2) and (3).

4.2 Finding Collisions

The problem with the polynomials considered above is that it seems impossible, given $n \in \mathbb{N}$ and a polynomial P generating $\mathbb{F}_{2^n}$, to provide a formal argument excluding collisions for any $i, j \in [2, t]$ for some integer $2 < t \leq 2^{n/2}$. One can note that we do not consider collisions in the set $\{X^i\}_{i=2}^t$, as X is a generator of $\mathbb{F}_{2^n}^*$ (since P is primitive) and we chose $t \leq 2^{n/2}$.

Actually, we show that trivial collisions can be found when the definition polynomial P has a special form, in particular when P is a trinomial or a pentanomial.

Case 1: $\mathbb{F}_{2^n}$ is generated by a trinomial $P(X) = X^n + X^j + 1$.

As explained in [Ser98], this is the standard choice for a majority of values $n \leq 10000$. In such a case, a collision in $\mathcal{S}_1$ is trivially given by P since $X^n = X^j + 1$ (this is thus a type-1 collision). Any encryption of a message M of m blocks such that $\lceil \frac{m}{2} \rceil \geq n - 1$ will then lead to the re-use of a mask and so to one of the attacks described in the next session.

One might argue that this can be avoided by generating $\mathbb{F}_{2^n}$ with a pentanomial instead of a trinomial. However, this unconventional choice will negatively impact the performances of the scheme and will not necessarily prevent collisions.

Case 2: $\mathbb{F}_{2^n}$ is generated by a pentanomial $P(X) = X^n + X^{j_1} + X^{j_2} + X^{j_3} + 1$. This case includes, for example, $n = 64$ and $n = 128$. Although there is no trivial collision as in the previous case, it is still necessary to check, for the chosen n and P, that $\mathcal{S}_1$ only contains distinct elements, which requires a significant amount of computations and storage space. We here describe the most popular cases:

- $n = 64$. The lexicographically first primitive pentanomial of degree 64 is $X^{64} + X^4 + X^3 + X + 1$ [Ser98]. It leads to a type-2 collision since $X^{64} = X^4 + X^3 + X + 1$.
- $n = 128$. Here again, the pentanomial generating $\mathbb{F}_{2^{128}}$ may give an obvious collision. For example, setting $P = X^{128} + X^{68} + X^{67} + X + 1$ leads to a type-2 collision $X^{128} = X^{68} + X^{67} + X + 1$, and setting $P = X^{128} + X^{127} + X^{61} + X^{60} + 1$ leads to a type-3 collision $X^{128} + X^{127} = X^{61} + X^{60} + 1$. However, this is not the case with the lexicographically first primitive pentanomial of degree 128, $P = X^{128} + X^7 + X^2 + X + 1$, that one generally uses to define $\mathbb{F}_{2^{128}}$. The latter therefore needs a more thorough study that we defer to Sect. 6.

5 Practical Attacks

One may wonder if the collisions found in the input masks simply invalidate the security proofs of OTR. Unfortunately, this is not the case and we show below that any kind of collision leads to attacks breaking privacy and/or authenticity. We recall that, for sake of simplicity, authenticated data are assumed to be empty in the following attacks. Attacks for non-empty authenticated data can easily be derived from them.

5.1 Type-1 Collisions

A type-1 collision occurs when there are i and j such that $X^i = X^j + 1$. We can assume, without loss of generality, that $j < i$ (since $X^i = X^j + 1 \Leftrightarrow X^j = X^i + 1$).

Breaking Authenticity. To break authenticity, one can make a query on an arbitrary message $M = M[1] \dots M[2i-3]$ for a nonce N, defining $\delta = E_K(N)$ and $L = X^2\delta$, and receive the ciphertext $C = C[1] \dots C[2i-3]$ along with the tag $T = TE$.

The message M has an odd number of blocks so $C[2i-3] = E_K(X^i\delta) \oplus M[2i-3]$.

Let $C' = C'[1] \dots C'[2i-3]$ such that $C'[k] = C[k]$ for $k \notin \{2j-3, 2j-2, 2i-3\}$, $C'[2j-3] = 0^n$, $C'[2j-2] = M[2j-3] \oplus C[2i-3] \oplus M[2i-3]$ and $C'[2i-3] = C[2i-3] \oplus C[2j-3]$.

Then, the pair (C', TE) is valid: OTR-$\mathcal{D}_{E,\tau}(N, \epsilon, C', T) = M'[1] \dots M'[2i-3] \neq \perp$. Indeed, by construction, we have $M'[k] = M[k]\ \forall k \notin \{2j-3, 2j-2, 2i-3\}$. Moreover, we have

$$\begin{aligned} M'[2j-3] &= E_K(C'[2j-3] \oplus (X^j+1)\delta) \oplus C'[2j-2] \\ &= E_K(0^n \oplus (X^j+1)\delta) \oplus M[2j-3] \oplus C[2i-3] \oplus M[2i-3] \\ &= E_K((X^j+1)\delta) \oplus M[2j-3] \oplus E_K(X^i\delta) \\ &= M[2j-3] \end{aligned}$$

and

$$\begin{aligned} M'[2j-2] &= E_K(M'[2j-3] \oplus X^j\delta) \oplus C'[2j-3] \\ &= E_K(M[2j-3] \oplus X^j\delta) \oplus 0^n \\ &= C[2j-3] \oplus M[2j-2]. \end{aligned}$$

Finally, we have $M'[2i-3] = M[2i-3] \oplus C[2j-3]$. Therefore:

$$\Sigma' = \Sigma \oplus C[2j-3] \oplus C[2j-3] = \Sigma$$

and the tag TE remains valid for C'.

For an adversary $\mathcal{A}$ following this procedure,

$$\mathtt{Adv}^{\mathtt{auth}}_{\mathrm{AE}[\tau]}(\mathcal{A}) = 1.$$

Breaking Privacy. We describe here a way that an adversary $\mathcal{A}$ can use to break privacy with advantage almost $1/4$ with a single query. To break privacy, $\mathcal{A}$ queries the encryption oracle with a random nonce N and a message $M = M[1] \dots M[2i-2]$ such that $|M[2i-2]| = 1$ and $M[2j-3] = 010^{n-2}$. $\mathcal{A}$ will receive $C = C[1] \dots C[2i-2]$ with $|C[2i-2]| = 1$. If $C[2i-2] = 1$ (which happens with probability $\frac{1}{2}$), $\mathcal{A}$ just picks its output bit at random (she does not try further up). Otherwise, we have $\underline{C[2i-2]} = 010^{n-2} = M[2j-3]$.

As a consequence, we get the following:

$$\begin{aligned} M[2i-3] &= E_K(\underline{C[2i-2]} \oplus (X^i+1)\delta) \oplus C[2i-3] \\ &= E_K(M[2j-3] \oplus X^j\delta) \oplus C[2i-3] \\ &= C[2j-3] \oplus M[2j-2] \oplus C[2i-3] \end{aligned}$$

and $M[2j-2] \oplus M[2i-3] = C[2j-3] \oplus C[2i-3]$, which defines an efficient distinguisher between the random encryption oracle and the real encryption oracle. More formally,

$$\mathtt{Adv}^{\mathtt{priv}}_{\mathrm{AE}[\tau]}(\mathcal{A}) = \frac{1}{2}\left(1 - \frac{1}{2^n}\right) - \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4} - \frac{1}{2^{n+1}}.$$

5.2 Type-2 Collisions

A type-2 collision occurs when there are i and j such that $X^i = X^{j+1} + X^j + r(X)$ with $r(X) \in \{0, 1, X, X+1\}$. We show below how one can break authenticity if $i \geq j$ and privacy if $i < j$.

Breaking Privacy for $i < j$. To break privacy, one submits a message $M = M[1] \dots M[m] = 0^n \dots 0^n M[2i-3] M[2i-2] 0^n \dots M[m-1] 0^{|M[m]|}$ where m, $|M[m]|$, $M[2i-3]$,$M[2i-2]$ and $M[m-1]$ are defined as follows:

- If $r(X) = X+1$, then one sets $m = 2(j-1)$, $|M[m]| = n-1$, $M[2i-3] = M[2i-2] \in \{0,1\}^n$ and $M[m-1] \in \{0,1\}^n$.
  Since the last block of M is 0^{n-1}, the $n-1$ most significant bits of $Z \oplus C[m]$ are 0^{n-1}. Therefore, if the last bit of Z is 1 (which occurs with probability $\frac{1}{2}$), $Z \oplus \underline{C[m]} = 0^n$. Also, in this case, $\Sigma = M[2i-2] = M[2i-3]$. If the last bit of Z is not 1, one simply submits new messages with different $M[m-1]$ until this condition is fulfilled.
  The authentication tag TE then verifies the following relation:

$$\begin{aligned} TE &= E_K(\Sigma \oplus \Delta_{*,m,1,0}) \\ &= E_K(M[2i-3] \oplus (X^{j+1} + X^j + X + 1)\delta) \\ &= E_K(M[2i-3] \oplus X^i \delta) \\ &= C[2i-3] \oplus M[2i-2] \end{aligned}$$

  Therefore, $TE \oplus C[2i-3] = M[2i-2]$, which breaks privacy.
- If $r(X) = X$, then one sets $m = 2(j-1)$, $|M[m]| = n$, $M[2i-3] = M[2i-2] \in \{0,1\}^n$ and $M[m-1] \in \{0,1\}^n$. In such a case, $\Sigma = M[2i-2] = M[2i-3]$ and the previous attack still applies.
- If $r(X) = 1$, then one sets $m = 2(j-1)-1$, $|M[m]| = n$, $M[2i-3] = M[2i-2] \in \{0,1\}^n$ and $M[m-1] = 0^n$. Here again, $\Sigma = M[2i-2] = M[2i-3]$ so the equality $TE \oplus C[2i-3] = M[2i-2]$ still holds.
- Else, $r(X) = 0$. One then sets $m = 2(j-1)-1$, $|M[m]| = n-1$, $M[2i-3] \in \{0,1\}^n$, $M[m-1] = 0^n$ and $M[2i-2]$ is equal to $M[2i-3]$ except on the last bit. We then have:

$$\begin{aligned} \Sigma &= M[2i-2] \oplus \underline{M[m]} \\ &= M[2i-2] \oplus 0^{|M[m]|}1 \\ &= M[2i-3] \end{aligned}$$

  and $TE \oplus C[2i-3] = M[2i-2]$, as before.

In all these cases, we have a distinguishing criteria between the truly random oracle and the real encryption oracle that can be trivially checked. An adversary $\mathcal{A}$ using this algorithm will break the privacy with advantage $\frac{1}{4} - \frac{1}{2^{n+1}}$ with a single encryption query.

Breaking Authenticity for $i \geq j$. The previous attacks against privacy shows that, for any $r(X)$, if there is a type-2 collision among the tweaks polynomials, with $i < j$, one can submit a message M such that its encryption (C, TE) satisfies the equation $TE = C[2i-3] \oplus M[2i-2]$. Informally, by taking this assertion backward, this means that one can compute a valid tag for some specific message from $C[2i-3]$ and $M[2i-2]$. The idea of the authenticity attacks is to query encryption for a message M such that $|M| > 2in$ to get these two bitstrings and then to truncate it to make TE a valid tag for a shorter message of size $\approx 2jn$.

More specifically, we distinguish the following cases:

- If $r(X) = X$, then $\Delta_{i-1,0} = \Delta_{*,j-1,1,1}$. $\mathcal{A}$ selects an integer $m > 2(i-1)$ and submits a message $M = M[1] \dots M[m]$ such that $M[k] = 0^n$ for $k \in [1, 2(j-2)]$, $M[2j-3], M[2j-2] \in \{0,1\}^n$, $M[2i-2] = M[2i-3] = M[2j-2]$ and $M[k] \in \{0,1\}^n$ otherwise. Let (C, TE) be the response to this encryption query. Then, the pair $(C', TE') \leftarrow (C[1] \dots C[2j-4]C[2j-2]C[2j-3], C[2i-3] \oplus M[2i-2])$ is valid (recall that the last two blocks of C are switched during the encryption process), and decrypts to $M' = M[1] \dots M[2j-3]$. Indeed, if M' is the decryption of C', $M'[k] = M[k]$ for $k \leq 2j-2$, $\Sigma' = M'[2j-2]$, the valid tag for C' should be

$$\begin{aligned}\widetilde{TE} &= E_K(\Sigma' \oplus \Delta_{*,j-1,1,1}) \\ &= E_K(M'[2j-2] \oplus \Delta_{*,j-1,1,1}) \\ &= E_K(M[2i-3] \oplus \Delta_{i-1,0}) \\ &= C[2i-3] \oplus M[2i-2] \\ &= TE'\end{aligned}$$

  This clearly breaks the authenticity of the scheme.
- If $r(X) = X + 1$ (and $\Delta_{i-1,0} = \Delta_{*,j-1,1,0}$), then one selects an integer $n > 2(i-1)$ and queries the message $M = M[1] \dots M[m]$ such that $M[k] = 0^n$ for $k \in [1, 2(j-2)]$, $M[2j-3], M[2j-2] \in \{0,1\}^n$, $M[2i-2] = M[2i-3] = M[2j-2]$ and $M[k] \in \{0,1\}^n$ are arbitrary strings otherwise.
  With probability $\frac{1}{2}$, the last bit of $C[2j-3]$ is 1. In this case, $\underline{\texttt{msb}_{n-1}(C[2j-3])} = C[2j-3]$. Let $(C', TE') = (C[1] \dots C[2j-4]C[2j-2]\texttt{msb}_{n-1}(C[2j-3]), C[2i-3] \oplus M[2i-2])$ and M' the decryption of C'. Again, for $k < 2j-3$, $M'[k] = M[k]$, but we also have $M'[2j-3] = M[2j-3]$ and $Z' = C[2j-3] \oplus M[2j-2]$:

$$\begin{aligned}M'[2j-3] &= E_K(\underline{C'[2j-2]} \oplus \Delta_{j-1,1}) \oplus C'[2j-3] \\ &= E_K(\underline{\texttt{msb}_{n-1}(C[2j-3])} \oplus \Delta_{j-1,1}) \oplus C[2j-2] \\ &= E_K(C[2j-3] \oplus \Delta_{j-1,1}) \oplus C[2j-2] \\ &= M[2j-3]\end{aligned}$$

$$\begin{aligned} Z' &= E_K(M'[2j-3] \oplus \Delta_{j-1,0}) \\ &= E_K(M[2j-3] \oplus \Delta_{j-1,0}) \\ &= C[2j-3] \oplus M[2j-2] \end{aligned}$$

As a direct consequence, we also have

$$\begin{aligned} \Sigma' = Z' \oplus \underline{C'[2j-2]} &= C[2j-3] \oplus M[2j-2] \oplus \underline{\mathtt{msb}_{n-1}(C[2j-3])} \\ &= M[2j-2]. \end{aligned}$$

As a consequence, using similar equalities to the $r(X) = X$ case, we can show that the authentication tag for C' should be $\widetilde{TE} = C[2i-3] \oplus M[2i-2] = TE'$. This attack produces a forgery with probability $\frac{1}{2}$.

– If $r(X) = 1$, $\Delta_{i-1,0} = \Delta_{*,j-1,0,1}$. $\mathcal{A}$ again selects $m \geq 2(i-2)$ and queries encryption of $M = M[1] \dots M[m]$ such that $M[k] = 0^n$ for $k \in [1, 2(j-1)]$, $M[2i-3] = 0^n$ and $M[k] \in \{0,1\}^n$ for $k > 2i-2$. Let $(C', TE') = (C[1] \dots C[2j-4]C[2j-3], C[2i-3] \oplus M[2i-2])$ and M' its decryption. Once again, we have $M[k] = M'[k]$ for $k < 2j-3$. Moreover, as the number of blocks in C' is odd,

$$\begin{aligned} M'[2j-3] &= C'[2j-3] \oplus E_K(\Delta_{j-1,0}) \\ &= C[2j-3] \oplus E_K(M[2j-3] \oplus \Delta_{j-1,0}) \\ &= M[2j-2] = 0^n \end{aligned}$$

and hence $\Sigma' = 0^n (= M[2i-3])$. Finally

$$\begin{aligned} TE' = C[2i-3] \oplus M[2i-2] &= E_K(M[2i-3] \oplus \Delta_{i-1,0}) \\ &= E_K(\Sigma' \oplus \Delta_{*,j-1,0,1}) = \widetilde{TE} \end{aligned}$$

where $\widetilde{TE}$ is the expected tag for C'. Again, we are able to produce a forgery.

– If $r(X) = 0$, then one proceeds as in the previous case except that $M[2i-3] = 0^{n-1}1$. We will still have $\Sigma' = M[2i-3]$ and the pair $(C', TE') = (C[1] \dots C[2j-4]\mathtt{msb}_{n-1}(C[2j-3]), C[2i-3] \oplus M[2i-2])$ is a valid forgery.

5.3 Type-3 Collisions

A type-3 collision occurs when there are ℓ and ℓ' such that $X^{\ell+2} + X^{\ell+1} = X^{\ell'+2} + X^{\ell'+1} + r(X)$, with $r(X) \in \{0, 1, X, X+1\}$. We assume, without loss of generality, that $\ell < \ell'$.

The input masks of the form $X^{k+2} + X^{k+1} + r(X)$ are the ones involved in the computation of the tag TE. So a type-3 collision informally means that the input mask used to compute TE for a message of length m' such that $\ell' = \lceil \frac{m'}{2} \rceil$ is the same than the one used to compute TE for a truncated message of length m verifying $\ell = \lceil \frac{m}{2} \rceil$. Again, this leads to a practical attack against authenticity.

Breaking Authenticity. As previously, the attack will slightly differ according to $r(X)$.

- If $r(X) = X$, $\Delta_{*,\ell,0,0} = \Delta_{*,\ell',1,1}$ $\mathcal{A}$ submits an encryption query for the message $M[1] \dots M[2\ell]M[2\ell+1] \dots M[2\ell'-1]M[2\ell']$ with $M[2\ell-1] = 0^n$, $M[2\ell]$ has its last bit set to 1 (in particular $\underline{\mathtt{msb}_{n-1}(M[2\ell])} = M[2\ell]$), and $M[i] = 0^n$ for $i \in [2\ell+1, 2\ell']$. Upon receiving $(C[1] \dots C[2\ell'], TE)$, $\mathcal{A}$ forges $(C', TE') = (C[1] \dots C[2\ell-2]\mathtt{msb}_{n-1}(C[2\ell-1]), TE)$, which is a valid ciphertext.
  Indeed, if Σ is the checksum corresponding to $(C[1] \dots C[2\ell'], TE)$ and Σ' is the one corresponding to the forged ciphertext, we have:

$$\begin{aligned}
\Sigma' &= M[2] \oplus \dots \oplus M[2\ell-2] \oplus \underline{\mathtt{msb}_{n-1}(E_K(\Delta_{\ell,0})) \oplus C'[2\ell-1]} \\
&= M[2] \oplus \dots \oplus M[2\ell-2] \oplus \underline{\mathtt{msb}_{n-1}(E_K(\Delta_{\ell,0}) \oplus C[2\ell-1])} \\
&= M[2] \oplus \dots \oplus M[2\ell-2] \oplus \underline{\mathtt{msb}_{n-1}(M[2\ell])} \\
&= M[2] \oplus \dots \oplus M[2\ell-2] \oplus M[2\ell] \\
&= \Sigma
\end{aligned}$$

  Therefore, $\widetilde{TE} = E_K(\Sigma' \oplus \Delta_{*,\ell,0,0}) = E_K(\Sigma \oplus \Delta_{*,\ell',1,1}) = TE$, so the tag TE is also valid for this truncated ciphertext C'.
- if $r(X) = X + 1$, one proceeds as in the previous case except that we take any value for $M[2\ell]$ and $(C', TE') = (C[1] \dots C[2\ell-2]C[2\ell-1], TE)$: we don't have to play with the padding. Therefore, $\widetilde{TE} = E_K(\Sigma' \oplus \Delta_{*,\ell,0,1}) = E_K(\Sigma \oplus \Delta_{*,\ell',1,1}) = TE$, and TE remains valid for this truncated ciphertext.
- If $r(X) = 1$, $\Delta_{*,\ell,0,0} = \Delta_{*,\ell',0,1}$, and $\mathcal{A}$ will proceed as in the first case $r(X) = X$, except that its first query will be with M with an odd number of blocks. $\mathcal{A}$ will query $M = M[1] \dots M[2\ell'-1]$ such that $M[2\ell-1] = 0^n$, $M[2\ell]$ has its last bit set to 1, and $M[i] = 0^n$ for $i \in [2\ell+1, 2\ell'-1]$. The forgery will be $(C', TE') = (C[1] \dots C[2\ell-2]\mathtt{msb}_{n-1}(C[2\ell-1]), TE)$.
  The proof that (C', TE') is a valid forgery proceeds exactly as for the $r(X) = X$ case.
- if $r(X) = 0$, $\Delta_{*,\ell,0,1} = \Delta_{*,\ell',0,1}$, and $\mathcal{A}$ submits an encryption query on $M = M[1] \dots M[2\ell'-1]$ such that $M[2\ell-1] = 0^n$, and $M[i] = 0^n$ for $i \in [2\ell+1, 2\ell'-1]$. The forgery will be $(C', TE') = (C[1] \dots C[2\ell-2]C[2\ell-1], TE)$. The validity of the forgery can be easily proven from the same arguments as before.

In every case, we are able to easily produce a valid forgery from a single encryption request. For an adversary $\mathcal{A}$ following this procedure,

$$\mathtt{Adv}^{\mathtt{auth}}_{\mathrm{AE}[\tau]}(\mathcal{A}) = 1.$$

6 Practical Security of OTR with 128 Bits Blocks

In the previous sections we exhibited tweak collisions on OTR breaking the security claim, in particular for non generic block sizes (sizes that are not divisible

by 8) and for 64 bits block ciphers. These collisions allow the adversary to break privacy and/or authenticity of the scheme in two encryption/decryption requests with a small number of blocks. Here, we focus on the case $n = 128$.

Also, note that for the sake of breaking OTR, we are only interested in collisions before the birthday bound, *i.e.* collisions for which the maximum index i of the polynomials defined by $\Delta_{i,a}$ or Δ_{*,i,b_1,b_2} is smaller than $2^{n/2}$. Higher order collisions are less interesting as OTR's proofs only guarantees security below the birthday bound.

6.1 Analytical Collisions

One strategy for quickly finding collisions could rely on the fact that $\mathbb{F}_{2^d} \subset \mathbb{F}_{2^{128}}$ for any d dividing 128. Indeed, any relation $Y^i = Y^j + 1$ for some $Y \in \mathbb{F}_{2^d}$ gives us a type-1 collision $X^{a \cdot i} = X^{a \cdot j} + 1$ with a such that $Y = X^a$ in $\mathbb{F}_{2^{128}}$. Such relations can easily be found in $\mathbb{F}_{2^d}$ for $d \in \{16, 32, 64\}$, for example by computing the discrete logarithm of $Y^j + 1$ in base Y. However, they do not lead to truly practical attacks because $Y^{2^d-1} = 1$ (as any element of $\mathbb{F}_{2^d}$) which implies that $2^{128} - 1 | a \cdot (2^d - 1)$ (recall that X generates $\mathbb{F}^*_{2^{128}}$) and so that $(2^{128} - 1)/(2^d - 1)$ divides a. Therefore, such relations will only give collisions for quite large indices $a \cdot i$ (since a is at least greater than $2^{64} + 1$) and so beyond the birthday bound.

6.2 Searching for Collisions Exhaustively

We also tried to algorithmically and exhaustively find collisions among tweaks polynomials. This can be done easily on a desktop computer for $n = 64$, but not for $n = 128$.

Indeed, to check collisions for tweak polynomials of index less than d, we need at least $2d \cdot 128$ bits of memory: the index i polynomials we are interested in are of the form $X^i(+1)$ and $X^i + X^{i-1}(+X)(+1)$, so to save memory, we can only store X^i and $X^i + X^{i-1}$ mod $P(X)$, and do the collision search on the 126 high degree bits. To exhibit a genuine collision, we then just have to recompute the different possibilities for the polynomials and find the matching ones. Also, for each polynomial, we have to store its 'index' i, adding $O(\log d)$ storage. So if we were to exhaustively search for all collisions for $d < 2^{64}$, we would need $2 \cdot 2^{64} \cdot 192$ bits, *i.e.* 24 exabytes.

On the computational point of view, the complexity of the algorithm is well-known, $O(d \log d)$, as we can generate all the $2d$ polynomials, sort them using the lexicographic order on their bits, and finally search a collision in $O(d)$.

It is also important to notice that the collision search is embarrassingly parallelizable: once generated, we can put the polynomials in some bins, depending on the value of the high degree bits, and limit the search to collisions inside each bin. This algorithm is described by Algorithm 2.

Algorithm 2 also offers a nice time/memory tradeoff: instead of keeping all bins in memory, we can instead limit ourself to the bins fitting in memory, and run the algorithms several times so that all the bins are spanned.

Algorithm 2. Our collision search algorithm

for $k = 0$ **to** $2^p - 1$ **do** ▷ *In parallel*
  $S_k \leftarrow \emptyset$ ▷ *Initialize bins*
end for
for $i = 0$ **to** d **do** ▷ *In parallel*
  $\alpha_i \leftarrow X^i \bmod P$
  $k_\alpha \leftarrow \mathtt{msb}_p(\alpha_i)$
  $S_{k_\alpha} \leftarrow S_{k_\alpha} \cup (\alpha_i, i)$
  $\beta_i \leftarrow X^{i+1} + X^i \bmod P$
  $k_\beta \leftarrow \mathtt{msb}_p(\beta_i)$
  $S_{k_\beta} \leftarrow S_{k_\beta} \cup (\beta_i, i)$
end for
for $k = 0$ **to** $2^p - 1$ **do** ▷ *In parallel*
  Lexicographically sort S_k
  Sequentially scan S_k for a collision
end for

We coded this algorithm in C, using OpenMP and SSE instructions, and we were able to show that there is no collisions among the tweak polynomials of index less than 2^{45} for $\mathbb{F}_{2^{128}}$ defined by $X^{128} + X^7 + X^2 + X + 1$, proving Proposition 3, which fixes Lemma 1 of [Min14].

Proposition 3. *For any adversary $\mathcal{A}$ making q queries on $\widetilde{E}$ as defined in Sect. 3, with tweak space $\mathcal{T} = \{0,1\}^{128} \times \{0,\ldots,2^{45}\} \times \{0,1\} \cup \{*\} \times \{0,1\}^{128} \times \{0,\ldots,2^{45}\} \times \{0,1\} \times \{0,1\}$,*

$$\mathtt{Adv}^{\widetilde{\mathtt{prp}}}_{\widetilde{E}}(\mathcal{A}) \leq 5q^2/2^{128}.$$

This exhaustive search took us around 15 CPU-years, using 3TB of RAM.

6.3 Probable Collision Before the Birthday Bound

The collisions exhibited earlier in the paper, for example for $n = 64$ or $n = 119$, use the special form of the polynomial. For the latter, we use the fact that it is a trinomial, directly giving a type-1 collision. For the former, as there are non zero coefficients of two consecutive degrees higher than 2, the polynomial gives a type-2 collision. One could wonder if, excepting these 'trivial' collisions, it is easy to find other before-birthday-bound collisions? Said otherwise, what is the repartition of the indices of colliding polynomials? We can also remember that if the tweak polynomials behaved randomly, we would expect a collision to be happening just before the birthday bound.

We ran experiments for $n = 16, 32$ and 64, using (respectively) irreducible polynomials $X^{16} + X^5 + X^3 + X + 1$, $X^{32} + X^7 + X^3 + X^2 + 1$ and $X^{64} + X^4 + X^3 + X + 1$. They are summarized in Table 1.

If we were to extrapolate, we would expect a collision for $n = 128$ using irreducible polynomial $X^{128} + X^7 + X^4 + X + 1$ to also happen slightly before

Table 1. Lower indices of colliding tweak polynomials (excepted trivial ones).

n	16	32	64
Polynomial	$(X+1)X^{105} = (X+1)X^{134}+X$	$(X+1)X^{30115} = X^{19743}+X$	$X^{2242000936} = X^{2302312163}+1$
log(degree)	7.07	14.88	31.10

the birthday bound. We support this claim with a few experiments we ran on smaller fields. Figures 2, 3 and 4 show the repartition of the smallest collisions of tweak polynomials (*i.e.* the collision with the lowest index) depending on the choice of the irreducible polynomial chosen to define $\mathbb{F}_{2^n}$.

The graphs not only show that the first collision is extremely likely to happen before the birthday bound, but also that it should not happen too early before: we cannot really hope for gaining more than a few bits.

In this case the security proof of [Min14] is only invalidated by a small amount. However, we do not have any formal argument to fill the gap between 2^{45} and 2^{64}.

7 Other Instantiations of Input Masks

The previous collisions do not exclude GF doublings to derive the offsets but simply show that this should be done differently. One of the most obvious solution consists in defining the input mask for the block $M[i]$ as $X^{i+2}\delta$ and Δ_* as

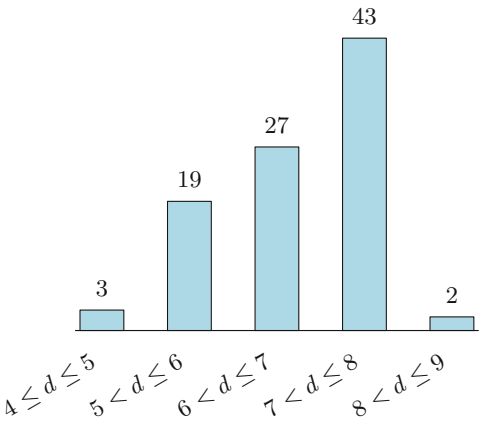

Fig. 2. Log of the lowest indices of colliding tweak polynomials for every $\mathbb{F}_{2^{16}}$ representations using the 94 degree 16 irreducible pentanomials over $\mathbb{F}_2$. In other words, among the 94 possible representations of $\mathbb{F}_{2^{16}}$, 3 leads to a collision between the 2^5 first tweak polynomials, 19 to a collision between polynomials of indices i and j such that $\mathtt{max}(i,j) \in]2^5, 2^6]$, and so on and so forth.

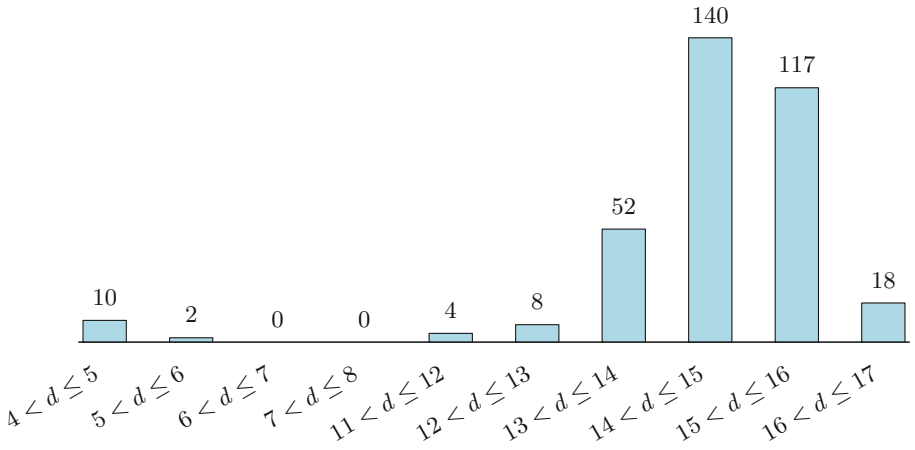

Fig. 3. Log of the lowest indices of colliding tweak polynomials for every $\mathbb{F}_{2^{32}}$ representations using the 351 degree 32 irreducible pentanomials over $\mathbb{F}_2$.

$X^m(X+1)^j\delta$ where m is the number of blocks of M and where j would depend on some properties of M, namely the parity and the number of bits of $M[m]$.

More specifically, the tweakable random function $\widetilde{\rho}$ (see Sect. 3) can be instantiated as follows:

$$\widetilde{E}_K^{N,i,a}(P) = E_K(\Delta_{i,a} + P) \text{ with } \Delta_{i,a} = X^{2(i-1)+a}L$$
$$\widetilde{E}_K^{*,N,i,b_1,b_2}(P) = E_K(\Delta_{*,i,b_1,b_2} + P) \text{ with } \Delta_{*,i,b_1,b_2} = (X+1)^{1+b_2+2b_1}X^{2(i-1)}L$$

where $\delta = E_K(N)$ and $L = X^2\delta$, as previously.

A collision then only occurs if there are some $i, j \in \mathbb{N}^*$ and $a, b_1, b_2 \in \{0,1\}$ such that:

$$X^{2(i-1)+a} = (X+1)^{1+b_2+2b_1}X^{2(j-1)} \quad \Leftrightarrow \quad X^{2(i-j)+a} = (X+1)^{1+b_2+2b_1}$$

However, [Rog04] shows that the latter relation cannot hold for $i, j \leq 2^{115}$ (resp. $i, j \leq 2^{51}$) when $\mathbb{F}_{2^{128}}$ (resp. $\mathbb{F}_{2^{64}}$) is generated by the standard polynomial. A collision attack would thus require to query encryption for a huge message M, whose number of blocks would be far greater than the birthday bound, which is impossible.

Unfortunately, such a solution entails a doubling of the number of multiplications, compared to the original construction. It is therefore preferable to construct $\widetilde{\rho}$ in a slightly different way:

$$\widetilde{E}_K^{N,i,a}(P) = E_K(\Delta_{i,a} + P) \text{ with } \Delta_{i,a} = (X+1)^a X^{i-1}L$$
$$\widetilde{E}_K^{*,N,i,b_1,b_2}(P) = E_K(\Delta_{*,i,b_1,b_2} + P) \text{ with } \Delta_{*,i,b_1,b_2} = (X+1)^{2+b_2+2b_1}X^{i-1}L.$$

Here again, the argument of [Rog04] formally excludes any practical collision attack. The point is that, since $\Delta_{i,1} = \Delta_{i,0} \oplus \Delta_{i+1,0}$, almost one half of the

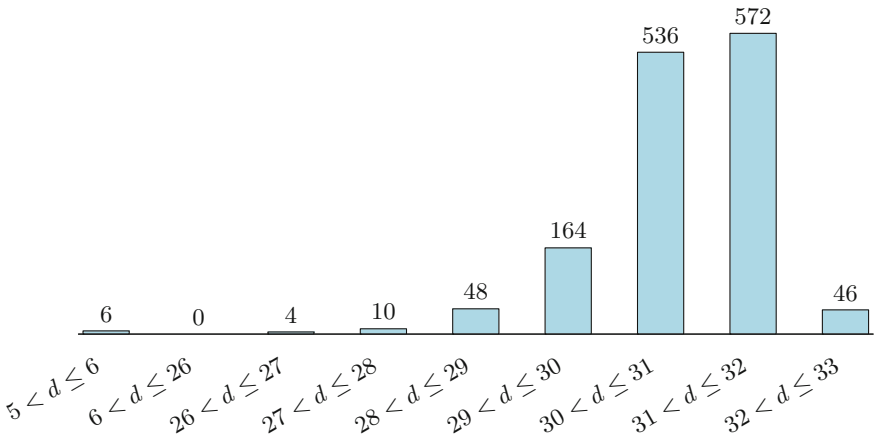

Fig. 4. Log of the lowest indices of colliding tweak polynomials for every $\mathbb{F}_{2^{64}}$ representations using the 1386 degree 64 irreducible pentanomials over $\mathbb{F}_2$.

offsets only require one xor to be computed. The cost is thus similar to the one of the original instantiation [Min14]. The last version of OTR [Min16] uses a similar method to generate tweaks and thus avoid our attack.

8 Conclusion

In this work, we have presented practical attacks against OTR resulting from collisions between the input masks. Although the occurrence of such collisions depend on both the blocksize n and on the polynomial generating $\mathbb{F}_{2^n}$, we argue that the large number of parameters concerned calls for another design of the input masks. We have therefore proposed some ways to immunize OTR to these attacks which do not affect efficiency while being provably secure.

Our results thus do not question the intrinsic security of OTR but simply point out a flaw in the current instantiation.

Acknowledgements. We thank Jean-Gabriel Kammerer for helpful discussions on the implementation of the collision search algorithm, and Julien Devigne for his help. We also thank anonymous reviewers for their useful remarks.

References

[ABB+14] Andreeva, E., Bilgin, B., Bogdanov, A., Luykx, A., Mendel, F., Mennink, B., Mouha, N., Wang, Q., Yasuda, K.: Primates. CAESAR 1st Round (2014). https://competitions.cr.yp.to/caesar-submissions.html

[AFF+15] Abed, F., Fluhrer, S., Forler, C., List, E., Lucks, S., McGrew, D., Wenzel, J.: Pipelineable on-line encryption. In: Cid, C., Rechberger, C. (eds.) FSE 2014. LNCS, vol. 8540, pp. 205–223. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46706-0_11

[AP13] AlFardan, N.J., Paterson, K.G.: Lucky thirteen: breaking the TLS and DTLS record protocols. In: 2013 IEEE Symposium on Security and Privacy, pp. 526–540. IEEE Computer Society Press, May 2013

[BDP+14] Bertoni, G., Daemen, J., Michaël Peeters, G.V., Assche, R.K.: Caesar submission: keyak v1. In: CAESAR 1st Round (2014). https://competitions.cr.yp.to/round1/keyakv1.pdf

[BN00] Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000). doi:10.1007/3-540-44448-3_41

[BRW04] Bellare, M., Rogaway, P., Wagner, D.: The EAX mode of operation. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 389–407. Springer, Heidelberg (2004). doi:10.1007/978-3-540-25937-4_25

[CAE14] Caesar: competition for authenticated encryption: security, applicability and robustness (2014). http://competitions.cr.yp.to/caesar.html

[DN14] Datta, N., Nandi, M.: ELmE: a misuse resistant parallel authenticated encryption. In: Susilo, W., Mu, Y. (eds.) ACISP 2014. LNCS, vol. 8544, pp. 306–321. Springer, Heidelberg (2014). doi:10.1007/978-3-319-08344-5_20

[Dwo04] Dworkin, M.J.: Recommendation for block cipher modes of operation: the CCM mode for authentication and confidentiality, sp. 800–38c. Technical report, National Institute of Standards and Technology (2004)

[Dwo07] Dworkin, M.J.: Recommendation for block cipher modes of operation: galois/counter mode (GCM) and GMAC, spp. 800–38d. Technical report, National Institute of Standards and Technology (2007)

[HKR15] Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 15–44. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46800-5_2

[IK03] Iwata, T., Kurosawa, K.: OMAC: one-key CBC MAC. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 129–153. Springer, Heidelberg (2003). doi:10.1007/978-3-540-39887-5_11

[IMGM15] Iwata, T., Minematsu, K., Guo, J., Morioka, S.: CLOC: authenticated encryption for short input. In: Cid, C., Rechberger, C. (eds.) FSE 2014. LNCS, vol. 8540, pp. 149–167. Springer, Heidelberg (2015). doi:10.1007/978-3-662-46706-0_8

[KR11] Krovetz, T., Rogaway, P.: The software performance of authenticated-encryption modes. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 306–327. Springer, Heidelberg (2011). doi:10.1007/978-3-642-21702-9_18

[Kra01] Krawczyk, H.: The order of encryption and authentication for protecting communications (or: How Secure Is SSL?). In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 310–331. Springer, Heidelberg (2001). doi:10.1007/3-540-44647-8_19

[LRW02] Liskov, M., Rivest, R.L., Wagner, D.: Tweakable block ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer, Heidelberg (2002). doi:10.1007/3-540-45708-9_3

[Min14] Minematsu, K.: Parallelizable rate-1 authenticated encryption from pseudorandom functions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 275–292. Springer, Heidelberg (2014). doi:10.1007/978-3-642-55220-5_16

[Min16] Kazuhiko Minematsu. AES-OTR v3. Technical report, NEC Corporation (2016). https://groups.google.com/group/crypto-competitions/attach/1290d45334f8a3/AESOTR_v3.pdf?part=0.1

[MV04] McGrew, D.A., Viega, J.: The security and performance of the galois/-counter mode (GCM) of operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 343–355. Springer, Heidelberg (2004). doi:10.1007/978-3-540-30556-9_27

[Nik14] Nikolic, I.: Tiaoxin-346. CAESAR Submission (2014)

[Rog04] 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). doi:10.1007/978-3-540-30539-2_2

[Ser98] Seroussi, G.: Table of low-weight binary irreducible polynomials. Technical report, HP (1998). http://www.hpl.hp.com/techreports/98/HPL-98-135.pdf?jumpid=reg_R1002_USEN