

Truncated Boomerang Attacks and Application to AES-based Ciphers*

Augustin Bariant¹, Gaëtan Leurent¹

Inria, Paris, France

Abstract. The boomerang attack is a cryptanalysis technique that combines two short differentials instead of using a single long differential. It has been applied to many primitives, and results in the best known attacks against several AES-based ciphers (Kiasu-BC, Deoxys-BC). In this paper, we introduce a general framework for boomerang attacks with truncated differentials.

We show that the use of truncated differentials provides a significant improvement over the best boomerang attacks in the literature. In particular, we take into account structures on the plaintext and ciphertext sides, and include an analysis of the key recovery step. On 6-round AES, we obtain a competitive structural distinguisher with complexity 2^{87} and a key recovery attack with complexity 2^{61} .

The truncated boomerang attack is particularly effective against tweakable AES variants. We apply it to 8-round Kiasu-BC, resulting in the best known attack with complexity 2^{83} (rather than 2^{103}). We also show an interesting use of the 6-round distinguisher on the full TNT-AES, a tweakable block-cipher using 6-round AES as a building block. Finally, we apply this framework to Deoxys-BC, using a MILP model to find optimal trails automatically. We obtain the best attacks against round-reduced versions of all variants of Deoxys-BC.

Keywords: Truncated differential · Boomerang attack · AES · Kiasu · Deoxys · TNT-AES · MILP

1 Introduction

The AES [DR02] is the most widely used block cipher today, and we have a good understanding of its security. Its round function is strongly byte-aligned; this simplifies the analysis with the wide-trail strategy, and many cryptanalysis techniques rely on truncated trails to take advantage of this property. After 20 years of analysis, we have a high confidence in the design, and many recent tweakable proposals reuse the AES round function with different tweakey schedules (Kiasu-BC [JNP14a], Deoxys-BC [JNPS21], and TNT-AES [BGGS20]). However, an attacker can introduce a difference in the additional tweak of these constructions, so they must be analysed in the related-tweak model, and by extension in the

* © IACR 2023. This article is the full version of the paper published by Springer-Verlag available at https://doi.org/10.1007/978-3-031-30634-1_1.

related-key model. In these models, the boomerang attack is particularly effective because both high-probability differentials composing the boomerang reach more rounds. In particular, the best known attacks against Kiasu-BC and Deoxys-BC are boomerang attacks.

In this work, we carefully and systematically analyse the interaction between truncated differentials and boomerang attacks. Our approach is similar to the analysis of impossible differential attacks in [BLNS18]: we aim at providing a unified formula taking into account many details of a broad class of attacks. We integrate and improve a set of techniques proposed in different variants of the boomerang attack, leading to the best boomerang attacks against several AES-based ciphers.

Our results. We present a generic framework to describe boomerang attacks based on truncated differentials (section 3). Instead of first building a boomerang distinguisher and then appending extra key recovery rounds, we consider the truncated boomerang attack as a whole, including the key recovery exploiting the first and last round transitions. The framework integrates and improves on previous analyses, including structures of plaintexts and ciphertexts [BDK02], and truncated differentials as introduced by Wagner [Wag99]. Our improvements come principally from the use of structures of ciphertexts, which were underused in recent works. We also consider boomerang trails with smaller probability than the random case, and mount attacks by gathering enough samples to detect the bias.

We first apply our framework to reduced AES (section 4). On 6-round AES, we obtain a distinguisher with complexity 2^{87} , and a key-recovery with complexity 2^{61} , improving the previous best boomerang attack with complexity 2^{71} [Bir04].

We adapt the key-recovery attack to 8-round Kiasu-BC (section 5) by revisiting a previous boomerang attack with complexity 2^{103} [DL17]. Using structures of ciphertexts, we obtain the best attack against Kiasu-BC, with complexity 2^{83} .

We also apply a variant of the 6-round attack to the full TNT-AES [BGG20], and obtain a marginal distinguisher with complexity slightly below 2^{128} (section 6). The attack is not competitive with the generic attack against TNT with complexity $\mathcal{O}(\sqrt{n} \cdot 2^{3n/4})$ [GGLS20], but it uses a lower memory (2^{32} instead of 2^{96}), and it can distinguish TNT with 6-round AES from TNT with a PRP. Moreover, this is the first property of 6-round AES that can be used to target a generic construction using 6-round AES as a building block (to the best of our knowledge). We also provide an attack on reduced TNT-AES, using a 5-round boomerang trail.

In section 7, we build a MILP model implementing our framework, to find good parameters for the full attack automatically. The model allows both fixed differences and truncated differences, and takes into account the complexity of the key recovery, instead of just optimizing a boomerang distinguisher. It confirms that our basic attack on AES is optimal within our framework.

Finally, we apply the MILP model to Deoxys-BC, and obtain improved attacks against most variants (section 8). Although the boomerang trails on AES and Deoxys-BC are quite different, the underlying analysis is the same.

Distinguishers and key-recovery attacks. In this work, we report distinguishers and key-recovery attacks, with key-recovery typically having a lower complexity on the same number of rounds. Obviously, a key-recovery attack can be used as a distinguisher, but we focus on structural distinguishers that only use statistical properties of the block cipher, without guessing subkey material (denoted as “independent of the secret key” in [GRR17]). Indeed, a series of recent works have proposed complex distinguishers on 5-round [GRR17] and 6-round AES [BGL20, BR19, RBH17], and we obtain similar results with simpler techniques. This notion of distinguisher is not clearly defined, but our distinguishers can be used with secret S-Boxes, which is not the case for most key-recovery attacks.

2 Preliminaries

2.1 The AES Round Function

AES was designed in 1998 by Daemen and Rijmen (as Rijndael) and won the NIST standardization competition in 2000 [DR02]. Three instances of the cipher exist, for key sizes of 128, 196 and 256 bits, but we only consider AES-128 in this paper. Since we do not exploit the AES key schedule, we only describe the round function. AES-128 operates on a 128-bit state, represented as a 4×4 -byte array, and iterates on 10 rounds a round function composed of the following operations:

- **SubBytes:** The AES S-Box is applied to each byte of the state.
- **ShiftRows:** The second row is shifted by 1 cell to the left, the third row by 2 cells, and the fourth row by 3 cells.
- **MixColumns:** Each column is multiplied by an MDS Matrix.
- **AddRoundKey:** Each byte is XORed with a byte of the round key.

There is one extra AddRoundKey operation before the first round, and the last round omits the MixColumns operation.

Due to the popularity of the AES, and its availability in hardware on several platforms, many constructions reuse its round function. In particular, Kiasu [JNP14b] and Deoxys [JNPS21] are two tweakable block ciphers that reuse the AES round function, with a modified tweakey schedule (combining the key and tweak) to compute the round (tweak)keys. Deoxys has been selected in the CAESAR portfolio. TNT-AES [BGS20] is another tweakable block cipher using the AES round function, where the tweak is only XORed to the internal state twice.

Kiasu-BC tweakey schedule. Kiasu-BC has a 128-bit key and 64-bit tweak, with 10 rounds. The round tweakeys are computed as $k_i + t$ where k_i is the round key following the AES key schedule, and t is the tweak (encoded in the first two rows). In particular, Kiasu-BC with the zero tweak is the same as the AES.

Table 1. AES distinguisher and key recovery attacks with known and secret S-Boxes. CP: chosen plaintexts / ACC: chosen plaintexts and adaptively-chosen ciphertexts

	Rounds	Type	Data	Time	Ref
AES Distinguishers	5	Multiple-of- n	2^{32}	CP	$2^{36.6}$ [GRR17]
	6	Yoyo	$2^{122.8}$	ACC	$2^{121.8}$ [RBH17]
	6	Exchange attack	$2^{88.2}$	CP	$2^{88.2}$ [BR19]
	6	Exchange attack	2^{84}	ACC	2^{83} [Bar19]
	6	Truncated differential	$2^{89.4}$	CP	$2^{96.5}$ [BGL20]
	6	Truncated boomerang	2^{87}	ACC	2^{87} subsection 4.1
AES Key-recovery	6	Square	2^{32}	CP	2^{71} [DKR97]
	6	Partial-sum	2^{32}	CP	2^{48} [FKL ⁺ 01]
	6	Boomerang	2^{71}	ACC	2^{71} [Bir04]
	6	Mixture	2^{26}	CP	2^{80} [BDK ⁺ 20]
	6	Retracing boomerang	2^{55}	ACC	2^{80} [DKRS20]
	6	Boomeyong	$2^{79.7}$	ACC	2^{78} [RSP21]
AES Secret S-Box KR	5	Square	2^{40}	CP	2^{40} [TKKL15]
	5	Multiple-of- n	$2^{53.3}$	CP	$2^{52.6}$ [Gra18]
	5	Retracing boomerang	$2^{25.8}$	ACC	$2^{25.8}$ [DKRS20]
	6	Square	2^{64}	CP	2^{90} [TKKL15]
	6	Truncated boomerang	2^{94}	ACC	2^{94} subsection 4.3

Table 2. Boomerang (B) and rectangle (R) attacks against variants of Deoxys-BC. Most attacks succeed with probability $1/2$.

Model	Rnd	Previous				New					
		Data	Time	Mem	Ref	Data	Time	Mem	Ref		
RTK1	8					B	2^{88}	2^{88}	2^{73}	Figure 9	
	9					B	2^{135}	2^{174}	2^{129}	Figure 10	
RTK2	8	B	2^{28}	2^{28}	2^{27}	[Sas18] ^a	B	2^{27}	2^{27}	2^{27}	Figure 11
	9	B	2^{98}	2^{112}	2^{17}	[Sas18]	B	$2^{55.2}$	$2^{55.2}$	$2^{55.2}$	Figure 12
	10	B	$2^{98.4}$	$2^{109.1}$	2^{88}	[ZDJ19]	B	$2^{94.2}$	$2^{95.2}$	$2^{94.2}$	Figure 5
	11	R	$2^{122.1}$	$2^{249.9}$	$2^{128.2}$	[ZDJ19]	B	2^{129}	$2^{223.9}$	2^{129}	Figure 5
RTK3	10	B	2^{22}	2^{22}	2^{17}	[Sas18]	B	$2^{19.4}$	$2^{19.4}$	2^{18}	Figure 13
	11	B	2^{100}	2^{100}	2^{17}	[Sas18]	B	$2^{32.7}$	$2^{32.7}$	$2^{32.7}$	Figure 14
	12	B	2^{98}	2^{98}	2^{64}	[ZDJ19]	B	$2^{67.4}$	$2^{67.4}$	2^{65}	Figure 15
	13	R	$2^{125.2}$	$2^{186.7}$	2^{136}	[ZDJM19]	B	$2^{126.7}$	$2^{170.2}$	$2^{126.7}$	Figure 6
	14	R	$2^{125.2}$	$2^{282.7}$	2^{136}	[ZDJM19]	B	2^{129}	$2^{278.8}$	2^{129}	Figure 16

^a The probability of Sasaki’s trail is 2^{-56} with structures, thus we believe that the complexity of the attack is actually 2^{30} in data and time and 2^{29} in memory.

Table 3. Attacks against Kiasu-BC and TNT-AES

	Rounds	Type	Data	Time	Ref
Kiasu-BC	7	Square (KR)	$2^{43.6}$	CP	$2^{48.5}$ [DEM16]
	8	Meet-in-the-Middle (KR)	2^{116}	CP	2^{116} [TAY16]
	8	Imposs. Diff (KR)	2^{118}	CP	2^{118} [DL17]
	8	Boomerang (KR)	2^{103}	ACC	$2^{103.1}$ [DL17]
	8	Truncated boomerang (KR)	2^{83}	ACC	2^{83} section 5
TNT-AES	*-5-*	Boomerang (dist.)	2^{126}	ACC	2^{126} [BGG20]
	5-*-*	Impossible differential (KR)	$2^{113.6}$	CP	$2^{113.6}$ [GGLS20]
	--*	Generic (dist.)	$2^{99.5}$	CP	$2^{99.5}$ [GGLS20]
	-5-	Truncated boomerang (dist.)	2^{76}	ACC	2^{76} Appendix B
	5-5-*	Truncated boomerang (KR)	2^{87}	ACC	2^{87} Appendix B
	-6-	Truncated boomerang (dist.)	$2^{127.8}$	ACC	$2^{127.8}$ section 6

Deoxys-BC tweakey schedule. Deoxys-BC has two variants: Deoxys-BC-256 has a 256-bit tweakey with 14 rounds, and Deoxys-BC-384 has a 384-bit tweakey with 16 rounds. The tweakey material is composed of a variable length key and tweak summing to 256 or 384 bits; for simplicity, we assume that the key length is a multiple of 128. The tweakey material is divided in words of 128 bits (denoted TK^i). Eventually, the round tweakey of round j is defined as:

$$STK_j = \begin{cases} RC_j + TK_j^1 + TK_j^2 & \text{For Deoxys-BC-256} \\ RC_j + TK_j^1 + TK_j^2 + TK_j^3 & \text{For Deoxys-BC-384} \end{cases}$$

TK_j^i is the tweakey state, initialized as $TK_0^i = TK^i$ and updated with

$$TK_{j+1}^1 = h(TK_j^1) \quad TK_{j+1}^2 = h(\text{LFSR}_2(TK_j^2)) \quad TK_{j+1}^3 = h(\text{LFSR}_3(TK_j^3))$$

where h is a byte permutation, and LFSR_2 and LFSR_3 are LFSRs that operate in parallel on each byte of the tweakey. This construction (the STK construction [JNP14b]) ensures that differences in subtweakey byte position may only cancel out up to $i - 1$ times every 15 rounds if differences are introduced in i tweakey words.

Notations. We denote E a block cipher operating on a state of n bits. In a 4×4 matrix, the bytes are numbered in the AES order (column-major):

$$\begin{pmatrix} 0 & 4 & 8 & 12 \\ 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \end{pmatrix}$$

A diagonal is one of the four sets of bytes with positions $(0, 5, 10, 15)$, $(4, 9, 14, 3)$, $(8, 13, 2, 7)$, $(12, 1, 6, 11)$ and the main diagonal is the first of the four. An anti-diagonal is one of the four sets of bytes $(0, 7, 10, 13)$, $(1, 4, 11, 14)$, $(2, 5, 8, 15)$, $(3, 6, 9, 12)$ and the main anti-diagonal is the first of the four.

When k_i is a sub(twea)key, we denote $k_i^{eq} = \text{MixColumns}^{-1}(k_i)$.




Fig. 1. Example of a truncated differential trail on 3-round AES.

2.2 Differentials and Truncated Differentials

We use $+$ to denote the XOR operation (the addition in $\mathbb{F}_{2^u}^n$). A differential is defined by an input difference $\Delta_{\text{in}} \in \{0, 1\}^n$ and an output difference $\Delta_{\text{out}} \in \{0, 1\}^n$. We use the notation $\Delta_{\text{in}} \xrightarrow{E} \Delta_{\text{out}}$ when a differential exists with probability p , where the probability is defined over a random plaintext P :

$$p = \Pr[\Delta_{\text{in}} \xrightarrow{E} \Delta_{\text{out}}] = \Pr[E(P) + E(P + \Delta_{\text{in}}) = \Delta_{\text{out}}]$$

Since E is a permutation, we have $\Pr[\Delta_{\text{in}} \xrightarrow{E} \Delta_{\text{out}}] = \Pr[\Delta_{\text{out}} \xrightarrow{E^{-1}} \Delta_{\text{in}}]$.

A truncated differential is defined by a set of input differences \mathcal{D}_{in} and a set of output differences \mathcal{D}_{out} . We use the notation $\mathcal{D}_{\text{in}} \xrightarrow{E} \mathcal{D}_{\text{out}}$ to denote the existence of a truncated differential with probability \vec{p} , defined as (with Avg denoting the average):

$$\vec{p} = \text{Avg}_{\Delta_{\text{in}} \in \mathcal{D}_{\text{in}}} \Pr[E(P) + E(P + \Delta_{\text{in}}) \in \mathcal{D}_{\text{out}}]$$

We also define the probability of the reverse truncated differential as

$$\vec{p} = \text{Avg}_{\Delta_{\text{out}} \in \mathcal{D}_{\text{out}}} \Pr[E^{-1}(P) + E^{-1}(P + \Delta_{\text{out}}) \in \mathcal{D}_{\text{in}}]$$

In general, the two probabilities are different, and related as follow:

$$\frac{\vec{p}}{|\mathcal{D}_{\text{out}}|} = \frac{\vec{p}}{|\mathcal{D}_{\text{in}}|} = \text{Avg}_{\Delta_{\text{in}} \in \mathcal{D}_{\text{in}}, \Delta_{\text{out}} \in \mathcal{D}_{\text{out}}} \Pr[E(P) + E(P + \Delta_{\text{in}}) = \Delta_{\text{out}}]$$

Figure 1 gives an example of a truncated differential on 3 rounds of AES, with respectively 4, 1, and 4 active S-Boxes in each round. On this truncated differential, \mathcal{D}_{in} corresponds to the vector space of elements which have zeros on all diagonals except the main one. \mathcal{D}_{out} is the vector space of states which inverses through the MixColumns operation have 0 values on every anti-diagonal except the main one. Therefore, $|\mathcal{D}_{\text{out}}| = |\mathcal{D}_{\text{in}}| = 2^{32}$. The probability of the truncated differential is $\vec{p} = 2^{-24}$ and the reverse probability is $\vec{p} = 2^{-24}$.

2.3 Boomerang Attacks

Boomerang attacks, introduced by Wagner in 1999 [Wag99], use adaptive plaintext and ciphertext queries to generate quartets with specific differences at an intermediate state of the cipher. The attacker decomposes the full cipher E into two subciphers E_0 (the upper part) and E_1 (the lower part), with $E = E_1 \circ E_0$, with high probability differentials on E_0 and E_1 (of probabilities p and q), denoted respectively $\Delta_{\text{in}} \xrightarrow{p} \Delta_{\text{out}}$ and $\nabla_{\text{in}} \xrightarrow{q} \nabla_{\text{out}}$. The attack proceeds as follows:

1. Generate pairs of plaintext (P, P') such that $P + P' = \Delta_{\text{in}}$, and query the corresponding ciphertexts $(C, C') = (E(P), E(P'))$.
2. Shift the ciphertexts pairs into new pairs $(\bar{C}, \bar{C}') = (C + \nabla_{\text{out}}, C' + \nabla_{\text{out}})$ and query their decryptions $(\bar{P}, \bar{P}') = (E^{-1}(\bar{C}), E^{-1}(\bar{C}'))$.
3. Look for pairs with $\bar{P} + \bar{P}' = \Delta_{\text{in}}$.

Analysis. We have $E_0(P) = E_1^{-1}(C)$ because $E = E_1 \circ E_0$. In particular,

$$E_0(\bar{P}) + E_0(\bar{P}') = E_1^{-1}(\bar{C}) + E_1^{-1}(\bar{C}') + E_0(P) + E_1^{-1}(C) + E_0(P') + E_1^{-1}(C')$$

Moreover, the differentials in E_0 and E_1 imply that:

$$\begin{aligned} \Pr[E_0(P) + E_0(P') = \Delta_{\text{out}}] &= p \\ \Pr[E_1^{-1}(C) + E_1^{-1}(\bar{C}) = \nabla_{\text{in}}] &= q \\ \Pr[E_1^{-1}(C') + E_1^{-1}(\bar{C}') = \nabla_{\text{in}}] &= q \end{aligned}$$

When the three events are satisfied, we obtain $E_0(\bar{P}) + E_0(\bar{P}') = \Delta_{\text{out}}$ and with an additional probability p , $\bar{P} + \bar{P}' = \Delta_{\text{in}}$. Finally, assuming that all events are independent, we compute the boomerang trail probability p_b as a lower bound of the probability of the boomerang relation $\bar{P} + \bar{P}' = \Delta_{\text{in}}$:

$$\Pr[E^{-1}(E(P) + \nabla_{\text{out}}) + E^{-1}(E(P + \Delta_{\text{in}}) + \nabla_{\text{out}}) = \Delta_{\text{in}}] \geq p_b = p^2 \times q^2$$

Figure 2 shows the construction of a boomerang quartet. When $p^2 \times q^2 \gg 2^{-n}$, this gives a distinguisher for the cipher using $\mathcal{O}(p^{-2} \times q^{-2})$ quartets because the probability of detecting a quartet is 2^{-n} for a random permutation. In most cases, the distinguisher can be converted into a key recovery by exploiting key dependencies in the distinguisher.

2.4 Improvements of the Boomerang Attack

Several variants and improvements of the boomerang attack have been proposed since Wagner's original work.

Analysis of the Connection Probability. The analysis above assumes that the four pairs involved in a boomerang quartet follow their corresponding differentials independently. In practice, we usually obtain a probability higher than $p^2 q^2$, but it is also possible for the four events to be incompatible [Mur11]. Several techniques have been proposed to improve this analysis.




Fig. 2. Construction of a boomerang quartet.

Multiple Differentials. Since the differences Δ_{out} and ∇_{in} are not used by the attacker, boomerang quartets can be detected with any internal difference, as long as the same difference is obtained with both pairs. Following the analysis of [Wag99,BDK01], this increases the probability to

$$p_b = \hat{p}^2 \hat{q}^2 \quad \hat{p} = \sqrt{\sum_{\Delta_{\text{out}}} \Pr[\Delta_{\text{in}} \xrightarrow{E_0} \Delta_{\text{out}}]^2} \quad \hat{q} = \sqrt{\sum_{\nabla_{\text{in}}} \Pr[\nabla_{\text{in}} \xrightarrow{E_1} \nabla_{\text{out}}]^2}$$

The Sandwich Attack. Instead of splitting the cipher E into two parts $E = E_1 \circ E_0$, Dunkelman, Keller and Shamir [DKS10] proposed to split it in three parts $E = E_1 \circ E_m \circ E_0$ with a small E_m in the middle. For the analysis, they evaluate the probability of the boomerang trail using the connection probability r of E_m :

$$\Pr[\bar{P} + \bar{P}' = \Delta_{\text{in}}] \geq p_b = p^2 q^2 r$$

$$r = \Pr[E_m^{-1}(E_m(X) + \nabla_{\text{in}}) + E_m^{-1}(E_m(X + \Delta_{\text{out}}) + \nabla_{\text{in}}) = \Delta_{\text{out}}]$$

The connection probability r can be evaluated experimentally, and some specific choices of E_m result in $r = 1$ (in particular, when E_m is the identity, we fall back to the standard analysis of boomerangs). The Boomerang Connectivity Table (BCT) was later introduced [CHP⁺18] to analyze the case where E_m is an S-Box layer. The case where E_m is composed of several rounds has been analyzed in further works [SQH19,WP19,DDV20]. Recent works also show that setting E_m to a single S-Box layer might lead to unaccurate connection probabilities [YSS⁺22].

Plaintext-only Attacks. The amplified boomerang attack [KKS01] and the rectangle attack [BDK01] are variants of the boomerang attack using only encryption queries (without adaptively chosen decryption queries). The complexity increases from $(pq)^{-2}$ to $2^{n/2}(pq)^{-1}$ (with the same condition that $pq \gg 2^{-n/2}$). In this paper, we focus on standard boomerang attacks.

Structures. Biham, Dunkelman and Shamir have introduced a variant of the boomerang attack using structures for the key recovery [BDK02]. They start from a boomerang distinguisher with fixed differences Δ_{in} and ∇_{out} , and add extra rounds at the beginning and at the end. By propagating the differences Δ_{in} and ∇_{out} , they obtain a set of possible input differences \mathcal{D}_{in} and output differences \mathcal{D}_{out} . In a typical SPN cipher, these sets are vector spaces.

The attacker builds a structure $P + \mathcal{D}_{\text{in}} = \{P + \delta : \delta \in \mathcal{D}_{\text{in}}\}$, and uses it as starting point for the attack. A structure of $|\mathcal{D}_{\text{in}}|$ elements defines $|\mathcal{D}_{\text{in}}|^2/2$ pairs, and $|\mathcal{D}_{\text{in}}|/2$ of them lead to the fixed difference Δ_{in} . Therefore, the use of structures covers additional rounds without increasing the data complexity.

Structures can also be used on the ciphertext side, by shifting each ciphertext with all differences in \mathcal{D}_{out} . However, many later works do not use structures on the ciphertext side.

Retracing Boomerang. Dunkelman *et al.* [DKRS20] proposed an improvement where the ciphertexts are chosen so that the two returning pairs (C, \bar{C}) and (C', \bar{C}') are dependant: if one passes the trail on E_1 , the other passes it too. This increases the probability of the trail to $p_b = p^2q$. In the same spirit, Rahman *et al.* [RSP21] recently proposed a boomerang attack embedding a yoyo distinguisher.

3 Truncated Boomerang Attacks

We consider boomerang attacks with truncated differentials, as introduced by Wagner in the original paper [Wag99]. We obtain a key-recovery attack, improving on the use of structures of Biham *et al.* [BDK02] by considering truncated differentials for the full cipher, instead of starting from a shorter boomerang distinguisher with fixed input/output differences and adding truncated trails for the key-recovery rounds. Some boomerang attacks of this type have been proposed on AES [Bir04] and Kiasu-BC [DL17], but they only use structures on the plaintext side. Our framework unifies and improves on previous boomerang attacks, using truncated differentials for E_0 and E_1 , and structures on both sides.

3.1 Truncated Boomerang Distinguisher

Let us consider two truncated differentials $\mathcal{D}_{\text{in}}^0 \xrightarrow{p, E_0} \mathcal{D}_{\text{out}}^0$ and $\mathcal{D}_{\text{in}}^1 \xrightarrow{q, E_1} \mathcal{D}_{\text{out}}^1$ with probabilities \bar{p}, \bar{p} and \bar{q}, \bar{q} on E_0 and E_1 . We assume that $\mathcal{D}_{\text{in}}^0$ is a vector subspace of $\{0, 1\}^n$ and $0 \notin \mathcal{D}_{\text{out}}^1$. The truncated boomerang attack proceeds as follows:

1. Choose a random plaintext P_0 , and query the encryption oracle over the structure $P_0 + \mathcal{D}_{\text{in}}^0$; for each $i \in \mathcal{D}_{\text{in}}^0$, we define $P_i = P_0 + i$ and $C_i = E(P_i)$.
2. For each ciphertext C_i , query the decryption oracle over the set $C_i + \mathcal{D}_{\text{out}}^1$: for each $j \in \mathcal{D}_{\text{out}}^1$, we define $\bar{C}_i^j = C_i + j$ and $\bar{P}_i^j = E^{-1}(\bar{C}_i^j)$.
3. Count the number of pairs $(\bar{P}_i^j, \bar{P}_{i'}^{j'})$ with $\bar{P}_i^j + \bar{P}_{i'}^{j'} \in \mathcal{D}_{\text{in}}^0$ (and $i \neq i'$). This can be done efficiently by projecting the plaintext values on the orthogonal complement of $\mathcal{D}_{\text{in}}^0$ in $\{0, 1\}^n$, and looking for collisions.
4. If needed, repeat steps 1 to 3 with different plaintext structures.

Analysis. We consider a potential quartet $(P, P', \bar{P}, \bar{P}')$ corresponding to $(C, C', \bar{C}, \bar{C}')$, with $P + P' \in \mathcal{D}_{\text{in}}^0$ and $C + \bar{C}, C' + \bar{C}' \in \mathcal{D}_{\text{out}}^1$. We have:

$$\begin{aligned}\Pr[E_0(P) + E_0(P') \in \mathcal{D}_{\text{out}}^0] &= \vec{p} \\ \Pr[E_1^{-1}(C) + E_1^{-1}(\bar{C}) \in \mathcal{D}_{\text{in}}^1] &= \vec{q} \\ \Pr[E_1^{-1}(C') + E_1^{-1}(\bar{C}') \in \mathcal{D}_{\text{in}}^1] &= \vec{q}\end{aligned}$$

Following the sandwich attack analysis (with $E_m = \text{id}$), we define the connection probability:

$$r = \Pr \left[E_0(\bar{P}) + E_0(\bar{P}') \in \mathcal{D}_{\text{out}}^0 \left| \begin{array}{l} E_0(P) + E_0(P') \in \mathcal{D}_{\text{out}}^0 \\ E_1^{-1}(C) + E_1^{-1}(\bar{C}) \in \mathcal{D}_{\text{in}}^1 \\ E_1^{-1}(C') + E_1^{-1}(\bar{C}') \in \mathcal{D}_{\text{in}}^1 \end{array} \right. \right]$$

If the four events hold, we have $\bar{P} + \bar{P}' \in \mathcal{D}_{\text{in}}^0$ with an additional probability \vec{p} . This analysis of the truncated boomerang distinguisher is the same as proposed by Wagner [Wag99], but our attack is more general with structures on both sides.

In general, we have $E_1^{-1}(C) + E_1^{-1}(\bar{C})$ and $E_1^{-1}(C') + E_1^{-1}(\bar{C}')$ in $\mathcal{D}_{\text{in}}^1$, therefore they are equal with probability $|\mathcal{D}_{\text{in}}^1|^{-1}$, and this implies $E_0(\bar{P}) + E_0(\bar{P}') \in \mathcal{D}_{\text{out}}^0$; hence $r \geq |\mathcal{D}_{\text{in}}^1|^{-1}$. Moreover, if $\mathcal{D}_{\text{in}}^1$ and $\mathcal{D}_{\text{out}}^0$ are vector subspaces, then $\Sigma = E_0(P) + E_0(P') + E_0(\bar{P}) + E_0(\bar{P}') \in \mathcal{D}_{\text{in}}^1$; in particular, $\Sigma \in \mathcal{D}_{\text{out}}^0$ with probability $|\mathcal{D}_{\text{out}}^0 \cap \mathcal{D}_{\text{in}}^1| / |\mathcal{D}_{\text{in}}^1|$; this implies $E_0(\bar{P}) + E_0(\bar{P}') \in \mathcal{D}_{\text{out}}^0$ hence $r \geq |\mathcal{D}_{\text{out}}^0 \cap \mathcal{D}_{\text{in}}^1| / |\mathcal{D}_{\text{in}}^1|$.

Assuming that all the events are independent, each quartet $(P_i, P_{i'}, \bar{P}_i^j, \bar{P}_{i'}^{j'})$, defined by a pair $(i, j), (i', j')$, follows the truncated boomerang trail with probability p_b , and randomly satisfies $\bar{P}_i^j + \bar{P}_{i'}^{j'} \in \mathcal{D}_{\text{in}}^0$ with probability p_{\S} :

$$p_b = \vec{p} \cdot \vec{p} \cdot \vec{q}^2 \cdot r \qquad r \geq |\mathcal{D}_{\text{in}}^1|^{-1} \qquad (1)$$

$$p_{\S} = |\mathcal{D}_{\text{in}}^0| / 2^n \qquad (2)$$

We assume that the boomerang probability can be well approximated as $p_b + p_{\S}$. We distinguish the cipher E from a random permutation when the expected number of remaining quartets (quartets with $\bar{P}_i^j + \bar{P}_{i'}^{j'} \in \mathcal{D}_{\text{in}}^0$) is significantly higher for E than for a random permutation. We define the signal-to-noise ratio:

$$\sigma = p_b / p_{\S} \qquad (3)$$

When $\sigma \gg 1$, we obtain a distinguisher using $Q = \mathcal{O}(p_b^{-1})$ quartets. More precisely, with $Q = \mu \cdot p_b^{-1}$ (μ a small constant) we expect μ remaining quartets with the cipher E , versus $\mu \cdot \sigma^{-1} \ll 1$ for a random permutation. A distinguisher that detects the presence of at least one quartet has a success rate of $1 - e^{-\mu}$.

When σ is smaller, we need to collect a large number of quartets, and compare the expected number of remaining quartets q_b for E and q_{\S} in the random case:

$$q_b = Q \times (p_{\S} + p_b) = Q \times p_{\S}(1 + \sigma) \qquad q_{\S} = Q \times p_{\S}$$




Fig. 3. A truncated boomerang trail on 6-round AES.

We detect the bias with $Q = \mathcal{O}(p_{\mathbb{S}}^{-1}\sigma^{-2}) = \mathcal{O}(p_b^{-1}\sigma^{-1})$ samples, following [MS02, Theorem 2]. Using $Q = c \times p_b^{-1}\sigma^{-1}$ with a small constant c and setting a threshold at $Q \times p_{\mathbb{S}}(1 + \sigma/2)$, the distinguisher has a success rate of $\Phi(\sqrt{c}/2)$, with Φ the cumulative distribution function of the standard normal distribution.

If Q is smaller than the number of quartets in a full structure ($|\mathcal{D}_{in}^0|^2 |\mathcal{D}_{out}^1|^2 / 2$), we use a partial structure with only $\sqrt{2Q}$ elements. Otherwise, we need $N = 2Q \times |\mathcal{D}_{in}^0|^{-2} |\mathcal{D}_{out}^1|^{-2}$ structures of $S = |\mathcal{D}_{in}^0| |\mathcal{D}_{out}^1|$ elements. Finally, we obtain a distinguisher with a constant probability of success with the following complexity in number of quartets, time, data, and memory:

$$Q = \mathcal{O}(\max(p_b^{-1}, \sigma^{-1} \cdot p_b^{-1})) \quad (4)$$

$$T = D = \max(\sqrt{2Q}, 2Q \times |\mathcal{D}_{in}^0|^{-1} |\mathcal{D}_{out}^1|^{-1}) \quad (5)$$

$$M = \min(D, |\mathcal{D}_{in}^0| |\mathcal{D}_{out}^1|) \quad (6)$$

Application to 6-round AES. To explain the truncated boomerang distinguisher in practice, we give a truncated boomerang trail on 6-round AES in figure 3, using the 3-round trail of figure 1 twice. $\mathcal{D}_{\text{in}}^0$ and $\mathcal{D}_{\text{in}}^1$ are the sets of all states that have zeros on all diagonals except the main one. $\mathcal{D}_{\text{out}}^0$ is the same as the output set of figure 1, and $\mathcal{D}_{\text{out}}^1$ is active on the main anti-diagonal (it differs because we omit the last MixColumns operation). We have

$$\begin{aligned} |\mathcal{D}_{\text{in}}^0| &= |\mathcal{D}_{\text{out}}^0| = 2^{32} & \vec{p} &= 2^{-24} & \bar{p} &= 2^{-24} \\ |\mathcal{D}_{\text{in}}^1| &= |\mathcal{D}_{\text{out}}^1| = 2^{32} & \vec{q} &= 2^{-24} & \bar{q} &= 2^{-24} \end{aligned}$$

Since $\mathcal{D}_{\text{out}}^0 \cap \mathcal{D}_{\text{in}}^1 = \{0\}$, we have $r = |\mathcal{D}_{\text{in}}^1|^{-1}$, and the analysis above gives the following parameters:

$$\begin{aligned} p_b &= \vec{p} \cdot \bar{p} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-128} & \sigma &= 2^{-32} \\ p_{\mathbb{S}} &= |\mathcal{D}_{\text{in}}^0|/2^n = 2^{-96} & Q &= c \cdot 2^{160} \end{aligned}$$

Algorithm 1. Truncated boomerang distinguisher on 6-round AES

```

 $q \leftarrow 0$ 
for  $0 \leq s < 2^{35}$  do ▷  $2^{35}$  structures
   $P_0 \leftarrow \mathbb{S}$ 
  for  $i \in \mathcal{D}_{\text{in}}^0$  do ▷ Iterate over the main diagonal
     $C_i \leftarrow E(P_0 + i)$ 
    for  $j \in \mathcal{D}_{\text{out}}^1$  do ▷ Iterate over the main anti-diagonal
       $\bar{P}_i^j \leftarrow E^{-1}(C_i + j)$ 
      Store  $\bar{P}_i^j$  in a hash table indexed by three diagonals
    Count collisions in the hash table, and increment  $q$ 
if  $q > 2^{66} + 2^{33}$  then
  return AES
else
  return  $\mathbb{S}$ 

```

Using $c = 4$ and the formulas of Equations (4), (5), and (6), we obtain a distinguisher with complexity:

$$T = D = 2^{99} \qquad M = 2^{64}$$

The distinguisher is detailed in Algorithm 1. It makes 2^{67} encryption queries and 2^{99} decryption queries, for a total data complexity of $D = 2^{67} + 2^{99} \approx 2^{99}$. In total, we have $Q = 2^{35} \times 2^{64} \times 2^{64}/2 = 2^{162}$ quartets $(P_i, P_{i'}, \bar{P}_i^j, \bar{P}_{i'}^{j'})$, so that the expected number of remaining quartets is:

$$q_{\mathbb{S}} = Q \times 2^{-96} = 2^{66} \qquad q_E = Q \times (2^{-96} + 2^{-128}) = 2^{66} + 2^{34}$$

The distinguisher returns the correct answer with probability $\Phi(\sqrt{c}/2) \approx 0.84$.

This distinguisher is interesting because it is very generic: it does not require knowledge of the S-Box or the MDS matrix, and it can be considered as “key-independent” in the sense of [GRR17]. As seen in Table 1, the complexity is slightly higher than previous distinguishers with similar properties. However, the simplicity of this distinguisher makes it more likely to be applicable when 6-round AES is used as a building block in a more complex structure, as shown in section 6 against TNT-AES.

3.2 Truncated Boomerang Key-recovery Attack

We now consider key-recovery attacks. As opposed to typical differential or linear attacks, we do not add rounds on top of the distinguisher. Instead, we assume that the truncated boomerang covers the full targeted cipher, and we design a key-recovery attack with smaller complexity than the corresponding distinguisher.

When $\sigma \geq 1$, the truncated boomerang distinguisher is easy to turn into a key-recovery attack, but we cannot reduce the complexity. Indeed, the bottleneck of the distinguisher is to have enough data so that a boomerang quartet exists. When a quartet with $\overline{P} + \overline{P}' \in \mathcal{D}_{\text{in}}^0$ is found, it has a high probability of following the boomerang, and standard methods can be used to recover key candidates. Therefore, we focus on the case $\sigma \ll 1$, where the distinguisher requires multiple quartets following the boomerang.

Given a candidate quartet with $\overline{P} + \overline{P}' \in \mathcal{D}_{\text{in}}^0$, we can extract some key information assuming that it follows the boomerang. If this is the case, we have two pairs of known plaintexts (P, P') and $(\overline{P}, \overline{P}')$ following the truncated differential $\mathcal{D}_{\text{in}}^0 \xrightarrow{p} \mathcal{D}_{\text{out}}^0$, and two pairs of known ciphertexts (C, \overline{C}) and (C', \overline{C}') following the truncated differential $\mathcal{D}_{\text{in}}^1 \xrightarrow{q} \mathcal{D}_{\text{out}}^1$. Using standard techniques from differential cryptanalysis, we can usually extract partial information about the first and last subkeys. We denote by κ the number of key bits that can be extracted, and by ℓ the average number of κ -bit key candidates suggested by a quartet. Note that the key information suggested by a quartet might be incompatible between both pairs of plaintexts following the upper differential (or between both pairs of ciphertexts following the lower differential), in this case the quartet is discarded.

We follow the standard approach to identify the most likely candidates for the κ bits of key: we build a table of 2^κ counters corresponding to key candidates, and we increment the counter of each key suggested by each quartet. With enough data, the right key is expected to be among the top $2^{\kappa-a}$ counters (a denotes the advantage of the attack).

Analysis. Following the previous analysis, we expect $Q \times (p_{\text{s}} + p_{\text{b}})$ quartets with $\overline{P} + \overline{P}' \in \mathcal{D}_{\text{in}}^0$: $Q \times p_{\text{b}}$ quartets following the boomerang trail (right quartets), and $Q \times p_{\text{s}}$ false positives. For a right quartet, the correct key is among the deduced key candidates, and for a wrong quartet, we expect that ℓ random key candidates are deduced. Assuming that all the quartets behave independently,

the wrong counters follow the binomial distribution $\mathcal{B}(Q, (p_s + p_b) \times \ell \times 2^{-\kappa})$ and the right counter follows the distribution $\mathcal{B}(Q, p_s \times \ell \times 2^{-\kappa} + p_b)$. We denote the probabilities of suggesting a wrong key and the right key as:

$$p_w = (p_s + p_b) \times \ell \times 2^{-\kappa} \approx p_s \times \ell \times 2^{-\kappa} \quad (7)$$

$$p_0 = p_s \times \ell \times 2^{-\kappa} + p_b \approx p_w + p_b \quad (8)$$

We obtain a higher signal-to-noise ratio $\tilde{\sigma}$ than previously:

$$\tilde{\sigma} = p_b/p_w = \sigma \times 2^\kappa/\ell \quad (9)$$

When $\tilde{\sigma} \gg 1$, only a handful of right quartets are necessary to have the right key ranked first, so that $Q = \mathcal{O}(p_b^{-1})$.

When $\tilde{\sigma} \ll 1$, the counters can be approximated by normal distributions, and we use the work of Selçuk [Sel08, Theorem 3] to evaluate the number of samples needed to have the right key among the top $2^{\kappa-a}$ key candidates (depending on the success rate). For a fixed value of a , we need Q proportional to $p_b^{-1}\tilde{\sigma}^{-1}$, and the complexity increases linearly in a . Finally, the increased signal-to-noise ratio $\tilde{\sigma} \gg \sigma$ reduces the data complexity to:

$$Q = \mathcal{O}(\max(p_b^{-1}, \tilde{\sigma}^{-1} \times p_b^{-1})) \quad (10)$$

$$D = \max(\sqrt{2Q}, 2Q \times |\mathcal{D}_{\text{in}}^0|^{-1} |\mathcal{D}_{\text{out}}^1|^{-1}) \quad (11)$$

The full attack is described in Algorithm 2. The time complexity is harder to evaluate; it can be bounded with T_E the cost of an oracle call (by convention, $T_E = 1$), and T_C the cost of deducing key candidates from a quartet:

$$T = D \times T_E + Q \times p_s \times T_C \quad (12)$$

When $T_C \ll 2^n \times |\mathcal{D}_{\text{in}}^0|^{-2} |\mathcal{D}_{\text{out}}^1|^{-1}$, we have $Q \times p_s \times T_C \ll D$ and the second term is negligible; the cost of the attack is thus dominated by the oracle queries. Otherwise, it is often possible to reduce the second term with more advanced filtering, but this requires a dedicated analysis for each attack.

After recovering $2^{\kappa-a}$ candidates for the κ -bit partial key, the full key can be recovered by exhaustive search of the remaining bits with complexity 2^{n-a} , or by launching a variant of the attack on a different set of key bits.

Success Probability. When $\tilde{\sigma} \ll 1$, the average values of right and wrong counters are high enough to approximate them with normal distributions. In that case, the success rate can be evaluated using the formula given by [Sel08], under additional assumptions about the independence of key counters, and order statistics:

$$P_S = \Phi\left(\frac{\sqrt{\mu\tilde{\sigma}} - \Phi^{-1}(1 - 2^{-a})}{\sqrt{\tilde{\sigma} + 1}}\right) \quad (13)$$

with $\mu = Q \times p_b$ the expected number of right quartets.

When $\tilde{\sigma}$ is high, the binomial distributions of right and wrong counters have their average values respectively $Q \times p_b \approx 1$ and $Q \times p_w \ll 1$. As discussed in [Sel08, section 3.2.1], the normal approximation is inaccurate in this case; we obtain a more accurate estimate of the success probability using a Poisson approximation.

Algorithm 2. Truncated boomerang key recover attack

Require: $\mathcal{D}_{\text{in}}^0, \mathcal{D}_{\text{out}}^1$
 $\mathcal{K} \leftarrow \text{InitKeyCounters}()$
for $i \leftarrow 1$ to N **do**
 $P_0 \leftarrow \text{Rand}()$
 $\mathcal{P} \leftarrow [P_0 + \Delta_{\text{in}}, \text{for } \Delta_{\text{in}} \in \mathcal{D}_{\text{in}}^0]$
 $\mathcal{C} \leftarrow [E(P), \text{for } P \in \mathcal{P}]$ $\triangleright N \times |\mathcal{D}_{\text{in}}^0|$ encryptions
 $\overline{\mathcal{P}} \leftarrow [E^{-1}(C + \Delta_{\text{out}}), \text{for } C \in \mathcal{C}, \Delta_{\text{out}} \in \mathcal{D}_{\text{out}}^1]$ $\triangleright N \times |\mathcal{D}_{\text{in}}^0| |\mathcal{D}_{\text{out}}^1|$ decryptions
 $\mathcal{H} \leftarrow \text{InitHashMap}()$
 for $\overline{P} \in \overline{\mathcal{P}}$ **do**
 Insert the projection of \overline{P} on $\{0, 1\}^n / \mathcal{D}_{\text{in}}^0$ in \mathcal{H} $\triangleright |\mathcal{D}_{\text{in}}^0| |\mathcal{D}_{\text{out}}^1|$ in memory
 if a collision occurs in \mathcal{H} between \overline{P} and \overline{P}' **then**
 Track back to the corresponding P and P'
 if $P \neq P'$ **then** $\triangleright N(p_{\text{rand}} + p_{\text{trunc}}) |\mathcal{D}_{\text{in}}^0|^2 |\mathcal{D}_{\text{out}}^1|^2 / 2$ such quartets
 for K in the ℓ key candidates induced by the quartet **do**
 Increment the counter K in \mathcal{K}
 Recover the key material K of the maximal counter of \mathcal{K}
 Recover the rest of the key by performing the attack with a different subspace $\mathcal{D}_{\text{in}}^0$

Extracting Key Candidates. When the truncated differentials are described by truncated trails (with a set of intermediate differences at each step), the parameters ℓ and κ can often be deduced directly from the trail. We assume that E_0 starts with the addition of a subkey K_0 , followed by an S-Box layer SB , and we denote the set of differences after the S-Box layer by $\mathcal{D}_{\text{mid}}^0$:

$$E_0 = \tilde{E}_0 \circ \text{SB} \circ \text{AK}_{K_0} \quad \mathcal{D}_{\text{in}}^0 \xrightarrow{\text{SB}} \mathcal{D}_{\text{mid}}^0 \quad \mathcal{D}_{\text{mid}}^0 \xrightarrow{\tilde{E}_0} \mathcal{D}_{\text{out}}^0$$

We also assume that $\mathcal{D}_{\text{in}}^0$ is a vector subspace aligned with the S-Box layer (each S-Box is either inactive, or active with all possible differences). $\mathcal{D}_{\text{mid}}^0$ is a subset of $\mathcal{D}_{\text{in}}^0$; typically it is constructed so that some parts of the state have fixed differences after the linear layer. For instance, in the AES trail of [Figure 1](#), $\mathcal{D}_{\text{mid}}^0$ corresponds to differences δ such that $\text{MixColumns}(\text{ShiftRows}(\delta))$ is active only on the first cell, with $|\mathcal{D}_{\text{mid}}^0| = 2^8$ and $\vec{p}_0 = 2^{-24}$. In general, we have:

$$\vec{p}_0 = |\mathcal{D}_{\text{mid}}^0| / |\mathcal{D}_{\text{in}}^0| \quad \vec{p} = \vec{p}_0 \times \vec{p}_1 \quad (14)$$

We consider a pair (P, P') , and assume that it follows the truncated trail, *i.e.* $\text{SB}(P + K_0) + \text{SB}(P' + K_0) \in \mathcal{D}_{\text{mid}}^0$. This constrains the partial subkey $K_{0|\mathcal{D}_{\text{in}}^0}$ corresponding to the active S-Boxes in SB . More precisely, for each difference δ in $\mathcal{D}_{\text{mid}}^0$, we expect on average 0.5 unordered pairs $\{X, X'\}$ such that $X + X' = P + P'$ and $\text{SB}(X) + \text{SB}(X') = \delta$ (restricted to the active bytes $\mathcal{D}_{\text{mid}}^0$). This pair can be recovered efficiently after pre-computing the DDT of the S-Box, and we deduce two possible keys $X + P$ and $X + P'$. Therefore, we have the following parameters when extracting key candidates from a pair (P, P') :

$$\ell^0 = |\mathcal{D}_{\text{mid}}^0| \quad \kappa^0 = \log_2(|\mathcal{D}_{\text{in}}^0|) \quad T_C^0 = \ell^0 = |\mathcal{D}_{\text{mid}}^0|$$

Starting from a candidate quartet, we have two different pairs (P, P') and (\bar{P}, \bar{P}') assumed to follow the upper differential. Therefore, we expect only $|\mathcal{D}_{\text{mid}}^0|^2/|\mathcal{D}_{\text{in}}^0|$ key candidates compatible with both pairs. We apply the same reasoning to the lower trail (using ciphertext pairs), and deduce the parameters ℓ and κ for a quartet in the general case:

$$\ell = |\mathcal{D}_{\text{mid}}^0|^2 \cdot |\mathcal{D}_{\text{mid}}^1|^2 \cdot |\mathcal{D}_{\text{in}}^0|^{-1} \cdot |\mathcal{D}_{\text{out}}^1|^{-1} \quad \kappa = \log_2(|\mathcal{D}_{\text{in}}^0| \cdot |\mathcal{D}_{\text{out}}^1|) \quad (15)$$

Using the probability \vec{p}_0 for the first round and \vec{q}_0 for the last round, we have

$$\ell \cdot 2^{-\kappa} = \vec{p}_0^2 \cdot \vec{q}_0^2 \quad (16)$$

For the lower trail, we only have to process a fraction $|\mathcal{D}_{\text{mid}}^0|^2/|\mathcal{D}_{\text{in}}^0|$ of the candidate quartets (with a key compatible with both pairs). In particular, when $|\mathcal{D}_{\text{mid}}^0|^2 \ll |\mathcal{D}_{\text{in}}^0|$, the time complexity is dominated by the first extraction step: $T_C = |\mathcal{D}_{\text{mid}}^0|$.

Application to 6-round AES. This attack can directly be applied to AES, using the same 3-round trails as in the previous section (see Figure 3):

$$\begin{aligned} |\mathcal{D}_{\text{in}}^0| = |\mathcal{D}_{\text{out}}^0| = 2^{32} & & |\mathcal{D}_{\text{mid}}^0| = 2^8 & & \vec{p} = 2^{-24} & & \bar{p} = 2^{-24} \\ |\mathcal{D}_{\text{in}}^1| = |\mathcal{D}_{\text{out}}^1| = 2^{32} & & |\mathcal{D}_{\text{mid}}^1| = 2^8 & & \vec{q} = 2^{-24} & & \bar{q} = 2^{-24} \end{aligned}$$

Using the parameters of the key extraction, our analysis results in

$$\begin{aligned} \ell &= |\mathcal{D}_{\text{mid}}^0|^2 \cdot |\mathcal{D}_{\text{mid}}^1|^2 \cdot |\mathcal{D}_{\text{in}}^0|^{-1} \cdot |\mathcal{D}_{\text{out}}^1|^{-1} = 2^{-32} & \kappa &= \log_2(|\mathcal{D}_{\text{in}}^0| \cdot |\mathcal{D}_{\text{out}}^1|) = 64 \\ p_b &= \vec{p} \cdot \bar{p} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-128} \\ p_w &= |\mathcal{D}_{\text{in}}^0| \times 2^{-n} \times \ell \times 2^{-\kappa} = 2^{-192} & \tilde{\sigma} &= 2^{64} \end{aligned}$$

Since $\tilde{\sigma} \gg 1$, we only need a few right quartets; with $\mu = 4$ we obtain

$$Q = \mu \times p_b^{-1} = 2^{130} \quad D = 2^{67}$$

Time complexity. With these parameters, the attack complexity is dominated by the oracle queries. We use 8 structures of 2^{64} elements; in each structure we detect $2^{64} \times 2^{63} \times p_{\mathcal{S}} = 2^{31}$ pairs with $\bar{P} + \bar{P}' \in \mathcal{D}_{\text{in}}^0$, resulting in $8 \times 2^{31} = 2^{34}$ candidate quartets in total. Each quartet suggests on average 2^{-32} candidates for 64 bits of key (for most of the quartets, there is no key compatible with both sides of the quartet). Finally, we expect 2^2 suggestions of wrong keys (each key is suggested 2^{-62} times on average), and $\mu = 4$ suggestions for the correct key. With high probability, the key with the most suggestions is the correct one.


We implemented the attack on a reduced AES with 4-bit S-Boxes, and it behaves as expected [BL23].

4 Optimized Boomerang Attacks on 6-round AES

As shown by Biryukov [Bir04], boomerang attacks on AES can be optimized using multiple trails. We now present improved versions of our attacks using this technique, including a 6-round key-recovery attack with complexity 2^{61} . The improvement compared to the attack of Biryukov with complexity 2^{71} is due to the use of structures on the ciphertext side.

4.1 Optimized distinguisher


Instead of only considering the trail of Figure 1 with fixed positions for all the active bytes, we consider a collection of four different trails for upper part:



The collection can be considered as a truncated differential $\mathcal{D}_{\text{in}}^0 \xrightarrow{p/E_0} \mathcal{D}_{\text{out}}^0$ with

$$\vec{p} = 2^{-22} \quad \vec{p} = 2^{-24} \quad |\mathcal{D}_{\text{in}}^0| = 2^{32} \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

Similarly, we consider four trails for the lower part:




Again, this can be considered as a truncated differential $\mathcal{D}_{\text{in}}^1 \xrightarrow{q/E_1} \mathcal{D}_{\text{out}}^1$ with

$$\vec{q} = 2^{-24} \quad \vec{q} = 2^{-22} \quad |\mathcal{D}_{\text{in}}^1| = 2^{34} \quad |\mathcal{D}_{\text{out}}^1| = 2^{32}$$

The analysis of the previous sections can be applied as-is with these trails. We obtain a better attack because we increase \vec{p} and \vec{q} by a factor 4, even though $|\mathcal{D}_{\text{in}}^1|$ increases by a factor 4; we obtain $p_b = 2^{-124}$ instead of 2^{-128} . The distinguisher is exactly the same because $\mathcal{D}_{\text{in}}^0$ and $\mathcal{D}_{\text{out}}^1$ are the same, but this improved analysis shows that the complexity of the distinguisher can be reduced to $T = D = 2^{91}$ (with $c = 4$, $\sigma = 2^{-28}$ and $Q = 2^{154}$).

Larger set $\mathcal{D}_{\text{out}}^1$. We further improve the distinguisher using a collection of 16 trails with the following input and output sets for the lower trail:



This collection can be considered as a truncated differential $\mathcal{D}_{\text{in}}^1 \xrightarrow{q/E_1} \mathcal{D}_{\text{out}}^1$ with

$$\vec{q} = 2^{-22} \quad \vec{q} = 2^{-22} \quad |\mathcal{D}_{\text{in}}^1| = 2^{34} \quad |\mathcal{D}_{\text{out}}^1| = 2^{34}$$

This does not affect the probability p_b , but generates larger structures; the complexity is reduced to $T = D = 2^{89}$ with $Q = 2^{154}$.

Different Set $\overline{\mathcal{D}}_{\text{in}}^0$ for Returning Pairs. Following Biryukov [Bir04], we use a higher probability differential for the returning pair $(\overline{P}, \overline{P}')$, different from for the initial pair (P, P') , and with a larger set $\overline{\mathcal{D}}_{\text{in}}^0$. We consider the same collection of 16 trails as above, corresponding to a truncated differential $\overline{\mathcal{D}}_{\text{in}}^0 \xrightarrow{E_0} \mathcal{D}_{\text{out}}^0$ with

$$\vec{p} = 2^{-22} \quad \vec{\bar{p}} = 2^{-22} \quad |\overline{\mathcal{D}}_{\text{in}}^0| = 2^{34} \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

This corresponds to keeping quartets with a single active diagonal in $\overline{P} + \overline{P}'$, but not necessarily the main one. We adapt our analysis to account for the two distinct upper differentials and we obtain

$$\begin{aligned} p_b &= \vec{p} \cdot \vec{\bar{p}} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-122} & \sigma &= 2^{-28} \\ p_{\mathbb{S}} &= |\overline{\mathcal{D}}_{\text{in}}^0|/2^n = 2^{-94} & Q &= 2^{152} \end{aligned}$$

Finally, we obtain a distinguisher with complexity $T = D = 2^{87}$ (with $c = 4$).

4.2 Optimized Key-recovery Attack

Using multiple trails. For a key-recovery attack, we use the trails above, but we keep the set $\mathcal{D}_{\text{out}}^1$ active only in the first column.

$$\begin{aligned} \vec{p} &= 2^{-22} & \vec{\bar{p}} &= 2^{-24} & |\mathcal{D}_{\text{in}}^0| &= 2^{32} & |\mathcal{D}_{\text{out}}^0| &= 2^{34} \\ \vec{\bar{p}} &= 2^{-22} & \vec{p} &= 2^{-22} & |\overline{\mathcal{D}}_{\text{in}}^0| &= 2^{34} & |\mathcal{D}_{\text{out}}^0| &= 2^{34} \\ \vec{q} &= 2^{-24} & \vec{\bar{q}} &= 2^{-22} & |\mathcal{D}_{\text{in}}^1| &= 2^{34} & |\mathcal{D}_{\text{out}}^1| &= 2^{32} \end{aligned}$$

When extracting the key, we recover some information on a diagonal of k_0 from (P, P') (but not $(\overline{P}, \overline{P}')$ since it not necessarily active on the same diagonal as (P, P')), and information about an anti-diagonal of k_6 from (C, \overline{C}) and (C', \overline{C}') :

$$\ell^0 = 2^{10} \quad \kappa^0 = 32 \quad \ell^1 = 2^{-14} \quad \kappa^1 = 32 \quad \ell = 2^{-4} \quad \kappa = 64$$

Therefore, we have the following parameters:

$$\begin{aligned} p_b &= \vec{p} \cdot \vec{\bar{p}} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-122} \\ p_w &= |\overline{\mathcal{D}}_{\text{in}}^0| \times 2^{-n} \times \ell \times 2^{-\kappa} = 2^{-162} & \tilde{\sigma} &= 2^{40} \end{aligned}$$

Since $\tilde{\sigma} \gg 1$, we only need a few right quartets; with $\mu = 4$ we obtain

$$Q = \mu \times p_b^{-1} = 2^{124} \quad D = 2^{62.5}$$

Time complexity. The complexity is dominated by the oracle queries: for each structure of 2^{64} plaintexts/ciphertexts, we filter $2^{64} \times 2^{63} \times |\overline{\mathcal{D}}_{\text{in}}^0| \times 2^{-n} = 2^{33}$ candidate quartets, and the time to extract the key candidates is negligible.

After recovering a candidate for 64 bits of key (32 bits of k_0 and 32 bits of k_6), we repeat the attack with $\mathcal{D}_{\text{in}}^0$ in a different diagonal and use the partial knowledge of k_6 to increase the probability \vec{q} (this has negligible complexity).

Using two active diagonals. In order to further reduce the complexity of this attack, we use truncated boomerang characteristics with lower signal-to-noise ratios, taking advantage of the additional filter provided by the key extraction. Following [Bir04], we modify the truncated trail on the returning side (\bar{P}, \bar{P}') to allow any combination of two active diagonals in input, leading to the following parameters:

$$\bar{\mathcal{D}}_{\text{in}}^0 = \left\{ \begin{array}{c} \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array}, \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array}, \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array}, \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array}, \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array}, \begin{array}{|c|c|c|c|} \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \blacksquare & \blacksquare & \blacksquare & \blacksquare \\ \hline \end{array} \right\}$$

$$\begin{array}{llll} \vec{p} = 2^{-22} & \vec{\bar{p}} = 2^{-24} & |\mathcal{D}_{\text{in}}^0| = 2^{32} & |\mathcal{D}_{\text{out}}^0| = 2^{34} \quad |\mathcal{D}_{\text{mid}}^0| = 2^{10} \\ \vec{\bar{p}} = 2^{-46} & \vec{\bar{p}} = 6 \times 2^{-16} = 2^{-13.4} & |\bar{\mathcal{D}}_{\text{in}}^0| = 6 \times 2^{64} & |\mathcal{D}_{\text{out}}^0| = 2^{34} \\ \vec{q} = 2^{-24} & \vec{\bar{q}} = 2^{-22} & |\mathcal{D}_{\text{in}}^1| = 2^{34} & |\mathcal{D}_{\text{out}}^1| = 2^{32} \quad |\mathcal{D}_{\text{mid}}^1| = 2^{10} \end{array}$$

When extracting the key, we recover information about the main diagonal of k_0 from (P, P') , and information about the first anti-diagonal of k_6 from (C, \bar{C}) and (C', \bar{C}') (note that (\bar{P}, \bar{P}') is not necessarily active in the main diagonal). Moreover, the key suggested by (C, \bar{C}) and (C', \bar{C}') must lead to the same active byte in z_4 , so that

$$\ell^0 = 2^{10} \quad \kappa^0 = 32 \quad \ell^1 = 2^{-14} \quad \kappa^1 = 32 \quad \ell = 2^{-4} \quad \kappa = 64$$

Using the previous analysis, we obtain

$$\begin{aligned} p_b &= \vec{p} \cdot \vec{\bar{p}} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-113.4} \\ p_w &= |\bar{\mathcal{D}}_{\text{in}}^0| 2^{-n} \times \ell \times 2^{-\kappa} = 2^{-129.4} \quad \tilde{\sigma} = 2^{16} \end{aligned}$$

Since $\tilde{\sigma} \gg 1$, a few right quartets are sufficient for the success of this attack; we use $\mu = 8$, this corresponds to $Q = 2^{116.4}$ and we use a partial structure of $D = 2^{58.7}$ elements.

Success probability. We assume that the attacker keeps key candidates with counter values of at least 5. With $\tilde{\sigma} \gg 1$, we approximate the wrong key counters by Poisson distributions with $\lambda = Q \times p_w = 2^{-13}$, each of which equal 5 or more with probability $1 - e^{-\lambda}(1 + \lambda + \lambda^2/2 + \lambda^3/6 + \lambda^4/24) \approx 2^{-71.9}$; we don't expect to keep any wrong keys. On the other hand, the counter for the right key follows a Poisson distribution with $\lambda = \mu = 8$. It reaches a value of 5 or more with probability ≈ 0.9 .

Time complexity. After recovering a candidate for 64 bits of key (32 bits of k_0 and 32 bits of k_6), we repeat the attack with $\mathcal{D}_{\text{in}}^0$ in a different diagonal and use the partial knowledge of k_6 to increase the probability \vec{q} . This step has a negligible complexity.

The time complexity is balanced between oracle queries and extracting key candidates. Indeed, we filter $2^{58.7} \times 2^{57.7} \times |\bar{\mathcal{D}}_{\text{in}}^0| \times 2^{-n} = 2^{55}$ candidates with

$\bar{P} + \bar{P}' \in \bar{\mathcal{D}}_{\text{in}}^0$ using 6 hash tables indexed by each combination of two active columns. The complexity T_C to generate key candidates for a given quartet is essentially 4×2^{10} accesses to a small table; we approximate it as $T_C \approx 2^{5.4} T_E$ (since one encryption has 6×16 S-Boxes). Finally, the time complexity is

$$T = 2^{58.7} T_E + 2^{55} T_C \approx 2^{60.8} T_E$$

4.3 Key-recovery with Secret S-Boxes

The techniques described in subsection 3.2 assume that the S-Box and MDS matrix are known to the attacker in order to extract key information. However, it is also possible to extract key information with an unknown S-Box under some conditions. Following [GRR16], we assume that all S-Boxes in a column are identical, and that the MDS matrix has two identical coefficients in each row.

As a concrete example, we consider the AES MixColumns matrix

$$\text{MC} = \begin{bmatrix} 2 & 3 & 1 & 1 \\ 1 & 2 & 3 & 1 \\ 1 & 1 & 2 & 3 \\ 3 & 1 & 1 & 2 \end{bmatrix}$$

We consider a pair (C, \bar{C}) following the truncated trail of Figure 3. According to the trail, the difference before the last round (w_4) is in a set of 2^8 differences; in particular, the difference in cell 1 is equal to the difference in cell 2:

$$w_4 + \bar{w}_4 \in \left\{ \begin{bmatrix} 2\delta & 0 & 0 & 0 \\ \delta & 0 & 0 & 0 \\ \delta & 0 & 0 & 0 \\ 3\delta & 0 & 0 & 0 \end{bmatrix} : \delta \in \{0, 1\}^8 \right\} = \left\{ \text{MC} \cdot \begin{bmatrix} \delta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} : \delta \in \{0, 1\}^8 \right\}$$

Moreover, we assume that the differences in cells 13 and 10 of the ciphertext are equal (they are moved to cell 1 and 2 by ShiftRows)

$$C + \bar{C} = \begin{bmatrix} \alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & \beta \\ 0 & 0 & \beta & 0 \\ 0 & \gamma & 0 & 0 \end{bmatrix}$$

In this case, S-Boxes 1 and 2 in the last round follow the same transition $\delta \rightarrow \beta$. With high probability, this implies that the unordered pairs of input/output are equal; in particular $\{C[13] + k_6[13], \bar{C}[13] + k_6[13]\} = \{C[10] + k_6[10], \bar{C}[10] + k_6[10]\}$. This suggests two key candidates:

$$k_6[13] + k_6[10] \in \{C[13] + C[10], C[13] + \bar{C}[10]\}$$

In order to use this property in a truncated boomerang attack, we use the multiple upper trails of subsection 4.1, and a single lower trail with a restricted $\mathcal{D}_{\text{out}}^1$ of size 2^{24} to ensure that $C + \bar{C}$ and $C' + \bar{C}'$ have the required properties for all quartets considered:

$$\mathcal{D}_{\text{out}}^1 = \left\{ \begin{bmatrix} a & 0 & 0 & 0 \\ 0 & 0 & 0 & b \\ 0 & 0 & b & 0 \\ 0 & c & 0 & 0 \end{bmatrix} : a, b, c \in \{0, 1\}^8 \right\}$$

The corresponding parameters are:

$$\begin{array}{llll}
\vec{p} = 2^{-22} & \vec{p} = 2^{-24} & |\mathcal{D}_{\text{in}}^0| = 2^{32} & |\mathcal{D}_{\text{out}}^0| = 2^{34} \\
\vec{p} = 2^{-22} & \vec{p} = 2^{-22} & |\overline{\mathcal{D}}_{\text{in}}^0| = 2^{34} & |\mathcal{D}_{\text{out}}^0| = 2^{34} \\
\vec{q} = 2^{-32} & \vec{q} = 2^{-24} & |\mathcal{D}_{\text{in}}^1| = 2^{32} & |\mathcal{D}_{\text{out}}^1| = 2^{24}
\end{array}$$

For each quartet, the pair (C, \overline{C}) suggests two values for $k_6[13] + k_6[10]$, and C', \overline{C}' also suggests two values. Therefore a quartet suggests on average $\ell = 2^{-6}$ values for $\kappa = 8$ bits of key. Using the analysis of [subsection 3.2](#), we obtain:

$$\begin{array}{ll}
p_b = \vec{p} \cdot \vec{p} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-124} & \bar{\sigma} = 2^{-16} \\
p_w = |\overline{\mathcal{D}}_{\text{in}}^0| \times 2^{-n} \times \ell \cdot 2^{-\kappa} = 2^{-108} & Q = \mathcal{O}(2^{140})
\end{array}$$


To obtain a high probability of success we use $Q = 2^{145}$, *i.e.* $D = 2^{90}$. Since $\bar{\sigma} \ll 1$, the counter distribution of the right key can be approximated to the normal distribution $\mathcal{N}(2^{37} + 2^{21}, 2^{37})$ while wrong key counters distributions can be approximated to $\mathcal{N}(2^{37}, 2^{37})$. We expect the correct key to be ranked first with very high probability ($P_S > 0.99$ using the formula from [\[Sel08\]](#)).

The time complexity is dominated by the oracle queries: for each structure of 2^{56} plaintexts/ciphertexts, we filter $2^{56} \times 2^{55} \times |\overline{\mathcal{D}}_{\text{in}}^0| \times 2^{-n} = 2^{17}$ candidate quartets with $\overline{P} + \overline{P}' \in \overline{\mathcal{D}}_{\text{in}}^0$, and the time to extract the key candidates is negligible. We can repeat the attack to recover up to 16 key bytes in different positions, with a complexity of $D = T = 2^{94}$ (but only 12 recovered bytes are linearly independent).

5 Application to 8-round Kiasu-BC

Kiasu-BC [\[JNP14a\]](#) is an instance of the TWEAKEY framework [\[JNP14b\]](#), reusing the AES round function in a tweakable block cipher. The 6-round boomerang attack on the AES can be extended to 8-round Kiasu-BC by taking advantage of the tweak input to cancel state differences in order to have one inactive round in the upper and lower trails. Indeed, the best known attack on Kiasu-BC is an 8-round attack with complexity 2^{103} in data and time [\[DL17\]](#) following this idea; the corresponding boomerang is represented in [Figure 4](#). Following our framework, we improve this attack with a better use of structures.

Truncated Boomerang. Since we use a tweak difference Δ_{tw} , we slightly generalize our truncated differential framework to allow a set of tweak differences \mathcal{D}_{tw} . We start from a 4-round truncated trail $(\mathcal{D}_{\text{in}}, \mathcal{D}_{\text{tw}}) \xrightarrow{P/E} \mathcal{D}_{\text{out}}$ with probability 2^{-32} , similar to the 3-round trail of previous sections:






Fig. 4. Truncated boomerang trail on 8-round Kiasu-BC, along with the key-recovery attack parameters. S is the structure size.

\mathcal{D}_{tw} is the set of differences active in the first cell of the tweak; following the tweaky schedule of Kiasu-BC, this results in a tweaky difference in \mathcal{D}_{tw} at each round. We obtain an 8-round boomerang with two 4-round differentials:

$$\begin{array}{ccccc} \vec{p} = 2^{-32} & \vec{p} = 2^{-32} & |\mathcal{D}_{\text{in}}^0| = 2^{32} & |\mathcal{D}_{\text{out}}^0| = 2^{32} & |\mathcal{D}_{\text{tw}}^0| = 2^8 \\ \vec{q} = 2^{-32} & \vec{q} = 2^{-32} & |\mathcal{D}_{\text{in}}^1| = 2^{32} & |\mathcal{D}_{\text{out}}^1| = 2^{32} & |\mathcal{D}_{\text{tw}}^1| = 2^8 \end{array}$$

Following the analysis of the AES attack in subsection 3.2, we deduce on average $\ell = 2^{-32}$ candidates of $\kappa = 64$ key bits per quartet. Therefore, we obtain

$$\begin{aligned} p_b &= \vec{p} \cdot \vec{p} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-160} \\ p_w &= |\mathcal{D}_{\text{in}}^0|/2^n \times \ell \times 2^{-\kappa} = 2^{-192} & \tilde{\sigma} &= 2^{32} \end{aligned}$$

Since $\tilde{\sigma} \gg 1$, we only need a few right quartets. Taking $\mu = 4$, we obtain an attack with $Q = 2^{162}$ quartets. We take advantage of the tweak to build larger structures (iterating over the tweak and data inputs), of size $|\mathcal{D}_{\text{in}}^0| \cdot |\mathcal{D}_{\text{tw}}^0| \cdot |\mathcal{D}_{\text{out}}^1| \cdot |\mathcal{D}_{\text{tw}}^1| = 2^{80}$. Thus we only need 8 structures, with data complexity $D = 2^{83}$. In each structure of 2^{80} elements, we expect 2^{63} quartets with $\bar{P} + \bar{P}' \in \mathcal{D}_{\text{in}}^0$, therefore the time complexity for the key recovery is negligible, and $T = D = 2^{83}$.

Success Probability. There are 2^{66} quartets with $\bar{P} + \bar{P}' \in \mathcal{D}_{\text{in}}^0$, suggesting on average 2^{-32} key candidates each; hence a total of 2^{34} candidates for 64 bits of key. We keep key candidates whose counter reaches 2 or more. Modeling counters for wrong keys with a Poisson distribution with $\lambda = 2^{-30}$, the probability for a specific wrong key counter to be at least 2 is $1 - e^{-\lambda}(1 + \lambda) \approx 2^{-61}$; therefore we expect to keep 8 wrong keys. On the other hand, the counter for the right key follows a Poisson distribution with $\lambda = 4$. It reaches a value of 2 or more with probability ≈ 0.9 .

As in the AES attacks, we recover the full key by repeating the attack with $\mathcal{D}_{\text{in}}^0$ in a different diagonal. Taking advantage of the recovered values of the last round key, this adds a negligible complexity.

6 Application to TNT-AES

TNT-AES is a tweakable block cipher reusing the AES round function published at Eurocrypt 2020 [BGG20]. It is part of the Tweak-aNd-Tweak framework, building a tweakable block cipher \tilde{E} from a block cipher E :

$$\tilde{E}_{K_0, K_1, K_2} : P, T \mapsto C = E_{K_2} \left(T + E_{K_1} (T + E_{K_0} (P)) \right)$$

In order to improve its efficiency, TNT-AES uses a 6-round AES as building block E . The designers of TNT proved its security up to $2^{2n/3}$ queries, and conjectured a higher security bound. Later work [GGLS20] proved the bound to be at least $\Omega(2^{3n/4})$ queries, and exhibited a distinguisher with $\mathcal{O}(\sqrt{n} \cdot 2^{3n/4})$ queries.

Truncated Boomerang. Our attack focuses on the middle cipher E_{K_1} , between both tweak additions. In order to skip the initial and final ciphers E_{K_0} and E_{K_2} , we introduce differences in the tweak, instead of introducing them in the plaintext and ciphertext. We fix a plaintext P , and consider four tweaks T, T', \bar{T}, \bar{T}' to create quartets as follows:

1. Query $C = \tilde{E}(P, T)$ and $C' = \tilde{E}(P, T')$
2. Query $\bar{P} = \tilde{E}^{-1}(C, \bar{T})$ and $\bar{P}' = \tilde{E}^{-1}(C', \bar{T}')$
3. Detect when $\bar{P} = \bar{P}'$

We denote the inputs and outputs of E_{K_1} as X and Y , with $Y = E_{K_1}(X)$:

$$\begin{aligned} X &= E_{K_0}(P) + T & X' &= E_{K_0}(P) + T' & \bar{X} &= E_{K_0}(\bar{P}) + \bar{T} & \bar{X}' &= E_{K_0}(\bar{P}') + \bar{T}' \\ Y &= E_{K_2}^{-1}(C) + T & Y' &= E_{K_2}^{-1}(C') + T' & \bar{Y} &= E_{K_2}^{-1}(C) + \bar{T} & \bar{Y}' &= E_{K_2}^{-1}(C') + \bar{T}' \end{aligned}$$

When $\bar{P} = \bar{P}'$, we have a boomerang quartet for E_{K_1} with differences


$$\begin{aligned} X + X' &= T + T' = \Delta_{\text{in}} & \bar{X} + \bar{X}' &= \bar{T} + \bar{T}' = \Delta'_{\text{in}} \\ Y + \bar{Y} &= T + \bar{T} = \nabla_{\text{out}} & Y' + \bar{Y}' &= T' + \bar{T}' = \nabla'_{\text{out}} \end{aligned}$$

When using a truncated boomerang trail (with a fixed P and a set of tweaks), there are two important limitations compared to the previous attacks:

- We only detect when the difference $\bar{X} + \bar{X}'$ matches exactly $\bar{T} + \bar{T}'$, instead of detecting a set of differences $\mathcal{D}_{\text{in}}^0$. The boomerang trail probability decreases.
- We necessarily have $\Delta_{\text{in}} + \Delta'_{\text{in}} = \nabla_{\text{out}} + \nabla'_{\text{out}}$. For the 6-round AES truncated boomerang trail of figure 3, this implies $\Delta_{\text{in}} = \Delta'_{\text{in}}$ and $\nabla_{\text{out}} = \nabla'_{\text{out}}$. Therefore, we cannot take advantage of structures on the ciphertext side.

Nonetheless, truncated boomerangs can be used with structures of tweaks on the plaintext side, and the analysis of the middle rounds as truncated differentials significantly reduces the complexity compared to the analysis of [BGG20].

Upper Differential. We use the same collection of 4 upper trails as in our optimized attack on AES:



We have the following parameters

$$\vec{p} = 2^{-22} \quad \vec{p} = 2^{-24} \quad |\mathcal{D}_{\text{in}}^0| = 2^{32} \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

For the return trail, we must hit a fixed $\bar{T} + \bar{T}' = \Delta_{\text{in}}^0$:

$$\vec{p} = 2^{-22} \quad \vec{p} = 2^{-56} \quad |\bar{\mathcal{D}}_{\text{in}}^0| = 2^0 \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

Lower Differential. Since we cannot use structures on the ciphertext side, we use a fixed value Δ_{out}^1 to maximize the probability of the trail. We observe that in an AES column, the transition $\delta \rightarrow (*, 0, 0, 0)$ through a layer of inverse S-Boxes followed by inverse MixColumns happens with probability $2272/2^{32} \approx 2^{-20.85}$ with $\delta = (L(\beta/2), L(\beta), L(\beta), L(\beta/3))$, with L the linear transform inside the AES S-Box (see Appendix A.1). We choose $\Delta_{\text{out}}^1 = \text{MixColumns}(\text{ShiftRows}(\delta))$:

$$\vec{q} = 2^{-52.85} \quad \bar{q} = 2^{-20.85} \quad |\mathcal{D}_{\text{in}}^1| = 2^{32} \quad |\mathcal{D}_{\text{out}}^1| = 2^0$$

Boomerang Trail Probability. We obtain:

$$p_b = \vec{p} \cdot \bar{p} \cdot \bar{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-151.7} \quad p_{\S} = |\overline{\mathcal{D}}_{\text{in}}^0|/2^n = 2^{-128}$$

As shown in Appendix A, we obtain a slightly better probability p_b by carefully analyzing the boomerang, and correlation between the sides:

$$p_b = 2^{-151.4}$$

It is not possible to recover actual key material with this attack because X is unknown. However, we can use $E_{K_0}(P) + K_1$ as an equivalent subkey if all queries are made with the same P . Using the pair (X, X') we extract $\ell = 2^{10}$ candidates for $\kappa = 32$ key bits. Unfortunately, we cannot use the pairs (Y, \bar{Y}) for filtering on the ciphertext side since the unknown value Y is different in each quartet. Similarly, the pair (\bar{X}, \bar{X}') is unusable for key extraction. Therefore,

$$p_b = 2^{-151.4} \quad p_w = p_{\S} \times \ell \times 2^{-\kappa} = 2^{-150} \quad \tilde{\sigma} = 2^{-1.4}$$

With $\tilde{\sigma} < 1$, we need $Q = c \cdot \tilde{\sigma}^{-1} \cdot p_b^{-1}$ with a small constant c ; we take $c = 64$, $Q = 2^{158.8}$. Since we have structures of size 2^{32} , this corresponds to $D = 2^{127.8}$.

Distinguisher. With $2^{127.8}$ queries we obtain a distinguisher between TNT-AES (using 6-round AES as the building block) and a PRP (or TNT using a PRP). This obviously does not threaten the security of TNT-AES, but we believe that it is an interesting use case showing that a 6-round boomerang distinguisher can be extended to a larger scheme, even if the attack is marginal.

In order to minimize the number of queries, we use the 255 possible values of $\Delta_{\text{out}}^1 = (L(\beta/2), L(\beta), L(\beta), L(\beta/3))$ with $\beta \in \mathbb{F}_{256} \setminus \{0\}$, so that each encryption query is amortized: we obtain $2^{158.8}$ quartets with $2^{127.8}/255$ encryption queries and $2^{127.8}$ decryption queries. After collecting the quartets, we expect that the counter corresponding to the right key follows the distribution $\mathcal{N}(2^{8.8} + 2^{7.4}, 2^{8.8})$ while counters for the wrong keys follow the distribution $\mathcal{N}(2^{8.8}, 2^{8.8})$ (the distance between the expected values is 8 times the standard deviation).

We obtain a distinguisher by observing whether the maximum counter is higher than a threshold $t = 2^{8.8} + 7 \times 2^{4.4}$. The probability that all counters for wrong keys are lower than t is $\Phi(7)^{2^{32}} \approx 0.995$, therefore the probability of false positive is 0.005. The probability that the counter for the right key is higher than t is $\Phi(1) = 0.84$ so the probability of false negative is 0.16.

Finally, we can increase the success rate by running three attacks in parallel using three input sets $\mathcal{D}_{\text{in}}^0$, $\mathcal{D}'_{\text{in}}{}^0$, $\mathcal{D}''_{\text{in}}{}^0$ on three different diagonals. Using superstructures of 2^{96} values, we run all three attacks with the same queries, and generate counters for three sets of 2^{32} equivalent keys. Using a threshold of $t = 2^{8.8} + 7.1 \times 2^{4.4}$, we keep the probability of false positive below 1%, while the probability that at least one of the three counters corresponding to right keys is higher than the threshold increases to 99%.

TNT with 5-round AES. A reduced version of TNT-AES with 5 AES rounds can be attacked more efficiently, using a probability-one truncated differential for one half of the cipher. In Appendix B we give a distinguisher with 2^{76} queries, and a key-recovery with complexity 2^{87} .

7 Modeling the framework using MILP

We now introduce a MILP model to search for truncated boomerang attacks on AES-based ciphers, based on [CHP⁺17]. It uses our analysis of truncated boomerangs to evaluate the data complexity of an attack, and minimizes it.

Mixed Integer Linear Programming (MILP) is a mathematical optimization tool that minimizes a linear objective function of constrained variables. The variables can be either integer or rational, and the constraints are given as linear inequalities. In the last years, it has proven to be a useful tool to evaluate the security of cryptographic primitives, due to the facility of encoding cryptographic properties as MILP problems, and the availability of high performance solvers. Differential trails can be modeled with a MILP program in order to obtain bounds on the probability of trails (as first used on SIMD [BFL11]), or to search for good trails (as first used by Mouha *et al.* [MWGP11]).

This method was applied to the search of boomerang distinguisher on Deoxys by Cid *et al.* [CHP⁺17]. Their MILP model encodes the activity of each state byte with a binary variable that equals 1 if its corresponding byte is active and 0 if not, and that is constrained depending on the activity pattern of Deoxys operations. In order to build a boomerang trail, their model includes two separate differential trails with two overlapping rounds in the middle (in order to account for the ladder switch and the BCT analysis). The objective function to minimize is roughly the number of active S-Boxes, *i.e.* the sum of all variables representing the activity of S-Box input (or output) bytes.

After generating the optimal boomerang template, they instantiate active bytes with concrete differences that maximize S-Box transition probabilities. An important contribution of their work is an analysis of the degrees of freedom of the tweakey differences. Their MILP model counts the number of linear relations between the tweakey differences and ensures that at least one degree of freedom remains in the final trail, otherwise it is unlikely to find concrete differences for the tweakey.

In 2019, Zhao *et al.* [ZDJ19,ZDJM19] improved this MILP model by adding two extra rounds at the end of the lower trail, containing truncated differences.

Our MILP model. Previous works [CHP⁺17,Sas18,ZDJM19] showed that the best attack is not always obtained with the best distinguisher. Therefore we follow the same high-level approach as in [QDW⁺21]: our main objective is to cover the full boomerang attack with the MILP model. We give an overview of the model in Appendix C, and for more details the full code is available [BL23]. Our model is based on [CHP⁺17]; the main improvement is to allow 4 possible states for each byte of the trail, instead of only 2 (active or inactive):

- inactive, with a zero difference, denoted as \square ;
- active with a *fixed* non-zero difference, denoted as \blacksquare ;
- active with an *unknown* (truncated) difference, denoted as \blacksquare .
- active with an *equal* (but unknown) difference for both pairs, denoted as \blacksquare .

From these possible states, we can deduce the parameters of the attacks, including the structure size, which allows us to ask the MILP solver to minimize the formula for the data complexity of the attack given in subsection 3.2.

Bytes with *equal* differences encode relations between the two different pairs that follow the same trail, rather than properties of a trail by itself. This allows the MILP model to capture trails like the 6-round AES boomerang of Figure 3 or the 8-round Kiasu boomerang of Figure 4; Figure 8 shows how they are captured by this model. The model does not encode linear relations between active bytes (*e.g.* the set of differences for w_2 of Figure 3 is active on all bytes but has size 2^{32}), but using this type of constraint is sufficient in many cases because it is propagated through the linear layer.

Our model generates boomerang trails without instantiating the differences. In order to derive concrete attacks from the trail, we instantiate the trail with differences that optimize the different S-Box transitions.

Limitations of the MILP model. Our model handles only fixed differences in the tweakey. More importantly, although our model takes into account the ladder switch to compute the boomerang connection probability, it does not accurately compute the connection probability. Some boomerang trails given in a previous version of this paper were even found incompatible by Song et al. [YSS⁺22]. More generally, some boomerang trails returned by the MILP are not instantiable. When that happens in practice, we modify the boomerang trail skeleton or we generate a different trail using the MILP solver.

To instantiate boomerang trails generated by the MILP, we use the tables DDT, BCT [CHP⁺18], UBCT, LBCT and EBCT, introduced in [SQH19,WP19,DDV20]. In order to ensure that the trails are not incompatible, we verified experimentally the probability of the middle rounds [BL23]. For each trail, we indicate the rounds that have been checked, with the theoretical probability p_{th} for the middle rounds, and the experimental value observed p_{exp} . In some attacks, the experimental probability slightly differs from the theoretical one; we deduce an adjusted trail probability \tilde{p}_b that is used to calculate the complexity of the attack.

7.1 Results on AES-128 and Kiasu-BC

We use the MILP model to search for attacks on AES-128 and Kiasu-BC and compare them with the results of the previous sections.

On AES-128, the model returns the trail of Figure 8 of Appendix D, which corresponds to Figure 3 with a full *equal* state in w_2 . This confirms the optimality of our truncated trail within our framework. However, the model does not handle multiple trails, so that it cannot suggest the improved attack of section 4.

On Kiasu-BC, the model cannot find the attack of section 5, because it does not handle truncated differences in the tweak. With fixed difference in the tweak, the best trail found by the solver is unfortunately not instantiable, because of incompatibilities in the middle rounds. The solver nevertheless ensures that no attack with complexity less than 2^{80} exist in that framework.

8 Application to Deoxys-BC

Deoxys-BC [JNPS21] is a tweakable block cipher of the TWEAKEY framework [JNP14b] based the AES round function, on which the best known attacks are based on boomerangs [CHP⁺17,Sas18,ZDJ19,ZDJM19]. Due to the large choice of tweak differences, finding the best truncated boomerang trails manually is a tedious work. Instead, we use our MILP model of section 7.

In the single tweakkey model, the analysis is exactly that of the AES, and the best known boomerang attack is given in section 4.

In the related tweakkey model, the attacker can insert differences in some of the tweakkey words TK^i . Depending on the tweak size and differences used, this can be either a single-key attack with chosen tweaks, or a related-key attack. We denote as RTK_r a model with differences in r 128-bit states, corresponding to:

- RTK1: single-key attacks on any variant with at least 128 bits of tweak.
- RTK2: single-key attacks on Deoxys-BC-384 with 256 bits of tweak, or related-key attacks on Deoxys-BC-256.
- RTK3: related-key attacks on Deoxys-BC-384.

For 13-round Deoxys-BC in the RTK3 model, we selected a non-optimal trail in terms of data complexity, which was better in time complexity. For 8-round and 9-round Deoxys-BC in the RKT1 model and 10-round Deoxys-BC in the RTK2 model, we modified the skeleton of the trail returned by the MILP because the original one was not instantiable. During other difference instantiations, we sometimes applied slight manual improvements. For instance, for minor gains, we introduced state changing bytes: fixed on the forward trail but truncated on the return trail.

Description of the attacks. In the related-tweakkey model, the attacker queries two sets of $|\mathcal{D}_{\text{in}}^0|$ plaintexts (under tweaks T and T'), and each ciphertext is shifted $|\mathcal{D}_{\text{out}}^1|$ times with a new tweak (\bar{T} and \bar{T}' respectively). In total a structure of

size $S = |\mathcal{D}_{\text{in}}^0| \times |\mathcal{D}_{\text{out}}^1|$ requires $2S$ decryption queries (or $4S$ if $|\mathcal{D}_{\text{out}}^1| = 1$) and generates S^2 quartets.

The values of ℓ and κ mentioned on the figures are the one used in our attacks, corresponding either to a 1-round or to a 2-round key-recovery. Each attack recovers a partial key, aiming for a success rate of $1/2$, comparable to previous analysis; we assume that the rest of the key can be recovered efficiently afterwards. When $\tilde{\sigma} \gg 1$, the number μ of right quartets required varies from 1 to 4. In particular, if $p_b \gg \ell \cdot p_s$, we expect no wrong quartet and $\mu = 1$ suffices, else several right quartets are needed to get the correct key ranked first.

10-round Deoxys-BC in the RTK2 model (Figure 5). Query one partial structure of $2^{93.2}$ ciphertexts, so that on average, $\mu = 2^{93.2+93.2} \cdot \tilde{p}_b = 2$ quartets follow the trail. For each element of the structure, deduce on average 1 candidate for 30 bits of key on the plaintext side: 1 candidate for $tk_0[4, 9, 14]$ and 1 representant of the 4 possible candidates each for $tk_0[3]$ ¹. In total, there are on average $2^{93.2+93.2-56-30} = 2^{100.4}$ candidate quartets matching on the ciphertext bytes with a known difference and on the key candidate.

For each quartet, retrieve 2^{-8} candidates for $tk_{10}[9]$ with 2 table accesses. This costs $2^{101.4}$ table accesses, and since an encryption makes $10 \times 16 = 2^{7.3}$ S-Box calls, this step costs $2^{94.1}$ equivalent encryptions. For each of the $2^{100.4-8} = 2^{92.4}$ remaining quartets, retrieve on average 2^{-32} candidates of $tk_{10}[0, 1, 2, 3, 4, 5, 6, 7]$. Finally, recover 2^{-16} candidates for $tk_9^{eq}[1, 6]$. There remains $2^{92.4-32-16} = 2^{44.4}$ quartets with a 118-bit key candidate. The only candidate suggested twice is expected to be the right candidate. The time complexity is $2^{94.2} + 2^{94.1} \approx 2^{95.2}$, thus $(D, T, M) = (2^{94.2}, 2^{95.2}, 2^{94.2})$.

11-round Deoxys-BC in the RTK2 model. The MILP solver did not return a pertinent trail for this key setting. Instead, we use the 10-round trail and append a round at the beginning. First, query the full encryption codebook with \bar{T}, \bar{T}' and store it. Then, guess the full tk_0 . Perform the 10-round attack, by using the same ciphertext structure for each guess of tk_0 and simulating encryption queries with fetches in the codebook. We chose $\mu = 4$ and for each key guess, the 10-round attack with partial structures of $2^{93.7}$ elements gives $2^{45.4}$ candidates for 118 bits, for a time complexity of $2^{95.1}$. If we suppose that a fetch to the codebook costs an encryption in time complexity, we end up with $T = 2^{128}(2^{94.7} + 2^{95.1}) = 2^{223.9}$. The probability that one of the counters is at least 4 is $2^{-295+116+128} = 2^{-51}$, so in average, the correct key is ranked first. This gives $(D, T, M) = (2^{129}, 2^{223.9}, 2^{129})$.

¹ S-Box 3 on the plaintext side has two pairs $(x, x + \delta), (x', x' + \delta)$ following the transition fixed by the trail. Instead of listing four key candidates, we identify one of the 2^6 cosets of $\langle \delta, x + x' \rangle$.

13-round Deoxys-BC in the RTK3 model (Figure 6). Query a partial structure of $2^{125.65}$ plaintexts. On average, $\mu = 2^{125.65} \cdot 2^{125.65} \cdot \tilde{p}_b = 4$ quartets follow the trail.

1. For each element of the structure, retrieve the representant k of the 2^6 possible key values of $tk_{13}[13, 14, 15]$ that satisfy the transition $y_{12} \rightarrow x_{12}$. k defines 18 key bits.
2. Guess the value of the tweakey material $tk_0[2, 7, 8, 13]$. Set $\delta = 0x7e42c465$ and $\delta_{in} = 0x00007a00$, and look for collisions between:

$$v = y_0[2, 7, 8, 13] \quad \| \overline{y_0}[2, 7, 8, 13] \quad \| \overline{P}[0, 5, 10, 15] \quad \| k$$

$$v' = y'_0[2, 7, 8, 13] + \delta \| \overline{y'_0}[2, 7, 8, 13] + \delta \| \overline{P'}[0, 5, 10, 15] + \delta_{in} \| k'$$

This step costs $2^{32} \cdot 2 \cdot 2^{125.65} = 2^{158.65}$ in time complexity. On average, $2^{125.65} \cdot 2^{125.65} \cdot 2^{-114} = 2^{137.3}$ quartets remain for each $tk_0[2, 7, 8, 13]$ ($2^{169.3}$ in total).

3. For each quartet, retrieve $2^{7+7-32} = 2^{-18}$ values of $tk_0[3, 4, 9, 14]$ such that the difference in $w_0[4]$ is compatible with the S-Box transition in the next round. In order to minimize the complexity, first deduce the $2^{7+7-8} = 2^6$ pairs of column differences compatible with a key candidate for $tk_0[3]$, by only checking the first S-Box. Then, deduce the $2^{6-8} = 2^{-2}$ pairs of columns compatible with a key candidate for $tk_0[4]$ with the second S-Box. Finally deduce $tk_0[9, 14]$.

This step requires $2^8 + 2^7 = 2^{8.6}$ table accesses per quartet, therefore a total of $2^{8.6+169.3} = 2^{177.9}$ accesses; and $2^{32+137.3-18} = 2^{151.3}$ quartets remain.

4. For each quartet, retrieve $2^{7+7-32} = 2^{-18}$ values of $tk_0[1, 6, 11, 12]$ and $2^{24+24-32} = 2^{16}$ key candidates for $tk_{13}[8, 9, 10, 11]$. Recover $x_{12}[8, 9, 10, 11]$ and the difference in $y_{11}[2, 7, 8]$. Retrieve 2^{-24} candidates for $tk_{12}^{eq}[2, 7, 8]$. $2^{151.3-18+16-24} = 2^{125.3}$ quartets remain.
5. For each quartet, recover the difference in $x_1[4, 14]$ and the value of $w_0[4, 14]$ from the known key bytes of tk_0 . Retrieve $2 \cdot 2 \cdot 2^{-8} = 2^{-6}$ values of $tk_1[4]$ and 2^{-6} values for $tk_1[14]$ (2 candidates are deduced per pair because the differences are already compatible). $2^{125.3} \cdot 2^{-12} = 2^{113.3}$ quartets remain.
6. Eventually, each of the $2^{113.3}$ quartets determines in average 1 candidate of $18 + 32 + 32 + 32 + 32 + 24 + 16 = 186$ bits. We model a wrong counter with a poisson distribution with $\lambda = 2^{-72.7}$. The probability that any wrong counter is at least 3 is $(1 - e^{-\lambda}(1 + \lambda + \lambda^2/2)) \cdot 2^{184} \approx 2^{-34.7}$. The correct counter follows the poisson distribution with $\lambda = 4$ and it is at least 3 with probability 0.76. Therefore, the success probability of this attack is 0.76.

Complexity analysis. The time complexity is dominated by the $2^{177.9}$ table accesses of step 3. An encryption of 13-round Deoxys-BC has 16×13 S-Boxes, so the time complexity is equivalent to $2^{177.9}/208 = 2^{170.2}$ encryptions. Thus $(D, T, M) = (2^{126.7}, 2^{170.2}, 2^{126.7})$.

Other variants. The attacks found on other variants of Deoxys-BC are detailed in Appendix D.




Fig. 5. Truncated boomerang attack on 10-round Deoxys-BC in the RTK2 model, starting from the ciphertext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with UBCT, LBCT and EBCT (probabilities on the trail).




Fig. 6. Truncated boomerang attacks on 13-round Deoxys-BC in the RTK3 model, starting from the plaintext side. This attack succeeds with probability 0.76. Middle rounds are analyzed with the ladder switch and single BCT (probability r).

9 Conclusion

In this paper, we develop a framework for truncated boomerang attacks. Instead of extending a distinguisher with additional key-recovery rounds, we integrate them inside the distinguisher. This results in a simple and generic formula for the data complexity of the attack, while the classical approach to add rounds strongly depends on the shape of input/output differences of the distinguisher. In particular, the formula can be integrated in a MILP model, leading to better results than a separate search for distinguishers and attacks.

Acknowledgement. We would like to thank the authors of [CHP⁺17] for providing the code they used to generate the MILP programs. This work was supported by the French Ministry of Defence (AID), and by the French Agence Nationale de la Recherche (ANR), under grant ANR-20-CE48-0017 (project SELECT).

References

- Bar19. Navid Ghaedi Bardeh. A key-independent distinguisher for 6-round AES in an adaptive setting. Cryptology ePrint Archive, Report 2019/945, 2019. <https://eprint.iacr.org/2019/945>.
- BDK01. Eli Biham, Orr Dunkelman, and Nathan Keller. The rectangle attack - rectangling the Serpent. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 340–357. Springer, Heidelberg, May 2001.
- BDK02. Eli Biham, Orr Dunkelman, and Nathan Keller. New results on boomerang and rectangle attacks. In Joan Daemen and Vincent Rijmen, editors, *FSE 2002*, volume 2365 of *LNCS*, pages 1–16. Springer, Heidelberg, February 2002.
- BDK⁺20. Achiya Bar-On, Orr Dunkelman, Nathan Keller, Eyal Ronen, and Adi Shamir. Improved key recovery attacks on reduced-round AES with practical data and memory complexities. *Journal of Cryptology*, 33(3):1003–1043, July 2020.
- BFL11. Charles Bouillaguet, Pierre-Alain Fouque, and Gaëtan Leurent. Security analysis of SIMD. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors, *SAC 2010*, volume 6544 of *LNCS*, pages 351–368. Springer, Heidelberg, August 2011.
- BGGS20. Zhenzhen Bao, Chun Guo, Jian Guo, and Ling Song. TNT: How to tweak a block cipher. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 641–673. Springer, Heidelberg, May 2020.
- BGL20. Zhenzhen Bao, Jian Guo, and Eik List. Extended truncated-differential distinguishers on round-reduced AES. *IACR Trans. Symm. Cryptol.*, 2020(3):197–261, 2020.
- Bir04. Alex Biryukov. The boomerang attack on 5 and 6-round reduced aes. In *International Conference on Advanced Encryption Standard*, pages 11–15. Springer, 2004.

- BK09. Alex Biryukov and Dmitry Khovratovich. Related-key cryptanalysis of the full AES-192 and AES-256. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 1–18. Springer, Heidelberg, December 2009.
- BL23. Augustin Bariant and Gaëtan Leurent. Truncated boomerang attacks and application to AES-based ciphers — Additional data. https://github.com/AugustinBariant/Truncated_boomerangs, 2023.
- BLNS18. Christina Boura, Virginie Lallemand, María Naya-Plasencia, and Valentin Suder. Making the impossible possible. *Journal of Cryptology*, 31(1):101–133, January 2018.
- BR19. Navid Ghaedi Bardeh and Sondre Rønjom. The exchange attack: How to distinguish six rounds of AES with $2^{88.2}$ chosen plaintexts. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 347–370. Springer, Heidelberg, December 2019.
- CHP⁺17. Carlos Cid, Tao Huang, Thomas Peyrin, Yu Sasaki, and Ling Song. A security analysis of Deoxys and its internal tweakable block ciphers. *IACR Trans. Symm. Cryptol.*, 2017(3):73–107, 2017.
- CHP⁺18. Carlos Cid, Tao Huang, Thomas Peyrin, Yu Sasaki, and Ling Song. Boomerang connectivity table: A new cryptanalysis tool. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 683–714. Springer, Heidelberg, April / May 2018.
- DDV20. Stéphanie Delaune, Patrick Derbez, and Mathieu Vavrille. Catching the fastest boomerangs application to SKINNY. *IACR Trans. Symm. Cryptol.*, 2020(4):104–129, 2020.
- DEM16. Christoph Dobraunig, Maria Eichlseder, and Florian Mendel. Square attack on 7-round kiasu-BC. In Mark Manulis, Ahmad-Reza Sadeghi, and Steve Schneider, editors, *ACNS 16*, volume 9696 of *LNCS*, pages 500–517. Springer, Heidelberg, June 2016.
- Der13. Patrick Derbez. *Meet-in-the-Middle Attacks on AES*. Theses, Ecole Normale Supérieure de Paris - ENS Paris, December 2013.
- DKR97. Joan Daemen, Lars R. Knudsen, and Vincent Rijmen. The block cipher Square. In Eli Biham, editor, *FSE'97*, volume 1267 of *LNCS*, pages 149–165. Springer, Heidelberg, January 1997.
- DKRS20. Orr Dunkelman, Nathan Keller, Eyal Ronen, and Adi Shamir. The retracing boomerang attack. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 280–309. Springer, Heidelberg, May 2020.
- DKS10. Orr Dunkelman, Nathan Keller, and Adi Shamir. A practical-time related-key attack on the KASUMI cryptosystem used in GSM and 3G telephony. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 393–410. Springer, Heidelberg, August 2010.
- DL17. Christoph Dobraunig and Eik List. Impossible-differential and boomerang cryptanalysis of round-reduced kiasu-BC. In Helena Handschuh, editor, *CT-RSA 2017*, volume 10159 of *LNCS*, pages 207–222. Springer, Heidelberg, February 2017.
- DR02. Joan Daemen and Vincent Rijmen. *The design of Rijndael*, volume 2. Springer, 2002.
- FKL⁺01. Niels Ferguson, John Kelsey, Stefan Lucks, Bruce Schneier, Michael Stay, David Wagner, and Doug Whiting. Improved cryptanalysis of Rijndael. In Bruce Schneier, editor, *FSE 2000*, volume 1978 of *LNCS*, pages 213–230. Springer, Heidelberg, April 2001.

- GGLS20. Chun Guo, Jian Guo, Eik List, and Ling Song. Towards closing the security gap of tweak-aNd-tweak (TNT). In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part I*, volume 12491 of *LNCS*, pages 567–597. Springer, Heidelberg, December 2020.
- Gra18. Lorenzo Grassi. MixColumns properties and attacks on (round-reduced) AES with a single secret S-box. In Nigel P. Smart, editor, *CT-RSA 2018*, volume 10808 of *LNCS*, pages 243–263. Springer, Heidelberg, April 2018.
- GRR16. Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. Subspace trail cryptanalysis and its applications to AES. *IACR Trans. Symm. Cryptol.*, 2016(2):192–225, 2016. <https://tosc.iacr.org/index.php/ToSC/article/view/571>.
- GRR17. Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. A new structural-differential property of 5-round AES. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part II*, volume 10211 of *LNCS*, pages 289–317. Springer, Heidelberg, April / May 2017.
- JNP14a. Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Kiasu v1. *Submitted to the CAESAR competition*, 2014.
- JNP14b. Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Tweaks and keys for block ciphers: The TWEAKEY framework. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 274–288. Springer, Heidelberg, December 2014.
- JNPS21. Jérémy Jean, Ivica Nikolic, Thomas Peyrin, and Yannick Seurin. The deoxys AEAD family. *Journal of Cryptology*, 34(3):31, July 2021.
- KKS01. John Kelsey, Tadayoshi Kohno, and Bruce Schneier. Amplified boomerang attacks against reduced-round MARS and Serpent. In Bruce Schneier, editor, *FSE 2000*, volume 1978 of *LNCS*, pages 75–93. Springer, Heidelberg, April 2001.
- MS02. Itsik Mantin and Adi Shamir. A practical attack on broadcast RC4. In Mitsuru Matsui, editor, *FSE 2001*, volume 2355 of *LNCS*, pages 152–164. Springer, Heidelberg, April 2002.
- Mur11. Sean Murphy. The return of the cryptographic boomerang. *IEEE Trans. Inf. Theory*, 57(4):2517–2521, 2011.
- MWGP11. Nicky Mouha, Qingju Wang, Dawu Gu, and Bart Preneel. Differential and linear cryptanalysis using mixed-integer linear programming. In *Inscrypt*, volume 7537 of *Lecture Notes in Computer Science*, pages 57–76. Springer, 2011.
- QDW⁺21. Lingyue Qin, Xiaoyang Dong, Xiaoyun Wang, Keting Jia, and Yunwen Liu. Automated search oriented to key recovery on ciphers with linear key schedule. *IACR Trans. Symm. Cryptol.*, 2021(2):249–291, 2021.
- RBH17. Sondre Rønjom, Navid Ghaedi Bardeh, and Tor Hellesest. Yoyo tricks with AES. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 217–243. Springer, Heidelberg, December 2017.
- RSP21. Mostafizar Rahman, Dhiman Saha, and Goutam Paul. Boomeyong: Embedding yoyo within boomerang and its applications to key recovery attacks on AES and Pholkos. *IACR Trans. Symm. Cryptol.*, 2021(3):137–169, 2021.
- Sas18. Yu Sasaki. Improved related-tweakey boomerang attacks on deoxys-BC. In Antoine Joux, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 18*, volume 10831 of *LNCS*, pages 87–106. Springer, Heidelberg, May 2018.

- Sel08. Ali Aydin Selçuk. On probability of success in linear and differential cryptanalysis. *Journal of Cryptology*, 21(1):131–147, January 2008.
- SQH19. Ling Song, Xianrui Qin, and Lei Hu. Boomerang connectivity table revisited. *IACR Trans. Symm. Cryptol.*, 2019(1):118–141, 2019.
- TAY16. Mohamed Tolba, Ahmed Abdelkhalek, and Amr M Youssef. A meet in the middle attack on reduced round kiasu-bc. *IEICE TRANSACTIONS on Fundamentals of Electronics, Communications and Computer Sciences*, 99(10):1888–1890, 2016.
- TKKL15. Tyge Tiessen, Lars R. Knudsen, Stefan Kölbl, and Martin M. Lauridsen. Security of the AES with a secret S-box. In Gregor Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 175–189. Springer, Heidelberg, March 2015.
- Wag99. David Wagner. The boomerang attack. In Lars R. Knudsen, editor, *FSE’99*, volume 1636 of *LNCS*, pages 156–170. Springer, Heidelberg, March 1999.
- WP19. Haoyang Wang and Thomas Peyrin. Boomerang switch in multiple rounds. *IACR Trans. Symm. Cryptol.*, 2019(1):142–169, 2019.
- YSS⁺22. Qianqian Yang, Ling Song, Siwei Sun, Danping Shi, and Lei Hu. New properties of the double boomerang connectivity table. *IACR Transactions on Symmetric Cryptology*, 2022(4):208–242, Dec. 2022.
- ZDJ19. Boxin Zhao, Xiaoyang Dong, and Keting Jia. New related-tweakey boomerang and rectangle attacks on deoxys-bc including BDT effect. *IACR Trans. Symm. Cryptol.*, 2019(3):121–151, 2019.
- ZDJM19. Boxin Zhao, Xiaoyang Dong, Keting Jia, and Willi Meier. Improved related-tweakey rectangle attacks on reduced-round Deoxys-BC-384 and Deoxys-I-256-128. In Feng Hao, Sushmita Ruj, and Sourav Sen Gupta, editors, *INDOCRYPT 2019*, volume 11898 of *LNCS*, pages 139–159. Springer, Heidelberg, December 2019.

A Optimizing the TNT-AES Boomerang

We can slightly reduce the complexity of the TNT-AES distinguisher by carefully analyzing the probability of the boomerang trail, and correlations between the sides.

In the lower trail, we have two pairs (C, \bar{C}) and (C', \bar{C}') and our analysis assumes that if they both follow the trail, then the differences $E_1^{-1}(C) + E_1^{-1}(\bar{C})$ and $E_1^{-1}(C') + E_1^{-1}(\bar{C}')$ are equal with probability $|\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-32}$. Actually, the differences are not uniformly distributed in $\mathcal{D}_{\text{in}}^1$, and this increases the probability that the differences are equal.

Indeed, we can split this analysis into two disjoint cases: either the differences in x_4 are equal, or they are different. If they are equal, then the 4 active S-Boxes in the fourth round have the same difference in y_3 for both pairs; therefore there are only 127 possible differences in x_3 .

$$\begin{aligned} & \Pr \left[E_1^{-1}(C) + E_1^{-1}(\bar{C}) = E_1^{-1}(C') + E_1^{-1}(\bar{C}') \right] \\ &= \tilde{q}^2 \times \left(\frac{254}{255} \times 2^{-32} + \frac{1}{255} \times 127^{-4} \right) \\ &\approx \tilde{q}^2 \times 2^{-31.915} \end{aligned}$$

In the upper trail, there are similar effects. Our analysis assumes that the pair (\bar{P}, \bar{P}') follows the truncated trail with probability \tilde{p} independently of the pair (P, P') . However, both pairs have the same differences at the input and output of the trail and the trail does not cover the sets $\mathcal{D}_{\text{in}}^0$ and $\mathcal{D}_{\text{out}}^0$ uniformly. Let us consider a pair \bar{C}, \bar{C}' with $E_1^{-1}(\bar{C}) + E_1^{-1}(\bar{C}') = E_1^{-1}(C) + E_1^{-1}(C')$. The difference in y_2 (after the S-Boxes of the third round) are the same for both pairs; the truncated trail allows a set 255 differences in x_2 (before the S-Boxes). In general the probability of a transition through four active S-Boxes is 2^{-32} , but we know that one of the 255 differences was followed by the pair P, P' , therefore all differences are compatible, and the probability increases to 127^{-4} . Finally the probability of having a difference in x_2 compatible with the trail is higher than 2^{-24} .

$$1 \times 127^{-4} + 254 \times 2^{-32} \approx 2^{-23.92}$$

In the first round, we have the same analysis as in the lower trail, and the probability to obtain the same difference as the pair P, P' is higher than 2^{-32}

$$\frac{254}{255} \times 2^{-32} + \frac{1}{255} \times 127^{-4} \approx 2^{-31.915}$$

Finally, we obtain

$$q_b = 2^{-22} \times 2^{-23.92} \cdot 2^{-31.915} \times 2^{-20.85 \times 2} \times 2^{-31.915} = 2^{-151.45}$$

Multiple trails. The previous analysis uses a single trail for the lower part of the cipher, with the active byte of z_4 in the cell 0. We can also consider




Fig. 7. Scheme of our boomerang attack on the full TNT-AES. Although this only shows a unique boomerang trail which would imply $p_b = 2^{-151.7}$, the attack parameters are shown for $p_b = 2^{-151.4}$, following the multi-trail analysis of Appendix A.

alternative trails in cell 1, 2, or 3. With this active byte in cell 1, we have the transition $\Delta_{\text{out}}^1 \rightarrow (0, *, 0, 0)$ through the layer of S-Boxes with probability $\bar{q} = 256/2^{32} = 2^{-24}$ (instead of $2^{-20.85}$ for cell 0), and the probability is the same in other positions.

Finally, we sum the probabilities of the four alternate boomerangs, and we obtain:

$$q_b = 2^{-151.45} + 2^{-157.75} + 2^{-157.75} + 2^{-157.75} \approx 2^{-151.4}$$

A.1 Choice of Δ_{out}^1

We have to choose a value Δ_{out}^1 to maximize the probability of having a single active byte in z_4 . Therefore we analyze the S-Box layer of the last round: the output difference is equal to $\text{ShiftRows}^{-1}(\text{MixColumns}^{-1}(\Delta_{\text{out}}^1))$, and we want the input difference to be of one of the following types:

$$\begin{bmatrix} 2\alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ 3\alpha & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} \alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ 3\alpha & 0 & 0 & 0 \\ 2\alpha & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} \alpha & 0 & 0 & 0 \\ 3\alpha & 0 & 0 & 0 \\ 2\alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 3\alpha & 0 & 0 & 0 \\ 2\alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \\ \alpha & 0 & 0 & 0 \end{bmatrix}$$

In particular, $\text{ShiftRows}^{-1}(\text{MixColumns}^{-1}(\Delta_{\text{out}}^1))$ must be active only on the first column. We experimentally tested all possible differences on the first column, and counted the number of pairs satisfying the transition. More precisely, we are interested in the joint probability that two different pairs reach an input difference of the same type, therefore we count the number of quartets of each type (using the DDT of the S-Box for each possible input and output difference).

We found that in the best case, there are $2^{22.35}$ quartets that satisfy this transition (with the same type for both pairs), compared to 2^{18} quartets expected on average (for each type of input difference, we expect on average one pair for each of the 2^8 differences, therefore 2^{8+8} quartets). There are 4×255 choices of Δ_{out}^1 reaching this maximum, and we found they all have a special form: $\text{ShiftRows}^{-1}(\text{MixColumns}^{-1}(\Delta_{\text{out}}^1))$ is of one of the following types, with L the linear transform inside the AES S-Box:

$$\begin{bmatrix} L(\beta/2) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \\ L(\beta/3) & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} L(\beta) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \\ L(\beta/3) & 0 & 0 & 0 \\ L(\beta/2) & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} L(\beta) & 0 & 0 & 0 \\ L(\beta/3) & 0 & 0 & 0 \\ L(\beta/2) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} L(\beta/3) & 0 & 0 & 0 \\ L(\beta/2) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \\ L(\beta) & 0 & 0 & 0 \end{bmatrix}$$

This special form can be explained by the structure of the AES S-Box; it is defined as $x \mapsto L(x^{254})$, with L an invertible linear mapping. It implies the following property:

Property 1. For any $\alpha, \beta \neq 0$, $\text{DDT}[\alpha, L(\beta)] = \text{DDT}[\alpha\beta, L(1)]$.

Proof.

$$\begin{aligned}
\text{DDT}[\alpha, L(\beta)] &= |\{x : L((x)^{254}) + L((x + \alpha)^{254}) = L(\beta)\}| \\
&= |\{x : (x)^{254} + (x + \alpha)^{254} = \beta\}| \\
&= |\{x : (\beta x)^{254} + (\beta x + \alpha\beta)^{254} = 1\}| \\
&= |\{x : (x)^{254} + (x + \alpha\beta)^{254} = 1\}| \\
&= |\{x : L((x)^{254}) + L((x + \alpha\beta)^{254}) = L(1)\}| \\
&= \text{DDT}[\alpha\beta, L(1)] \quad \square
\end{aligned}$$

Because of this property, S-Box transitions of the form

$$(2\alpha, \alpha, \alpha, 3\alpha) \rightarrow (L(\beta/2), L(\beta), L(\beta), L(\beta/3))$$

have a higher probability than expected because the transition probability of the S-Boxes are not independent: for a given α and β either all transitions are possible simultaneously, or none of the transitions are possible.

For instance, with $\beta = 1$, we obtain the following transitions though an inverse layer of S-Boxes followed by an inverse MixColumns:

$$\begin{aligned}
\Pr[(14, 1f, 1f, 18) \rightarrow (*, 0, 0, 0)] &= 2272/2^{32} \approx 2^{-20.85} \\
\Pr[(14, 1f, 1f, 18) \rightarrow (0, *, 0, 0)] &= 256/2^{32} \\
\Pr[(14, 1f, 1f, 18) \rightarrow (0, 0, *, 0)] &= 256/2^{32} \\
\Pr[(14, 1f, 1f, 18) \rightarrow (0, 0, 0, *)] &= 256/2^{32}
\end{aligned}$$

The probability that two pairs follow the same trail is therefore:

$$(2272/2^{32})^2 + 3 \times (256/2^{32})^2 \approx 2^{22.35}/2^{64}$$


B Attacking TNT with 5-round AES

If TNT is used with 5-round AES, we can build efficient boomerang attacks with a probability-1 truncated differential in one part.

B.1 Distinguisher

For a distinguisher, we split E_{K_1} into a 2-round E_0 and 3-round E_1 .

Upper Differential. We use a 2-round truncated trail with probability 1:



We have the following parameters

$$\vec{p} = 1 \quad \bar{p} = 1 \quad |\mathcal{D}_{\text{in}}^0| = 2^{32} \quad |\mathcal{D}_{\text{out}}^0| = 2^{32}$$

For the return trail, we must hit a fixed $\bar{T} + \overline{T^v} = \Delta_{\text{in}}^0$:

$$\vec{\bar{p}} = 1 \quad \bar{\bar{p}} = 2^{-32} \quad |\overline{\mathcal{D}}_{\text{in}}^0| = 2^0 \quad |\overline{\mathcal{D}}_{\text{out}}^0| = 2^{32}$$

Lower Differential. We use the same trail as in the 6-round attack, with a fixed difference $\delta = (L(\beta/2), L(\beta), L(\beta), L(\beta/3))$:

$$\vec{q} = 2^{-52.85} \quad \bar{q} = 2^{-20.85} \quad |\mathcal{D}_{\text{in}}^1| = 2^{32} \quad |\mathcal{D}_{\text{out}}^1| = 2^0$$

Boomerang Trail Probability. We obtain:


$$p_b = \vec{p} \cdot \bar{p} \cdot \bar{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-105.7} \quad p_{\S} = |\bar{\mathcal{D}}_{\text{in}}^0|/2^n = 2^{-128} \quad \sigma = 2^{22.3}$$

With $\sigma \gg 1$, we don't expect any false positive with $\bar{P} = \bar{P}'$, and we obtain a simple distinguisher with $Q = \mu \times p_b^{-1}$ for a small μ . For instance, with $Q = 2^{107}$, corresponding to $D = 2^{76}$, we have a success rate of 0.91.

B.2 Key recovery

For a key-recovery attack, we use the same ideas as in the 6-round attack, in order to recover $E_{K_0}(P) + K_1$. This provides known plaintext/ciphertext pairs for E_{K_0} , and we use a low data complexity attack on 5-round AES to recover the master key. However, we have to split E_{K_1} into a 3-round E_0 and 2-round E_1 in order to extract information about $E_{K_0}(P) + K_1$.

Upper Differential. We use the same collection of 4 upper trails as in the 6-round attack:




We have the following parameters

$$\vec{p} = 2^{-22} \quad \bar{p} = 2^{-24} \quad |\mathcal{D}_{\text{in}}^0| = 2^{32} \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

For the return trail, we must hit a fixed $\bar{T} + \bar{T}' = \Delta_{\text{in}}^0$:

$$\vec{p} = 2^{-22} \quad \bar{p} = 2^{-56} \quad |\bar{\mathcal{D}}_{\text{in}}^0| = 2^0 \quad |\mathcal{D}_{\text{out}}^0| = 2^{34}$$

Lower Differential. We use a 2-round truncated trail with reverse probability 1, with a fixed (arbitrary) output difference:



We have the following parameters

$$\vec{q} = 2^{-32} \quad \bar{q} = 1 \quad |\mathcal{D}_{\text{in}}^1| = 2^{32} \quad |\mathcal{D}_{\text{out}}^1| = 2^0$$

Using a single active output byte ensures that we don't have quartets that follow an alternative boomerang with a 2-round trail in the upper part with probability 1, and a 3-round trail in the lower part.

Boomerang Trail Probability. We obtain:

$$p_b = \vec{p} \cdot \vec{p} \cdot \vec{q}^2 \times |\mathcal{D}_{\text{in}}^1|^{-1} = 2^{-110} \quad p_s = |\overline{\mathcal{D}}_{\text{in}}^0|/2^n = 2^{-128} \quad \sigma = 2^{18}$$

Key recovery. As in the 6-round attack, we recover $\ell = 2^{10}$ candidates for the main diagonal of $E_{K_0}(P) + K_1$ from each candidate quartet ($\kappa = 32$ bits of information). If we have at least two candidate quartets, the correct equivalent key is ranked first with high probability. Therefore we use $Q = 2^{113}$, so that the probability to have at least two candidate quartets is higher than 0.99. This corresponds to an attack with $D = 2^{82}$.

In order to recover the master key, we repeat the boomerang attack on the four diagonals to recover the full value of $E_{K_0}(P) + K_1$. Therefore, we obtain a known plaintext/ciphertext pair for E_{K_0} , and we use a low data complexity attack on 5-round AES. Using the attack from [Der13], we only need 8 chosen plaintexts to attack 5-round AES with a complexity of 2^{64} . Therefore, the total attack has a complexity of $8 \times 4 \times 2^{82} = 2^{87}$.

C Description of the MILP Model

On a high level, the model is given by a set of variables, a set of constraints, and an objective function. Our model is not symmetric when switching plaintext and ciphertext², but for simplicity we only describe the attack starting from the plaintext, though it works similarly the other way around.

State Variables. The model of Cid *et al.* considers only two types of internal differences: either it is inactive with a zero difference, or it is active with a fixed non-zero difference. In order to model truncated trails, we consider four types of differences for all intermediate state variables (for the tweakable variables, we use only the first two types for simplicity):

- inactive, with a zero difference, denoted as \square ;
- active with a *fixed* non-zero difference, denoted as \blacksquare ;
- active with an *unknown* (truncated) difference, denoted as \blacksquare .
- active with an *equal* (but unknown) difference for both pairs, denoted as \blacksquare .

Equal Differences. Bytes with *equal* differences encode relations between the two different pairs that follow the same trail, rather than properties of a trail by itself. This allows the MILP model to capture trails like the 6-round AES boomerang of Figure 3 or the 8-round Kiasu boomerang of Figure 4; the model does not encode linear relations between active bytes (*e.g.* the set of differences for w_2 of Figure 3 is active on all bytes but has size 2^{32}), but using this type of

² The inverse AES round can be re-written with the same form as the encryption, but this requires to use equivalent round keys, and interacts badly with the introduction of sparse differences through the tweakable schedule.

constraint is sufficient in many cases because it is propagated through the linear layer. Figure 8 depicts the boomerang trail on 6-round AES as interpreted by the MILP model. In terms of analysis, we treat them specially: for the first pair they are considered as truncated bytes, but for the second pair they are considered as fixed differences (fixed to the value given by the first pair). Therefore, we explain the rest of the MILP model assuming that each trail has been duplicated, and *equal* differences have been replaced.

Constraints. We have constraints for each operation:

SubBytes. With truncated differences, the variables before and after the S-Box layer do not necessarily have the same type. However, an S-Box output is active (truncated or not) if and only if the input is active.

ShiftRows. ShiftRows only moves the bytes, so we have trivial equalities between the corresponding state variables.

MixColumns. The MixColumns operation operates on each column, multiplying it with an MDS matrix. Because of the MDS property, each column is either completely inactive, or has at least 5 active bytes (truncated or not) on input and output. Moreover, we have the same property with truncated bytes: either no byte is truncated, or at least 5 bytes are truncated.

For Deoxys-BC, we also reuse the constraints for the degree of freedom given in [CHP⁺17] to avoid over-defining the fixed differences (for these constraints, the *equal* bytes are considered as truncated).

AddTweakKey. The AddTweakKey operation is just a XOR with the subkey. In our model, the subkey is not truncated, so it is either inactive, or active with a fixed difference. Therefore, the input state is truncated if and only if the output state is truncated. Otherwise, we use the constraints of Cid *et al.* to model XOR for the active bytes and the activity of the tweakkey bytes.

Key Schedule. For Deoxys-BC, we follow the approach of [CHP⁺17] to model the key schedule. The permutation h just moves the bytes, and the LFSR construction ensures that each byte is either completely inactive, or inactive in at most $i - 1$ rounds in the TK_i model. For Kiasu-BC, the 64-bit tweak is represented with 8 activity variables. For AES-128, the subkeys are inactive.

Additional Constraint. To avoid the silly truncated trail with probability 1, we constrain the trail to have at maximum 3 active truncated columns in a single state, except for the first and last rounds.

Objective Function. Using the results from the previous sections, we estimate the data complexity of an attack as:

$$\begin{aligned}
 D &= \max(\sqrt{2Q}, 2Q \times |\mathcal{D}_{\text{in}}^0|^{-1} \times |\mathcal{D}_{\text{out}}^1|^{-1}) && \text{with} \\
 Q &\approx \max(p_b^{-1}, \tilde{\sigma}^{-1} \times p_b^{-1}) && \tilde{\sigma} \approx p_b / (p_{\S} \times \ell \times 2^{-\kappa}) \\
 p_b &= \vec{p} \cdot \vec{p} \cdot \vec{q}_1 \cdot \vec{q}_2 \times r && p_{\S} = |\mathcal{D}_{\text{in}}^0| / 2^n
 \end{aligned}$$

Since all variables are represented logarithmically by the MILP model, these formulas only involve additions and maximums, and are easily expressed with the MILP variables:

Table 4. Transition probability (DDT) and connection probability (BCT) for the AES S-Box. For the transition probability of *equal* differences, we distinguish the first and the second pair. We omit the cases where the probability is 0.

Legend	Transition probability		1st pair	2nd pair
\square $\delta = 0$	$\Pr(\square \rightarrow \square) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{1}{=} 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{2}{=} 2^{-7}$	
\blacksquare $\delta \neq 0$ fixed	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-6}$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{1}{=} 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{2}{=} 2^{-8}$	
\blacksquare unknown δ	$\Pr(\blacksquare \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{1}{=} 2^{-8}$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{2}{=} 2^{-7}$	
\blacksquare <i>equal</i> δ	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{1}{=} 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{2}{=} 1$	
	$\Pr(\blacksquare \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{1}{=} 1$	$\Pr(\blacksquare \rightarrow \blacksquare) \stackrel{2}{=} 2^{-7}$	
Boomerang connection probability				
$\Pr(\square \rightarrow \square) = 1$	$\Pr(\blacksquare \rightarrow \square) = 1$	$\Pr(\blacksquare \rightarrow \square) = 1$	$\Pr(\blacksquare \rightarrow \square) = 1$	
$\Pr(\square \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-6}$	$\Pr(\blacksquare \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	
$\Pr(\square \rightarrow \blacksquare) = 2^{-8}$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	$\Pr(\blacksquare \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	
$\Pr(\square \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	$\Pr(\blacksquare \rightarrow \blacksquare) = 1$	$\Pr(\blacksquare \rightarrow \blacksquare) = 2^{-8}$	

- $|\mathcal{D}_{\text{in}}^0|$ and $|\mathcal{D}_{\text{out}}^1|$ are obtained by counting truncated input/output bytes;
- \bar{p} , \tilde{p} , \tilde{q}_1 and \tilde{q}_2 are computed from S-Box and MixColumns transitions. Due to the *equal* bytes, the probability of the lower trail is computed twice: \tilde{q}_1 after replacing them with truncated bytes, and \tilde{q}_2 after replacing them with fixed bytes. Similarly \bar{p} is computed after replacing them with truncated bytes, and \tilde{p} after replacing them with fixed bytes (see Table 4);
- r is computed from the boomerang connection probability of each S-Box;
- $\ell \times 2^{-\kappa}$ is estimated as $\bar{p}_0^2 \cdot \tilde{q}_0^2$ following section 3.2, where the probability \bar{p}_0 for the first round and \tilde{q}_0 for the last round are evaluated from the trail.

Trail Probability. With truncated differences, both MixColumns and SubBytes operations contribute to the probability of a trail. If at least one byte of a column is truncated, the probability of a MixColumns transition is 2^{-8t} where t is the number of non truncated bytes (active or not) in the output of MixColumns.

The SubBytes probability is computed by multiplying the transition probability of all the S-Boxes, following the transition probabilities given in Table 4.

Boomerang Connection Probability. In section 3, we evaluated the connection probability as $r = |\mathcal{D}_{\text{in}}^1|^{-1}$. This is sufficient for AES boomerangs, but when targeting Deoxys we obtain better results using the Boomerang Connectivity Table (BCT) [CHP⁺18]. Instead of splitting the cipher in two parts $E = E_1 \circ E_0$, we split it in three parts $E = E_1 \circ E_m \circ E_0$ with E_m an S-Box layer. Given fixed differences on both sides of E_m , the probability that the boomerang connects is given by multiplying the BCT probabilities of each S-Box. Because the MILP does not handle fixed difference instantiation, our MILP program uses the optimistic probabilities given in Table 4.

This analysis can be improved using the ladder switch [BK09]. Instead of splitting the cipher with E_m a full S-Box layer, we use the Super-Box representation of the two middle rounds: from one S-Box layer to the next, the AES round operates as four independent parallel transformations. Each of those four transformations can independently be considered as part of the upper or lower trail. We obtain E_0 and E_1 with partial rounds in the middle, and E_m corresponds to S-Boxes of different rounds. We model the ladder switch using a binary variable for each Super-Box in the middle, encoding whether it is part of the upper trail or the lower trail. Depending on these variables, the S-Boxes and MDS matrices in the middle are counted as part of E_0 , E_m , or E_1 .

Resource Consumption. For the search, we use gurobi with 96 threads on a machine with two AMD EPIC 7352 CPUs and 256 GB of RAM. The solve time varies from several minutes to several days, but the best trail is in general found much faster than the optimality proof. Memory shortage (exceeding 256 GB) is an issue for some of the largest models.

D Truncated Boomerang Attacks from the MILP Model

8-round Deoxys-BC in the RTK1 model (Figure 9). Query 2^{15} structures of 2^{72} elements. On average, $\mu = 2^{15} \cdot 2^{72} \cdot 2^{72} \cdot \tilde{p}_b = 2$ quartets follow the trail and $2^{15} \cdot 2^{72} \cdot 2^{72} \cdot p_s = 2^{79}$ random quartets are detected. For each quartet, retrieve in average 2^{-16} values of $tk_8[4, 6]$ in a few table accesses. $2^{79} \cdot 2^{-16} = 2^{63}$ quartets remain. For each remaining quartet, retrieve 2^{-24} values of $tk_0[1, 6, 12]$. For each of the $2^{63} \times 2^{-24} = 2^{39}$ remaining quartets, deduce in average 2^{-16} key candidate for $tk_8[0, 1, 2, 3]$ and $tk_7^{eq}[1, 3]$. For each of the 2^{23} remaining quartets, increase the counter of the retrieved 88-bit key candidate. Only the right counter is expected to be greater than 2. The complexity is dominated by the encryption queries. This gives $(D, T, M) = (2^{88}, 2^{88}, 2^{73})$.

9-round Deoxys-BC in the RTK1 model (Figure 10). Query 2^6 full structures of 2^{128} plaintexts. Each structure is built with a different tweak. If we are in the related key model rather than the related tweak model, we ask for structures under different linked keys. On average, $\mu = 2^6 \cdot 2^{128} \cdot 2^{128} \cdot \tilde{p}_b = 4$ quartets follow the trail. For each element of the structure, deduce on average 1 candidate k for 38 bits of key on the ciphertext side: 1 candidate for $tk_9[0, 1, 2, 4]$ and 1 representant of the 4 possible candidates for $tk_9[6]$.

Guess 5 key bytes: $tk_0[0, 5, 10, 15]$ and $tk_1[1]$, so that $w_0[0, 1, 2, 3]$ and $y_1[1]$ can be evaluated from each plaintext. Set $\delta = 0x3b000000\|0x000000\|0x000000\|0x80\|0x80\|0$ and look for collisions between:

$$\begin{aligned} v &= (\overline{P}[2, 7, 8, 13] \parallel w_0[0, 2, 3] \parallel \overline{w}_0[0, 2, 3] \parallel y_1[1] \parallel \overline{y}_1[1] \parallel k) \\ v' &= (\overline{P}'[2, 7, 8, 13] \parallel w'_0[0, 2, 3] \parallel \overline{w}'_0[0, 2, 3] \parallel y'_1[1] \parallel \overline{y}'_1[1] \parallel k') + \delta \end{aligned}$$

Since we match on 134 bits, we expect $2^{6+128+128-134} = 2^{128}$ remaining quartets for each guess, or 2^{168} in total, with a complexity of $2^{6+128+40} = 2^{174}$.

Extract more key information from the remaining quartets. First, retrieve 2^{-16} candidates for $tk_0[3, 4, 9, 14]$; this requires about 2^9 table accesses per quartet, which is less than to 2^6 encryptions, and therefore does not dominate the time complexity; 2^{152} quartet remain. Then retrieve 2^{-16} candidates for $tk_0[1, 6, 11, 12]$, and 2^{-16} candidates for $tk_1[6, 12]$. 2^{120} candidate quartets remain. For each remaining quartet, for all possible $tk_9[3]$, deduce 2^{-6} representant of the 4 possible $tk_8^{eq}[1]$ and 2^{-7} representant for the 2 possible $tk_8^{eq}[3]$. This step allows to deduce $2^{8-6-7} = 2^{-5}$ 21-bit key candidates. We end up with 2^{115} candidates for 179 bits of key.

We model the counters for wrong keys as a Poisson distribution with $\lambda = 2^{-64}$; they reach 4 or more with probability $2^{-260.6}$, therefore the right key is expected to be ranked first. This gives $(D, T, M) = (2^{135}, 2^{174}, 2^{129})$.

8-round Deoxys-BC in the RTK2 model (Figure 11). Query a partial structure of 2^{25} ciphertexts. The only detected quartet is a right quartet. Thus, $(D, T, M) = (2^{27}, 2^{27}, 2^{27})$.

9-round Deoxys-BC in the RTK2 model (Figure 12). Query a partial structure of $2^{54.2}$ ciphertexts. On average, $\mu = 2^{54.2} \cdot 2^{54.2} \cdot \tilde{p}_b = 1$ quartets follow the trail and $2^{54.2} \cdot 2^{54.2} \cdot p_s = 2^{20.4}$ random quartets are detected. For each quartet, retrieve 2^{-52} values of $tk_0[0, 5, 10, 15]$ and $tk_9[0, 1, 2, 3, 5]$. This step is of negligible complexity, and with high probability no wrong quartet remains. Thus, $(D, T, M) = (2^{55.2}, 2^{55.2}, 2^{55.2})$

10-round Deoxys-BC in the RTK3 model (Figure 13). Query $2^{1.4}$ structures of 2^{16} ciphertexts. The only detected quartet is a right quartet. $(D, T, M) = (2^{19.4}, 2^{19.4}, 2^{18})$. This attack is equivalent to the attack given in [Sas18], but the complexity was wrongly estimated as 2^{22} .

11-round Deoxys-BC in the RTK3 model (Figure 14). Query a partial structure of $2^{30.7}$ elements. The only detected quartet is a right quartet. $(D, T, M) = (2^{32.7}, 2^{32.7}, 2^{32.7})$.

12-round Deoxys-BC in the RTK3 model (Figure 15). Query $2^{2.4}$ structures of 2^{64} ciphertexts. On average, $\mu = 2^{2.4} \cdot 2^{64} \cdot 2^{64} \cdot \tilde{p}_b = 2$ quartets follow the trail and $2^{2.4} \cdot 2^{64} \cdot 2^{64} \cdot p_s = 2^{58.4}$ random quartets are detected. For each quartet, retrieve in average 2^{-32} key candidates for $tk_0[4]$ and $tk_{12}[0, 2, 3]$ in a few table accesses. Then, for each of the $2^{58.4} \cdot 2^{-32} = 2^{26.4}$ remaining quartets, deduce in average $2^{24} \cdot 2^{24} \cdot 2^{-32} = 2^{16}$ candidates for $tk_{12}[12, 13, 14, 15]$. For each candidate, compute the values of $x_{11}[12, 13, 14, 15]$ and the differences in state y_{10} . From the transition $x_{10} \rightarrow y_{10}$, retrieve 2^{-24} key candidates for $tk_{11}^{eq}[1, 11, 12]$.

Thus, $2^{26.4} \cdot 2^{16} \cdot 2^{-24} = 2^{18.4}$ quartets remain with 1 average key candidate of 88 bits. For each remaining quartet, increase the counter of the corresponding key candidate. Only the right counter is expected to be greater than 2. This gives $(D, T, M) = (2^{67.4}, 2^{67.4}, 2^{65})$.

14-round Deoxys-BC in the RTK3 model (Figure 16). We did not manage to find a 14-round trail with the MILP solver, but the 13-round attack can be extended by adding a round at the end, and guessing most of the last subkey.

We start by querying the decryption oracle over the full codebook with tweaks \bar{T} and \bar{T}' , and storing the results in memory. Then we guess 104 bits of the last round key: $k_{14}[0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12]$.

This allows us to essentially simulate the 13-round attack using queries from the 14-round oracle. Since 13 bytes of k_{14} are known, we can partially decrypt the corresponding S-Boxes in the last round, and the MixColumns in the first three columns (after replacing tk_{13} by an equivalent key).

We build the same type of structure as used in the 13-round attack. For each plaintext P_i , we query $C_i = E(P_i, T)$ and we partially decrypt the final round. Then we generate the set of 2^{40} values with the required difference in y_{12} , and partially encrypt them to obtain the corresponding ciphertexts \bar{C}_i^j . Finally, we use the stored decryption values to obtain the corresponding \bar{P}_i^j .

We start with a full structure of 2^{128} plaintexts, so that we expect $\mu = 2^{128+128} \cdot \tilde{p}_b = 2^{2.7}$ good quartets. As in the 13-round attack (steps 1 and 2), we match elements on 114 bits: we expect $2^{128+128-114} = 2^{142}$ candidate quartets for each guess of $104 + 32 = 136$ bits of key, or 2^{278} quartets in total.

Following the 13-round attack (steps 3, 4, and 6), we extract on average 2^{-48} candidates for 80 additional key bits. Finally, we use the constraints of the MixColumns operation of round 11: there are only 2^{16} possible differences in w_{11} from which we deduce $2^{16+16-24} = 2^8$ candidates for $tk_{13}^{eq}[2, 5, 8]$. We end up with 2^{238} suggestions for 258 bits of key.

We keep all key candidates suggested at least 6 times. Modeling the counters for wrong keys as following a Poisson distribution with $\lambda = 2^{-20}$, we expect wrong keys to be kept with probability $2^{-129.5}$. We expect $2^{2.7} = 6.5$ good quartets, so the right key candidate is at least 6 with probability more than 1/2. Finally we do an exhaustive search over the 126 key bits remaining, for a cost of $2^{126+128.5} = 2^{254.5}$.

The bottleneck of the attack is the extraction of key candidates for 2^{278} quartets. Following the analysis of the 13-round attack, we estimate that it requires about $2^{8.6}$ table accesses, equivalent to $2^{0.8}$ encryption. The full attack has complexity $(D, T, M) = (2^{129}, 2^{278.8}, 2^{129})$.

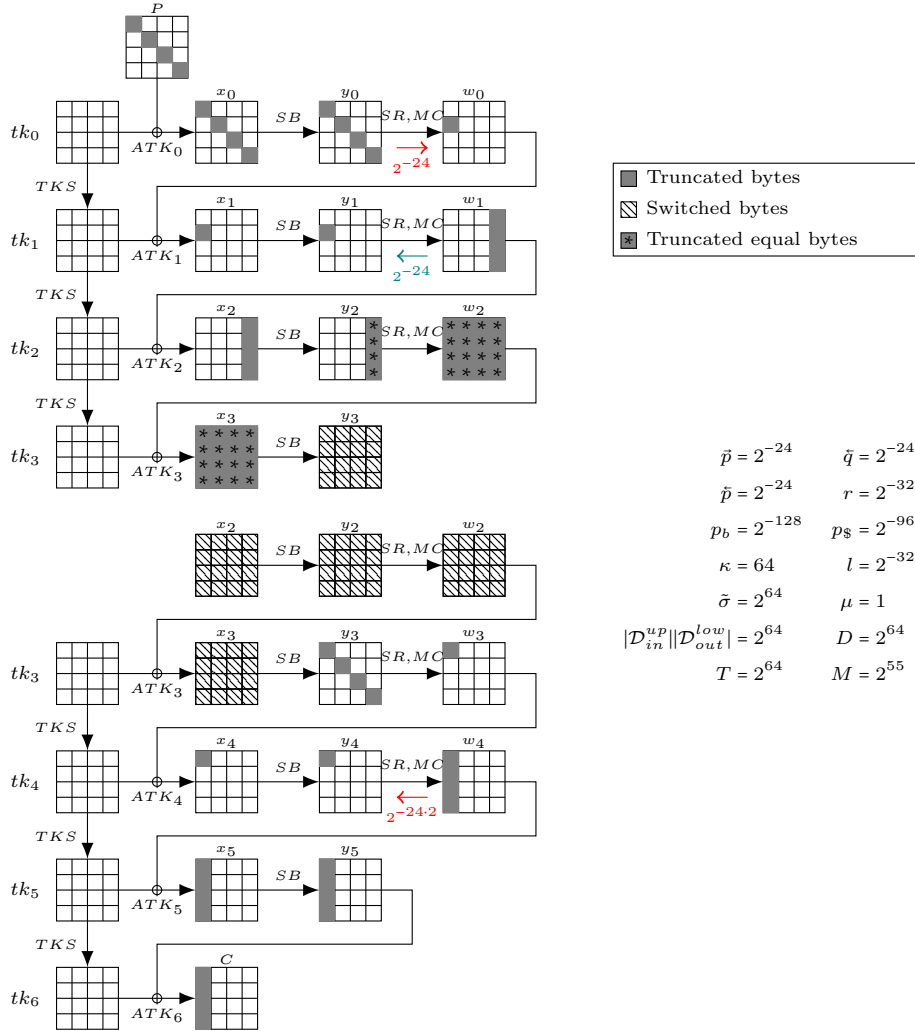


Fig. 8. Truncated boomerang trail on 6-round AES, starting from the plaintext side, returned by the MILP model with *equal* differences in the middle.

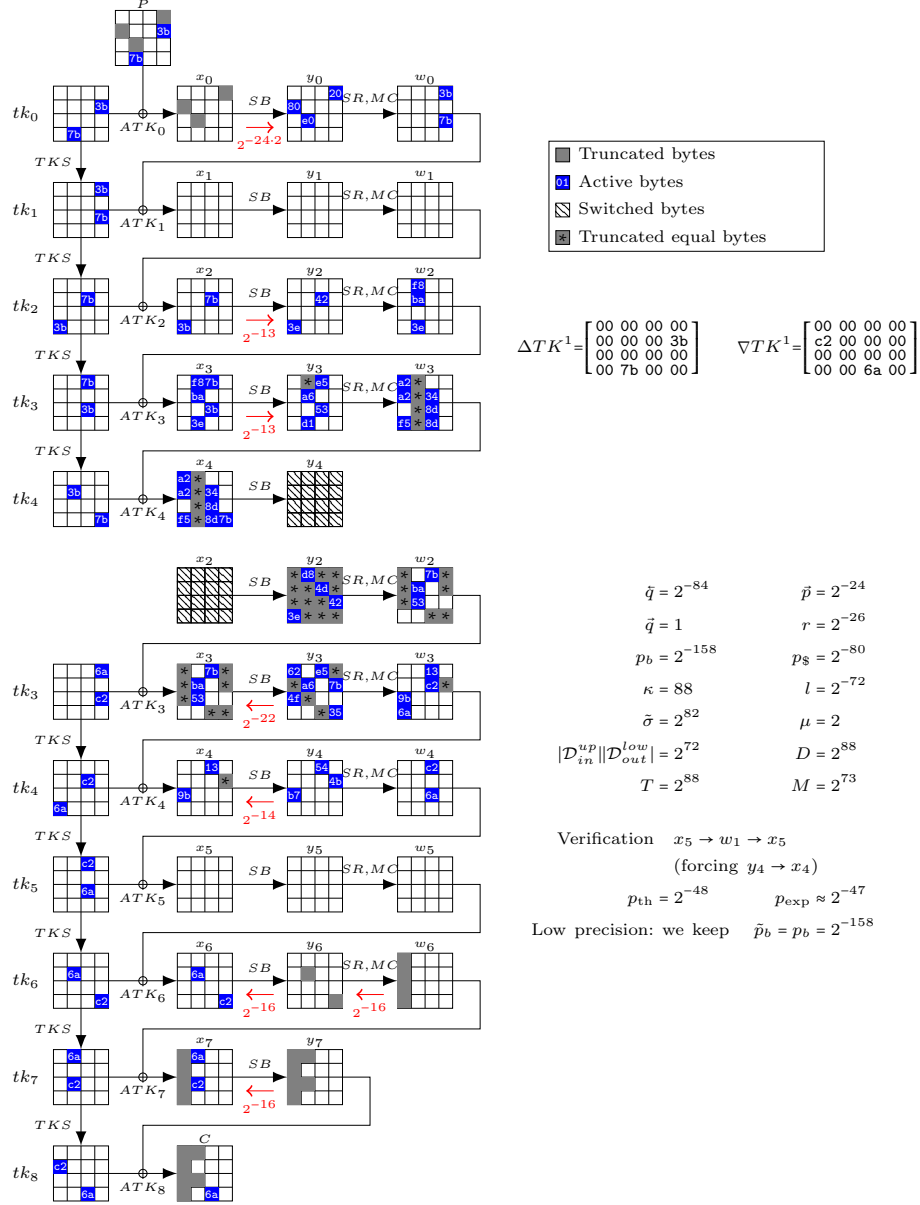


Fig. 9. Truncated boomerang attack on 8-round Deoxys-BC in the RTK1 model, starting from the ciphertext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with UBCT, LBCT and EBCT (probabilities on the trail).




Fig. 10. Truncated boomerang attack on 9-round Deoxys-BC in the RTK1 model, starting from the plaintext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with UBCT, LBCT and EBCT (probabilities on the trail).




Fig. 11. Truncated boomerang attack on 8-round Deoxys-BC in the RTK2 model, starting from the ciphertext side. This attack succeeds with probability $1/2$. Middle rounds are analyzed with the ladder switch and single BCT (probability r).




Fig. 12. Truncated boomerang attack on 9-round Deoxys-BC in the RTK2 model, starting from the ciphertext side. This attack succeeds with probability $1/2$. Middle rounds are analyzed with the ladder switch and single BCT (probability r).




Fig. 13. Truncated boomerang attack on 10-round Deoxys-BC in the RTK3 model, starting from the ciphertext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with the ladder switch and single BCT (probability r).




Fig. 14. Truncated boomerang attack on 11-round Deoxys-BC in the RTK3 model, starting from the ciphertext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with the ladder switch and single BCT (probability r).

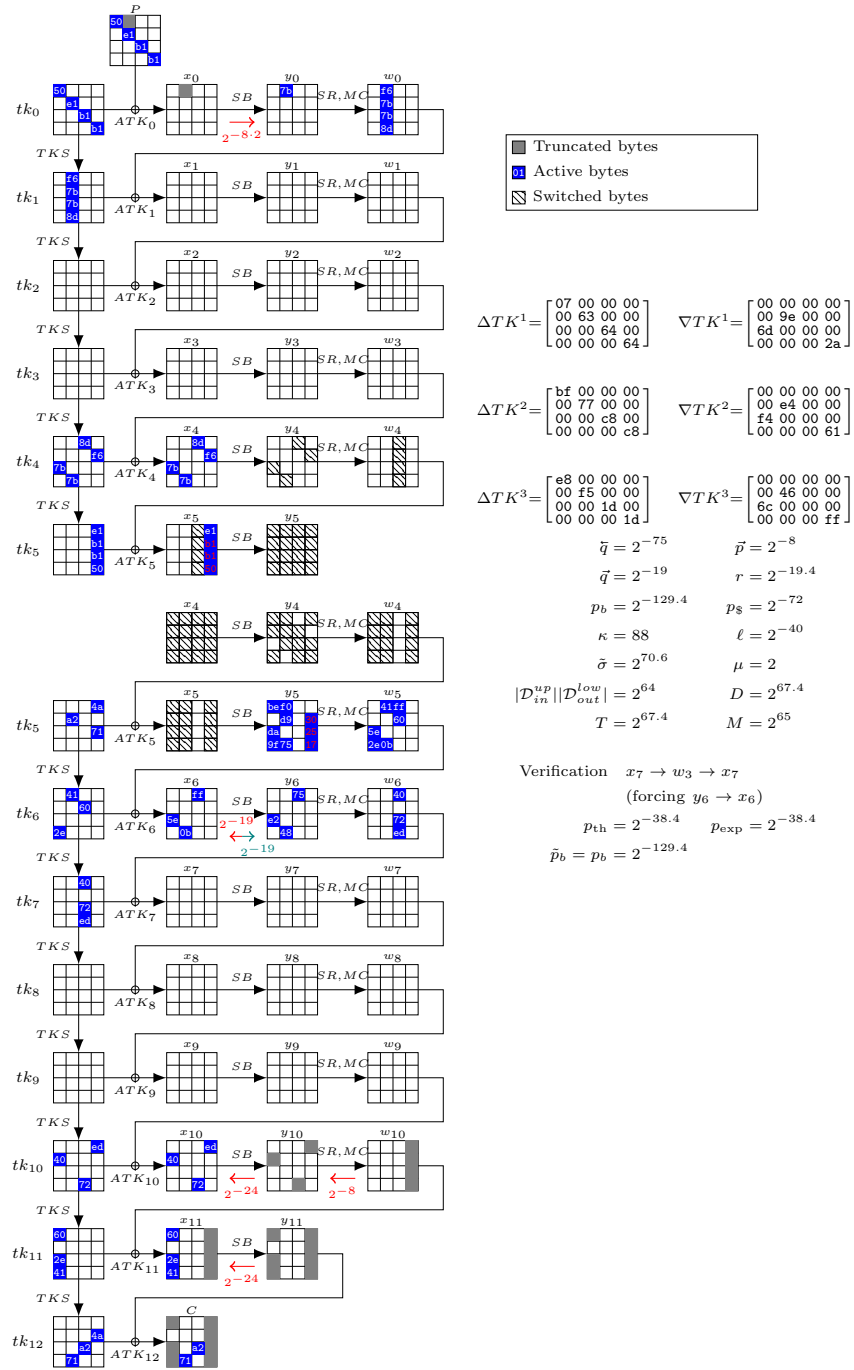


Fig. 15. Truncated boomerang attack on 12-round Deoxys-BC in the RTK3 model, starting from the ciphertext side. This attack succeeds with probability 1/2. Middle rounds are analyzed with the ladder switch and single BCT (probability r).

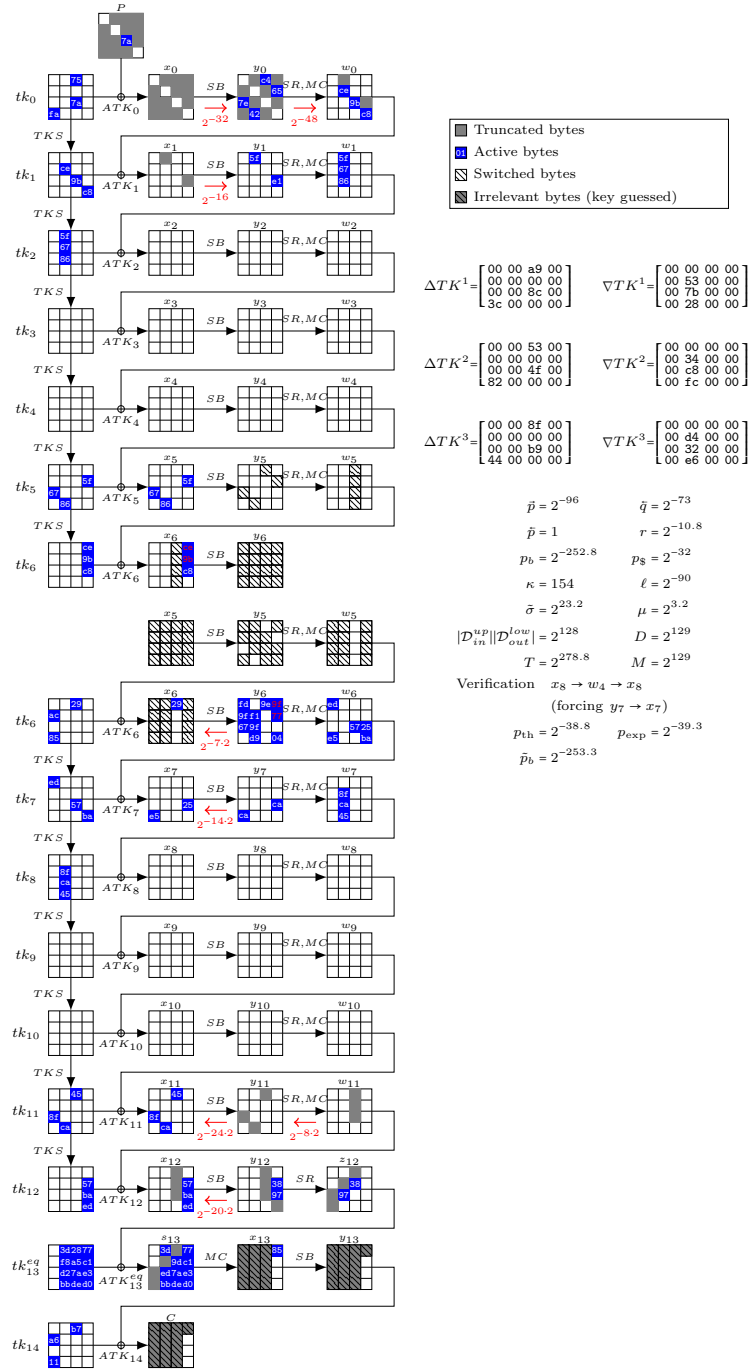


Fig. 16. Truncated boomerang attack on 14-round Deoxys-BC in the RTK3 model, starting from the plaintext side. Middle rounds are analyzed with the ladder switch and single BCT (probability r).