IACR Transactions on Symmetric Cryptology
ISSN 2519-173X, Vol. 2022, No. 1, pp. 88–109. DOI:10.46586/tosc.v2022.i1.88-109

A Formal Analysis of Boomerang Probabilities

Andreas B. Kidmose and Tyge Tiessen

Technical University of Denmark, Kgs. Lyngby, Denmark {abki,tyti}@dtu.dk

Abstract. In the past 20 years since their conception, boomerang attacks have become an important tool in the cryptanalysis of block ciphers. In the classical estimate of their success probability, assumptions are made about the independence of the underlying differential trails that are not well-founded. We underline the problems inherent in these independence assumptions by using them to prove that for any boomerang there exists a differential trail over the entire cipher with a higher probability than the boomerang.

While cryptanalysts today have a clear understanding that the trails can be dependent, the focus of previous research has mostly gone into using these dependencies to improve attacks but little effort has been put into giving boomerangs and their success probabilities a stronger theoretical underpinning. With this publication, we provide such a formalization.

We provide a framework which allows us to formulate and prove rigorous statements about the probabilities involved in boomerang attacks without relying on independence assumptions of the trails. Among these statements is a proof that two-round boomerangs on SPNs with differentially 4-uniform S-boxes always deviate from the classical probability estimate to the largest degree possible.

We applied the results of this formalization to analyze the validity of some of the first boomerang attacks. We show that the boomerang constructed in the amplified boomerang attack on Serpent by Kelsey, Kohno, and Schneier has probability zero. For the rectangle attack on Serpent by Dunkelman, Biham, and Keller, we demonstrate that a minuscule fraction of only $2^{-43.4}$ of all differential trail combinations used in the original attack have a non-zero probability. In spite of this, the probability of the boomerang is in fact a little higher than the original estimate suggests as the non-zero trails have a vastly higher probability than the classical estimate predicts.

Keywords: boomerang attack · cryptanalysis · independence · Serpent

1 Introduction

One of the most important developments in block cipher cryptanalysis was the invention of differential cryptanalysis by Biham and Shamir [BS91]. Any block cipher proposed today must be argued secure against differential attacks. Several ways to do this have been tried over the years, mainly focused on bounding the maximal probability of a single differential trail. The idea was that if the maximal probability was p then at least p^{-1} texts would be needed. Several ciphers with guaranteed security against simple differential attacks were proposed and later broken by more sophisticated methods.

One example is the KN-cipher [NK95], which is a 6-round Feistel cipher that uses x^3 as the nonlinear part of the round function,. The authors had proven this construction secure against ordinary differentials. However, the low degree meant it could be broken by higher-order differentials [JK97] with very low complexity.

Another example is COCONUT98 [Vau98] which used a decorrelation technique to separate the upper and lower halves of the cipher. The upper and lower halves are weak 4-round Feistel networks and the decorrelation module is an addition and a multiplication

Received: 2021-11-23 Accepted: 2022-01-23 Published: 2022-03-11

with key material in a finite field. The multiplication with a secret value makes it impossible to push a difference through the module.

This design caused Wagner [Wag99] to propose the boomerang attack as a clever way to connect unrelated high-probability differentials for the top and bottom half. The basic idea is to "throw" a pair of plaintexts through the cipher, add a difference to the resulting ciphertexts, and observe how they return (see Figure 1). A second-order differential in the middle of the cipher connects the differentials for the upper and lower halves causing the boomerang to return. The classical analysis of the probability of the boomerang returning (see Subsection 2.1) assumes the differentials act independently.

This attack worked exceptionally well on the COCONUT98 cipher since the 4-round Feistel cipher admit differentials with probability of $4^{-4.3}$. The important observation is that for a fixed key the decorrelation module is affine. This means there is a probability-one transition, however, predicting the output difference is impossible without knowledge of the key. Knowing the exact difference is not needed for the attack, only that any second-order derivative is zero with probability one.

Some variants to the basic boomerang attack have been proposed over the years. The idea of the amplified boomerang attack by Kohnu et al. [KKS00] is to turn the boomerang attack into a chosen-plaintext attack instead the original adaptively chosen-plaintext/ciphertext attack. The rectangle attack by Biham et al. [BDK01] builds up on the amplified boomerang by making use of the fact that the differences in the middle need not be fixed but can take any value, as long as the sum is 0. The sandwich attack by Dunkelman et al. [DKS10, DKS14] proposes a framework where the two differentials are separated in the middle, like two pieces of bread with a thin slice of meat hence the name. The differentials for the upper and lower halves are then connected via ad-hoc methods through the middle round. As a framework it has proven to be a good basis for investigating the dependencies of the differentials involved in a boomerang attack.

The boomerang attack and its variants have proven themselves to be effective on a wide variety of ciphers. Notable examples include the attack on AES by Biryukov and Khovratovich [BK09], and on KASUMI by Dunkelman et al. [DKS10]. More recently the retracing boomerang attack was introduced by Dunkelman et al. [DKRS20], which improved the best attack on 5-round AES by discarding some data and forcing the boomerang back along the same trajectory.

Related work

The differentials in a boomerang attack were usually assumed to be independent, however there is no a priori justification for that assumption. Several techniques, commonly known as boomerang switches, have been proposed to take advantage of dependencies to boost the probability of the boomerang. The *Feistel switch*, which bypasses a round for free, was already implemented by Wagner in [Wag99] in the attack on Khufu. The *ladder switch* and the *S-box switch* were introduced by Biryukov and Khovratovich [BK09]. In the ladder switch the attacker chooses the boundary of the two differentials such that it does not necessarily align with the rounds of the cipher. When putting an S-box in a differential where it is inactive, instead of active, it does not add to the probability. The S-box switch is the fact that, if the output difference for the upper differential matches the difference from the lower differential, then the pairs are just swapped and we only pay for the probability in one direction.

While the switches were used to aid the attacker, Murphy [Mur11] pointed out that the differentials might in fact be incompatible. In the middle of the boomerang, where the upper and lower differentials meet, we have a pair of pairs. The upper differential defines the distance between the pairs and the lower defines the difference for the differential transition. It may be the case that there are several pairs that follow the transition for the lower differential but none with the distance dictated by the upper differential. Murphy in particular showed an example for DES, where the required transition over S_2 in round 4 is impossible, and therefore that particular boomerang never comes back.

Cid et al. [CHP+18] proposed the boomerang connectivity table (BCT) as a unified approach to calculate the boomerang switching probability in SPN-ciphers when the middle part is one S-box layer. It includes the previously mentioned switches and as well as a new switching property they discovered, and it can be used to show when two differentials are incompatible.

The BCT depends on the inverse S-box, however since the inverse S-box is not used in a Feistel cipher, a different approach is needed here. The FBCT, for Feistel boomerang connectivity table, was introduced by Boukerrou et al. [BHL+20] to solve this issue. The authors also show the invalidity of the related-key boomerang attack on LBlock by Liu et al. [LGW12].

The BCT and FBCT work at the S-box level and as such only consider one round; the dependency can, however, span many rounds. Two concurrent papers looked at this problem with different approaches. Wang and Peyrin [WP19] took a table-based approach, where the DDT and BCT are combined in a table which they call the boomerang difference table (BDT). This table can be used to evaluate the probability of a boomerang switch over 2 rounds. Song et al. [SQH19] instead proposed a way to determine the length of the middle part, where the dependency exists.

The classical way to find a good boomerang distinguisher would be to choose the best differentials for the upper and lower part separately, and then just hope that they are compatible. The problem is that the best differentials might not have a high probability of connecting in the middle, and therefore choosing a lower probability differential might result in a higher probability boomerang. Recently several MILP models have been proposed to search for boomerang distinguishers, e.g., [DDV20] and [HBS21], which will take the switching probability into account for multiple rounds.

Motivation

Since the inception of the boomerang attack, we have come to appreciate some of the difficulties involved in estimating the probabilities of boomerang distinguishers. While there is a general understanding that the naive method of estimating boomerang probabilities as the product of the individual involved trails is incorrect, and while dependencies between the trails have been put to good use in attacks such as the sandwich attack, we still lack a consistent model of describing boomerang probabilities.

With this work, we want to fill this gap by creating a mathematical model that allows us to precisely formulate the probabilities of boomerang attacks. The only assumptions that we want to rely on are those commonly made in differential cryptanalysis.

Our contribution

In this paper, we take a close look at the probability estimates classically made in boomerang attacks and which assumptions are being made. We show that using these assumptions we can prove that for any boomerang there would necessarily exists a differential over the entire cipher with higher probability than the boomerang. While this is clearly not the case, it underlines the need for a better formal underpinning of boomerang success probabilities.

Building up on a notation that extends the notions from differential cryptanalysis to take a quartet of messages into account, we are able to rigorously prove several results regarding the probability of boomerang attacks. Among these are compact expressions of the boomerang probabilities as well as results on the applicability and limitation of the classical estimates of boomerang probabilities. In particular we are able to prove that

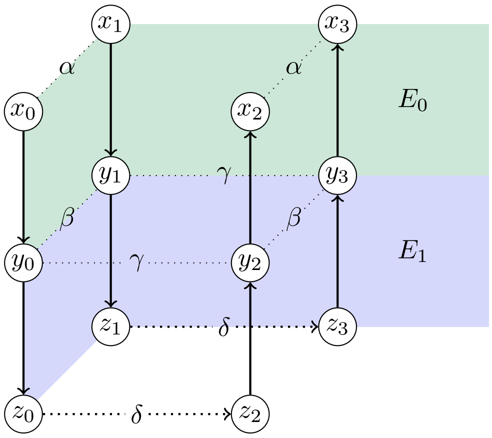

Figure 1: Outline of a basic boomerang attack.

two-round boomerangs on SPNs with differentially 4-uniform S-boxes never adhere to the classical probability estimate.

We furthermore apply our results to two classic boomerang attacks on the block cipher Serpent. The first one, the amplified boomerang attack on Serpent [KKS00], is shown to be invalid as stated. For the second one, the rectangle attack [BDK05], we show that of all boomerang trails considered in the attack only a fraction of $2^{-43.4}$ have a non-zero probability. For the remaining boomerang trails we demonstrate that their probability is much higher than a classical estimate would suggest, leaving the total combined probability of all boomerang trails close to the original estimate. By including even more trails in a refined analysis we are able to improve this original estimate by a factor of $2^{4.3}$.

Outline

In Section 2 we introduce the basic boomerang attack. In Section 3 we prove that the independence assumption is an inherently flawed assumption. In Section 4 we introduce d-differences which will be used in Section 5 to prove some statements about boomerang probabilities. In Section 7 we look at two boomerang attacks on Serpent, and finally Section 8 concludes the paper.

2 Boomerang attacks

Before we can start our investigation into formalizing the boomerang attack probabilities, we need a proper exposition of this attack and its variants.

2.1 Basic boomerangs

Let us start by properly introducing boomerang attacks, developed by David Wagner [Wag99] and set the notation which is used throughout this paper.

Let Enc denote a block cipher that maps n-bit plaintexts bijectively to n-bit ciphertexts. For the purpose of the attack, we assume that the cipher can be decomposed into two parts E_0 and E_1 such that

$$\mathrm{Enc} = E_1 \circ E_0 \ . \tag{1}$$

We are now interested in the scenario where we have a good differential $\alpha \xrightarrow{E_0} \beta$ over E_0 that holds with probability p, as well as a good differential $\gamma \xrightarrow{E_1} \delta$ over E_1 that holds with probability q. These two differentials can now be used to construct a distinguisher over the whole cipher as follows.

To construct the distinguisher, we start with a pair of plaintexts x_0 and x_1 with a difference α. When encrypting these two plaintexts, we expect the corresponding intermediate texts $y_0 := E_0(x_0)$ and $y_1 := E_0(x_1)$ to have a difference β with probability p (at least according to the standard assumptions of differential cryptanalysis). For orientation, see also Fig. 1.

With $z_0 := E_1(y_0)$ and $z_1 := E_1(y_1)$ being the respective ciphertexts, we now construct two more ciphertexts $z_2 := z_0 \oplus \delta$ and $z_3 := z_1 \oplus \delta$ by adding the difference δ to each of z_0 and z_1. Then the pairs (z_0, z_2) and (z_1, z_3) both have a difference of δ, the ciphertext difference in the second differential.

Decrypting these two ciphertexts, provides us with two more intermediate texts, $y_2 := E_1^{-1}(z_2)$ and $y_3 := E_1^{-1}(z_3)$, and two more plaintexts, $x_2 := E_0^{-1}(y_2)$ and $x_3 := E_0^{-1}(y_3)$. Assuming independence of the two ciphertext pairs (z_0, z_2) and (z_1, z_3), both of their respective intermediate pairs (y_0, y_2) and (y_1, y_3) will have a differences of γ with probability q^2.[1]

Combining this with the probability that (x_0, x_1) follows the first differential, we have with probability pq^2 that $y_0 \oplus y_1 = \beta$, $y_0 \oplus y_2 = \gamma$, and $y_1 \oplus y_3 = \gamma$. This forces the difference between y_2 and y_3 to be β. Again assuming independence from the other pairs, the pair (y_2, y_3) will follow the first differential with probability p, resulting in a plaintext difference of α between x_2 and x_3.

Taking all of these steps together, we estimate that the probability to see a difference α between x_2 and x_3 is equal to p^2q^2.

Over a random permutation, the probability for x_2 and x_3 to have a difference of α is $(2^n - 1)^{-1}$. We therefore expect such a boomerang distinguisher to be successful as long as p^2q^2 is sufficiently larger than $(2^n - 1)^{-1}$, since we can then use the above technique to distinguish the cipher from a random permutation in a chosen plaintext/chosen ciphertext attack.

In accordance with the above notation, we make the following definition:

Definition 1. We call a tuple of four plaintexts $(x_0, x_1, x_2, x_3) \in \mathbb{F}_2^{4n}$ a *right quartet* with respect to the input difference $\alpha \in \mathbb{F}_2^n$ and the output difference $\delta \in \mathbb{F}_2^n$ and a fixed key if and only if $x_0 \oplus x_1 = \alpha$, $x_2 \oplus x_3 = \alpha$, and $\mathrm{Enc}(x_0) \oplus \mathrm{Enc}(x_2) = \delta$ and $\mathrm{Enc}(x_1) \oplus \mathrm{Enc}(x_3) = \delta$.

Since the only non-deterministic part of the method is our initial choice of x_0, it is straightforward to see that the probability to detect the output difference α between x_2 and x_3 using the boomerang attack is equal to the number of right quartets divided by 2^n.

As described above, classically this probability is estimated to be p^2q^2 by making mentioned independence assumptions. A focal point of this paper is to shed light upon how and why these independence assumptions fail and what the consequences are for the estimate of the boomerang probability.

2.1.1 Taking more differentials into consideration

When estimating the number of expected right quartets with the classical assumptions as $2^np^2q^2$, we clearly are estimating a lower bound as we are only considering two particular differentials. To get a more accurate classical estimate of the boomerang probability we need to consider all possible values for the differences β and γ in the middle of the cipher:

$$\sum_{\beta \in \mathbb{F}_2^n, \gamma \in \mathbb{F}_2^n} \Pr\left(\alpha \xrightarrow{E_0} \beta\right)^2 \Pr\left(\gamma \xrightarrow{E_1} \delta\right)^2. \tag{2}$$

[1] The probability of a differential is the same both in the forward and backward direction for any bijective function. Note that this is not true for truncated differentials.

As a matter of fact, there is no reason to restrict the differences $y_0 \oplus y_1$ and $y_2 \oplus y_3$ to be the same, and likewise the differences $y_0 \oplus y_2$ and $y_1 \oplus y_3$ to be the same. Thus an even better classical approximation of the boomerang probability is thus

$$\sum_{\substack{\beta_0,\beta_1,\gamma_0,\gamma_1 \in \mathbb{F}_2^n, \\ \beta_0 \oplus \beta_1 \oplus \gamma_0 \oplus \gamma_1 = 0}} \Pr\left(\alpha \xrightarrow{E_0} \beta_0\right)\Pr\left(\alpha \xrightarrow{E_0} \beta_1\right)\Pr\left(\gamma_0 \xrightarrow{E_1} \delta\right)\Pr\left(\gamma_1 \xrightarrow{E_1} \delta\right). \tag{3}$$

For references of these classical estimates, see for example Wagner [Wag99, Section 6] or Biham et al. [BDK01, Section 4].

2.1.2 Practical restrictions

To find a good approximation of the boomerang probability, one would ideally like to determine the expressions in Eq. (3) or Eq. (2). This is in most cases impossible in practice. As a matter of fact, even when we only work with two differentials for each cipher part, it is usually computationally infeasible to determine the probabilities of just these differentials. What we tend to do in practice then, is to restrict ourselves to the most promising differential trails for the upper and lower parts of the cipher and take their probabilities to determine a good approximation for the boomerang probability.

2.2 Amplified boomerangs

One inherent limitation of the original boomerang attack is that it requires both chosen plaintexts and adaptively-chosen ciphertexts. The attack can be adapted to only requiring chosen plaintexts albeit with a much higher data complexity. This method was developed by Kelsey, Kohno, and Schneier [KKS00]. The complexity of this method also inherently depends on the number of expected right quartets, i.e., the boomerang probability.

The idea is to choose two values x_0, x_2 and generate $x_1 = x_0 \oplus \alpha$ and $x_3 = x_2 \oplus \alpha$. Then like in the standard boomerang attack the differences in the middle will be β for both pairs with probability p^2, that is, $y_0 \oplus y_1 = y_2 \oplus y_3 = \beta$. The distance will be γ with a probability of 2^{-n} since if $y_0 \oplus y_2 = \gamma$ then so is $y_1 \oplus y_3 = \gamma$. Finally the transition for the lower half is again q^2 and therefore the probability that $z_0 \oplus z_2 = z_1 \oplus z_3 = \delta$ is $p^2q^22^{-n}$. By creating a large set of pairs of input texts with difference α, the large number of possible combinations of input pairs allows the success probability to be higher than 2^{-n} then.

3 Independence assumptions in boomerang attacks

3.1 Independence of rounds and trails

The assumptions used in boomerang attacks can generally be put into two categories: those assumptions that stem from the theory of differential cryptanalysis and those assumptions that are specific to the boomerang attack.

The most important assumption from differential cryptanalysis is the assumption that the rounds of a cipher can be treated independently when determining differential probabilities, thereby allowing us to multiply the probabilities of differential round transitions to obtain the probability of a trail. We will briefly discuss this assumption when creating our model for calculating boomerang probabilities, so for now let us leave it by saying that this assumption works sufficiently well in practice.

The additional assumption made in the classical estimate of the boomerang probability though is about the independence of trails and cannot be derived from the standard assumptions of differential cryptanalysis. In the classical estimate of the boomerang probability, we simply multiply the probabilities of the four individual differential transitions.

To be able to do this, we implicitly assume that these four differentials can be regarded as independent. We can state this assumption as follows:

Assumption 1. For any two text pairs (x_0, x_1) and (x_2, x_3) that both have a difference α, i.e., $x_0 \oplus x_1 = x_2 \oplus x_3 = \alpha$, the probability that both pairs are mapped to a difference β is the same as the square of the probability of the differential $\alpha \longrightarrow \beta$ for any choice of α and β.

Clearly this assumption does not hold when the text pairs coincide as the second pair then always follows the same differential as the first. In the following, we show that even assuming that this assumption holds closely for the case where the text pairs are distinct leads to contradictions.

3.2 An inherent problem

The idea and purpose of boomerang attacks is to provide an attack in scenarios where we might not find a good differential that covers the entire cipher. There seems to be little reason to use a boomerang attack if we already have a differential of higher probability over the entire cipher. The more surprising it might be that we can prove that there always exists a differential with probability higher than the boomerang probability when we rely on the assumption that we can treat the trail probabilities independently, as we do in the classical estimate.

Using the notation for the boomerang attack established in Section 2.1, we thus would expect that no differentials $\beta \xrightarrow{E_1} \varepsilon$ or $\eta \xrightarrow{E_0} \gamma$ exist for which

$$\Pr\left(\beta \xrightarrow{E_1} \varepsilon\right) \geq q^2 \quad \text{or} \quad \Pr\left(\eta \xrightarrow{E_0} \gamma\right) \geq p^2. \tag{4}$$

If this were not the case, at least one of the differential trails $\alpha \xrightarrow{E_0} \beta \xrightarrow{E_1} \varepsilon$ and $\eta \xrightarrow{E_0} \gamma \xrightarrow{E_1} \delta$ would have a probability higher than p^2q^2 and thus would be better suited as a distinguisher than the boomerang distinguisher.

Theorem 1. *Assume that we have a boomerang as described above of probability p^2q^2 and assume that the assumption of the independence of differentials holds (Assumption 1). Then there exist differentials $\alpha \xrightarrow{\text{Enc}} \varepsilon$ and $\eta \xrightarrow{\text{Enc}} \delta$ over the whole cipher with probabilities at least pq^2 and qp^2, respectively.*

Proof. We only prove here that there exists a differential $\alpha \xrightarrow{\text{Enc}} \varepsilon$ of probability at least pq^2. To show that a differential $\eta \xrightarrow{\text{Enc}} \delta$ of probability at least p^2q exists goes analogously.

We start by showing that if Assumption 1 holds for E_1, then

$$\sum_{\varepsilon} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right)^2 = \sum_{\delta} \Pr\left(\gamma \xrightarrow{E_1} \delta\right)^2. \tag{5}$$

By choosing some arbitrary fixed γ, we get that

$$\begin{aligned}
&\sum_{\varepsilon} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right)^2 \\
&= \sum_{\varepsilon} \Pr\Big(E_1(x_0 \oplus \beta) \oplus E_1(x_0) = \varepsilon \text{ and } E_1(x_0 \oplus \gamma \oplus \beta) \oplus E_1(x_0 \oplus \gamma) = \varepsilon\Big) \\
&= \Pr\Big(E_1(x_0 \oplus \beta) \oplus E_1(x_0) = E_1(x_0 \oplus \gamma \oplus \beta) \oplus E_1(x_0 \oplus \gamma)\Big) \\
&= \Pr\Big(E_1(x_0 \oplus \gamma) \oplus E_1(x_0) = E_1(x_0 \oplus \gamma \oplus \beta) \oplus E_1(x_0 \oplus \beta)\Big) \\
&= \sum_{\delta} \Pr\Big(E_1(x_0 \oplus \gamma) \oplus E_1(x_0) = \delta \text{ and } E_1(x_0 \oplus \gamma \oplus \beta) \oplus E_1(x_0 \oplus \beta) = \delta\Big) \\
&= \sum_{\delta} \Pr\left(\gamma \xrightarrow{E_1} \delta\right)^2,
\end{aligned}$$

where we used Assumption 1 in the first and last steps of the derivation. This concludes the proof of Eq. (5).

The probability that both these β differences are mapped to the same value thus gives us an upper bound for the probability that both δ differences are mapped to γ differences:

$$\sum_{\varepsilon \in \mathbb{F}_2^n} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right)^2 \geq \Pr\left(\gamma \xrightarrow{E_1} \delta\right)^2 = q^2. \tag{6}$$

Let s now be the maximal value for any of the differentials $\beta \xrightarrow{E_1} \varepsilon$:

$$s := \max_{\varepsilon \in \mathbb{F}_2^n} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right). \tag{7}$$

Using the fact that the trail probabilities sum to one $\left(\sum_{\varepsilon \in \mathbb{F}_2^n} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right) = 1\right)$, we then have that

$$q^2 \leq \sum_{\varepsilon \in \mathbb{F}_2^n} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right)^2 \leq \sum_{\varepsilon \in \mathbb{F}_2^n} s \cdot \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right) = s \cdot \sum_{\varepsilon \in \mathbb{F}_2^n} \Pr\left(\beta \xrightarrow{E_1} \varepsilon\right) = s \tag{8}$$

Thus there exists a differential $\beta \xrightarrow{E_1} \varepsilon$ with probability at least q^2 and thus there exists a differential $\alpha \xrightarrow{\text{Enc}} \varepsilon$ of probability at least pq^2.

□

Can we conclude from this that there always exist differentials that beat boomerang distinguishers and that we only need to find them? Certainly not (see for example Corollary 1). It much rather demonstrates that we must be very careful when unconditionally assuming independence of differentials as done in the classical estimate of the boomerang probability.

4 Generalized differences and their transitions

Before we look in more detail into the probabilities of boomerangs, let us introduce some notation and a model that allows us to formally discuss boomerang attacks. Parts of the notation that we are using are taken from [Tie16].

In differential cryptanalysis, we are usually not interested in the absolute position of texts in the state space but only in their relative positions, i.e., their relative differences.

The relative positions of a tuple of $d+1$ texts are uniquely defined by the differences of the d last texts with respect to the first text. To capture this notion, we define:

Definition 2 (d-difference)**.** For a tuple of $(d+1)$ texts $(m_0, m_1, m_2, \ldots, m_d)$, we describe their relative differences by the d-tuple

$$(m_1 \oplus m_0, m_2 \oplus m_0, \ldots, m_n \oplus m_0). \tag{9}$$

We refer to such a d-tuple as the *d-difference* of the $(d+1)$-tuple of messages. We refer to the first text of the $(d+1)$-tuple of messages as the *anchor* of the d-difference.

Note that d-differences thus describe the translation-invariant equivalence class of $(d+1)$-tuples in the state space. Thus a $(d+1)$-tuple is uniquely identified by its d-difference together with its anchor. For the remainder, we will refer to a $(d+1)$-tuple only as a tuple if the value d is clear from the context.

Extending the notion of differentials to d-differences, we make the following definition:

Definition 3 (Transition with fixed anchor)**.** Let $f : \mathbb{F}_2^n \to \mathbb{F}_2^n$. Let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_d)$ and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_d)$ be two d-differences over $\mathbb{F}_2^n$. By the transition $\boldsymbol{\alpha} \xrightarrow[x]{f} \boldsymbol{\beta}$, we denote the event that f maps the tuple of messages corresponding to the d-difference $\boldsymbol{\alpha}$ with anchor x to a tuple of messages with d-difference $\boldsymbol{\beta}$. More precisely, we say that $\boldsymbol{\alpha} \xrightarrow[x]{f} \boldsymbol{\beta}$ holds if and only if

$$\begin{gathered} f(x \oplus \alpha_1) \oplus f(x) = \beta_1, \\ f(x \oplus \alpha_2) \oplus f(x) = \beta_2, \\ \cdots \\ f(x \oplus \alpha_d) \oplus f(x) = \beta_d. \end{gathered}$$

Example 1. Let (m_0, m_1, m_2, m_3) be a plaintext tuple and let (c_0, c_1, c_2, c_3) be the corresponding tuple of ciphertexts. Set $\boldsymbol{\alpha} = (m_1 \oplus m_0, m_2 \oplus m_0, m_3 \oplus m_0)$ and let $\boldsymbol{\beta}$ be some 3-difference. Then $\boldsymbol{\alpha} \xrightarrow[m_0]{f} \boldsymbol{\beta}$ holds if and only if $\boldsymbol{\beta} = (c_0 \oplus c_1, c_0 \oplus c_2, c_0 \oplus c_3)$.

To define the probability of such transitions where the anchor is not fixed, we take the same route as standard differential cryptanalysis and assume that the first text is chosen uniformly at random.

Definition 4 (Probability of transitions)**.** Let f, $\boldsymbol{\alpha}$, and $\boldsymbol{\beta}$ again be as in Definition 3. The probability of the transition $\boldsymbol{\alpha} \xrightarrow{f} \boldsymbol{\beta}$ is then defined as:

$$\Pr\left(\boldsymbol{\alpha} \xrightarrow{f} \boldsymbol{\beta}\right) := \Pr_{\mathbf{X}}\left(\boldsymbol{\alpha} \xrightarrow[\mathbf{X}]{f} \boldsymbol{\beta}\right) \tag{10}$$

where $\mathbf{X}$ is a random variable, distributed uniformly on $\mathbb{F}_2^n$.

For simple differences (1-differences) this definition coincides with the definition of differentials.

Example 2. Let f be the AES S-box, and consider the 3-difference transition $(\alpha_1, \alpha_2, \alpha_3) \xrightarrow{f} (\beta_1, \beta_2, \beta_3)$, where $\alpha_1 = 7$, $\alpha_2 = 25$, and $\alpha_3 = \alpha_1 \oplus \alpha_2$ and $\beta_1 = 166$, $\beta_2 = 183$, and $\beta_3 = \beta_1 \oplus \beta_2$. To calculate the probability we simple count all values of x for which $f(x \oplus \alpha_1) \oplus f(x) = \beta_1$ and $f(x \oplus \alpha_2) \oplus f(x) = \beta_2$ and $f(x \oplus \alpha_3) \oplus f(x) = \beta_3$. There are exactly 4 values (0, 7, 25, and 30) which means that the probability is $2^2 \cdot 2^{-8} = 2^{-6}$. If we instead change $\beta_2 = 1$, then there are no values for x for which it holds and the probability is therefore 0. Note that this also illustrates Lemma 2.

To allow us to lower bound the probability of a transition over several rounds of a cipher, we need the ability to split transition into a collection of trails. To this end, we define what a trail is in this context.

Definition 5 (A d-difference trail)**.** Let $f = f_l \circ \cdots \circ f_2 \circ f_1$ be a composition of n-bit to n-bit functions, let $\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_l \in \mathbb{F}_2^{dn}$ be a sequence of d-differences, and let $x \in \mathbb{F}_2^n$. We refer to the sequence $(\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_l)$ as a trail over f and denote the event that this trail is followed as $\boldsymbol{\alpha}_0 \xrightarrow[x]{f_0} \boldsymbol{\alpha}_1 \xrightarrow[f_0(x)]{f_1} \boldsymbol{\alpha}_2 \cdots \xrightarrow[f_{l-1}\circ\cdots\circ f_0(x)]{f_l} \boldsymbol{\alpha}_l$. That is, we say that the $(d+1)$-tuple corresponding to $\boldsymbol{\alpha}_0$ with the anchor x follows the trail $(\boldsymbol{\alpha}_0, \ldots, \boldsymbol{\alpha}_l)$ if and only if all the transitions

$$\begin{gathered}\boldsymbol{\alpha}_0 \xrightarrow[x]{f_0} \boldsymbol{\alpha}_1, \\ \boldsymbol{\alpha}_1 \xrightarrow[f_0(x)]{f_1} \boldsymbol{\alpha}_2, \\ \ldots \\ \boldsymbol{\alpha}_{l-1} \xrightarrow[f_{l-1}\circ\cdots\circ f_0(x)]{f_l} \boldsymbol{\alpha}_l\end{gathered}$$

are adhered to.

To be able to efficiently determine the probability of these trails, we run into the same problem that one encounters when formalizing the probability of differentials trails, namely that the transitions that make up a trail are generally not independent. To solve this issue we make the same assumption that is conventionally made in differential cryptanalysis, namely that we can reasonably well approximate the probability of a trail by considering the individual transitions as independent. This independence is achieved by assuming that the anchor used in each of the transitions is distributed uniformly randomly. This can for example be modelled by saying that a uniformly random constant is added onto the state after every found/round (for the classic example see [LMM91]).

Assumption 2 (Hypothesis of stochastic equivalence)**.** We assume that treating the individual transitions of a trail as independent gives a reasonably good approximation of the real trail probability. Using the previously established notation, we write this as

$$\Pr_{\mathbf{X}}\Big(\boldsymbol{\alpha}_0 \xrightarrow[\mathbf{X}]{f_0} \boldsymbol{\alpha}_1 \xrightarrow[f_0(\mathbf{X})]{f_1} \boldsymbol{\alpha}_2 \cdots \xrightarrow[f_{l-1}\circ\cdots\circ f_0(\mathbf{X})]{f_l} \boldsymbol{\alpha}_l\Big) \approx \Pr\left(\boldsymbol{\alpha}_0 \xrightarrow{f_0} \boldsymbol{\alpha}_1\right) \cdots \Pr\left(\boldsymbol{\alpha}_{l-1} \xrightarrow{f_l} \boldsymbol{\alpha}_l\right) .$$

Note that for the case of $d = 2$, this corresponds exactly to the standard assumption made in differential cryptanalysis.

We will now state some rules without proof that are useful when working with these transitions (or when working with standard differentials for that matter). For the purpose of readability, and as we are mostly concerned with quartets of messages here, we fix $d = 3$. Let now $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3)$ and $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)$ be two 3-differences and let f be a bijective function from $\mathbb{F}_2^n$ to $\mathbb{F}_2^n$.

Rule 1. $\boldsymbol{\alpha} \xrightarrow[x]{f} \boldsymbol{\beta} \iff \boldsymbol{\beta} \xrightarrow[f(x)]{f^{-1}} \boldsymbol{\alpha}$

Rule 2.

$$\begin{aligned}&(\alpha_1, \alpha_2, \alpha_3) \xrightarrow[x]{f} (\beta_1, \beta_2, \beta_3) \\ &\qquad \iff (\alpha_1, \alpha_1 \oplus \alpha_2, \alpha_1 \oplus \alpha_3) \xrightarrow[x\oplus\alpha_1]{f} (\beta_1, \beta_1 \oplus \beta_2, \beta_1 \oplus \beta_3) \quad (11)\end{aligned}$$

In accordance with the common definition of truncated differentials, we also define:

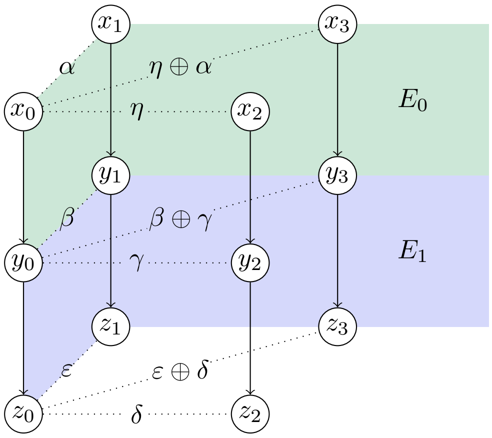

Figure 2: Outline of a simple boomerang attack in the d-difference view. The differences v and w are allowed to take any value here.

Definition 6 (Transitions of truncated d-differences)**.** A truncated d-difference is an affine subspace in the linear space of d-differences. A truncated d-difference transition is a pair (A, B) of truncated d-differences denoted as $A \xrightarrow{f} B$. The probability of such a truncated d-difference transition $A \xrightarrow{f} B$ is then defined as the probability that an input d-difference chosen uniformly at random from A maps to a d-difference in B:

$$\Pr\left(A \xrightarrow{f} B\right) := |A|^{-1} \sum_{\substack{\boldsymbol{\alpha} \in A \\ \boldsymbol{\beta} \in B}} \Pr\left(\boldsymbol{\alpha} \xrightarrow{f} \boldsymbol{\beta}\right). \tag{12}$$

From this definition and Rule 1 follows immediately the following rule:

Rule 3.

$$|A| \Pr\left(A \xrightarrow{f} B\right) = |B| \Pr\left(B \xrightarrow{f^{-1}} A\right). \tag{13}$$

5 Rigorous statements for boomerang probabilties

As we will see in the following, d-differences and truncated d-difference transitions provide a notation that allows us to formalize statements about the probability of boomerangs in a consistent model.

Let us start by looking at a tuple of four texts (x_0, x_1, x_2, x_3) where the differences of the pairs (x_0, x_1) and (x_2, x_3) are α each. Using 3-differences, we can say that this 4-tuple (x_0, x_1, x_2, x_3) corresponds to a 3-difference $(\alpha, \eta, \alpha \oplus \eta)$ for some $\eta \in \mathbb{F}_2^n$. With this view, we can think of the relationships between the texts in the boomerang attack as 3-differences. This alternative view is depicted in Fig. 2 (see also Fig. 1 for comparison).

Looking at the boomerang attack from this 3-difference perspective, we can state the probability of the return of the boomerang:

Theorem 2. *Let A be the affine subspace of all* 3*-differences which correspond to an input quartet:*

$$A := \left\{(\alpha, \eta, \alpha \oplus \eta) \in \mathbb{F}_2^{3n} \;\middle|\; \eta \in \mathbb{F}_2^n\right\}. \tag{14}$$

Let B be the set of all 3*-differences which correspond to a right ciphertext quartet:*

$$B := \left\{(\varepsilon, \delta, \varepsilon \oplus \delta) \in \mathbb{F}_2^{3n} \;\middle|\; \varepsilon \in \mathbb{F}_2^n\right\}. \tag{15}$$

The probability of the return of the boomerang is then equal to the probability of the truncated 3-difference transition $A \xrightarrow{\text{Enc}} B$ multiplied by 2^n:

$$\Pr\left(\textit{Boomerang returns}\right) = 2^n \cdot \Pr\left(A \xrightarrow{\text{Enc}} B\right). \tag{16}$$

Proof. In this proof $\mathbf{X}$ denotes a random variable which is uniformly distributed over $\mathbb{F}_2^n$. The random variable $\mathbf{Y}$ denotes the image of $\mathbf{X}$ under Enc. We know that the boomerang returns if and only if both δ differences are mapped to the same difference in the decryption direction. Using this in the second step and rules 1 and 2 in subsequent steps, we get:

$$\Pr\left(\text{Boomerang returns}\right) \tag{17}$$

$$= \Pr_{\mathbf{X}}\Big(\text{Enc}^{-1}\big(\text{Enc}(\mathbf{X}) \oplus \delta\big) \oplus \text{Enc}^{-1}\big(\text{Enc}(\mathbf{X} \oplus \alpha) \oplus \delta\big) = \alpha\Big) \tag{18}$$

$$= \sum_{\eta \in \mathbb{F}_2^n} \Pr_{\mathbf{X}}\left(\delta \xrightarrow[\text{Enc}(\mathbf{X})]{\text{Enc}^{-1}} \eta \;\text{ and }\; \delta \xrightarrow[\text{Enc}(\mathbf{X}\oplus\alpha)]{\text{Enc}^{-1}} \eta\right) \tag{19}$$

Using the fact that $\text{Enc}(\mathbf{X} \oplus \alpha) = \mathbf{Y} \oplus \varepsilon$ whenever $\varepsilon \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \alpha$, we continue with

$$= \sum_{\substack{\eta \in \mathbb{F}_2^n \\ \varepsilon \in \mathbb{F}_2^n}} \Pr_{\mathbf{Y}}\left(\delta \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \eta \;\text{ and }\; \delta \xrightarrow[\mathbf{Y}\oplus\varepsilon]{\text{Enc}^{-1}} \eta \;\text{ and }\; \varepsilon \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \alpha\right) \tag{20}$$

Using the third expression, we can now simplify the second by changing the anchor

$$= \sum_{\substack{\eta \in \mathbb{F}_2^n \\ \varepsilon \in \mathbb{F}_2^n}} \Pr_{\mathbf{Y}}\left(\delta \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \eta \;\text{ and }\; \delta \oplus \varepsilon \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \eta \oplus \alpha \;\text{ and }\; \varepsilon \xrightarrow[\mathbf{Y}]{\text{Enc}^{-1}} \alpha\right) \tag{21}$$

Applying Rule 1 to all three subexpressions

$$= \sum_{\substack{\eta \in \mathbb{F}_2^n \\ \varepsilon \in \mathbb{F}_2^n}} \Pr_{\mathbf{X}}\left(\eta \xrightarrow[\mathbf{X}]{\text{Enc}} \delta \;\text{ and }\; \eta \oplus \alpha \xrightarrow[\mathbf{X}]{\text{Enc}} \delta \oplus \varepsilon \;\text{ and }\; \alpha \xrightarrow[\mathbf{X}]{\text{Enc}} \varepsilon\right) \tag{22}$$

Collecting everything into 3-differences

$$= \sum_{\substack{\eta \in \mathbb{F}_2^n \\ \varepsilon \in \mathbb{F}_2^n}} \Pr_{\mathbf{X}}\left((\alpha, \eta, \eta \oplus \alpha) \xrightarrow[\mathbf{X}]{\text{Enc}} (\varepsilon, \delta, \delta \oplus \varepsilon)\right) \tag{23}$$

$$= \sum_{\eta \in \mathbb{F}_2^n} \Pr_{\mathbf{X}}\left((\alpha, \eta, \eta \oplus \alpha) \xrightarrow[\mathbf{X}]{\text{Enc}} B\right) \tag{24}$$

$$= 2^n \cdot \Pr\left(A \xrightarrow{\text{Enc}} B\right) \tag{25}$$

which concludes the proof. □

The sum (23) contains a single term where $\varepsilon = \delta$ and $\eta = \alpha$ which corresponds to the probability of the ordinary differential $\alpha \xrightarrow{\text{Enc}} \delta$. This yields the following:

Corollary 1. *The probability of the return of the boomerang defined in Theorem 2 is greater than or equal to the probability of the ordinary differential with input difference α and output difference δ.*

We would like to point out that we did not need to use the Assumption 2 for the proof of Theorem 2. Using that assumption now though, we can make statements about the case where we split the encryption function into two parts E_0 and E_1.

Theorem 3. *Let A and B be again as in Eqs. (14) and (15). The probability of the boomerang to follow the differentials $\alpha \xrightarrow{E_0} \beta$ and $\gamma \xrightarrow{E_1} \delta$ for the respective text pairs in the upper and lower halves is then equal to*

$$\Pr\left((\beta, \gamma, \beta \oplus \gamma) \xrightarrow{E_0^{-1}} A\right) \cdot \Pr\left((\beta, \gamma, \beta \oplus \gamma) \xrightarrow{E_1} B\right) \ . \tag{26}$$

Proof. Along similar lines as the proof for Theorem 2. □

Comparing with the classical estimate of the boomerang probability, we see that we classically estimate that the 3-difference $(\beta, \gamma, \beta \oplus \gamma)$ is mapped by E_0^{-1} to a 3-difference in A as p^2 and likewise estimate the probability for this 3-difference to be mapped by E_1 into B as q^2.

How well do these approximations hold? The following lemma sheds some light on that:

Lemma 1. *The average of the probability for a 3-difference $(\beta, \gamma, \beta \oplus \gamma)$ to be mapped by a function f to a 3-difference of type $(\alpha, \eta, \alpha \oplus \eta)$ for some $\eta \in \mathbb{F}_2^n$ over all $\gamma \in \mathbb{F}_2^n$ is equal to the square of the probability of the differential $\beta \xrightarrow{f} \alpha$:*

$$2^{-n} \sum_{\gamma, \eta \in \mathbb{F}_2^n} \Pr\left((\beta, \gamma, \beta \oplus \gamma) \xrightarrow{f} (\alpha, \eta, \alpha \oplus \eta)\right) = \left(\Pr\left(\beta \xrightarrow{f} \alpha\right)\right)^2 . \tag{27}$$

Proof. Let $\mathbf{X}$ and $\mathbf{Y}$ denote two independent, uniformly distributed random variables on $\mathbb{F}_2^n$. We then have

$$2^{-n} \sum_{\gamma, \eta \in \mathbb{F}_2^n} \Pr\left((\beta, \gamma, \beta \oplus \gamma) \xrightarrow{f} (\alpha, \eta, \alpha \oplus \eta)\right) \tag{28}$$

$$= 2^{-n} \sum_{\gamma, \eta \in \mathbb{F}_2^n} \Pr_{\mathbf{X}}\left(\beta \xrightarrow[\mathbf{X}]{f} \alpha \text{ and } \gamma \xrightarrow[\mathbf{X}]{f} \eta \text{ and } \beta \oplus \gamma \xrightarrow[\mathbf{X}]{f} \alpha \oplus \eta\right) \tag{29}$$

$$= 2^{-n} \sum_{\gamma \in \mathbb{F}_2^n} \Pr_{\mathbf{X}}\left(\beta \xrightarrow[\mathbf{X}]{f} \alpha \text{ and } \beta \xrightarrow[\mathbf{X} \oplus \gamma]{f} \alpha\right) \tag{30}$$

$$= \Pr_{\mathbf{X}, \mathbf{Y}}\left(\beta \xrightarrow[\mathbf{X}]{f} \alpha \text{ and } \beta \xrightarrow[\mathbf{Y}]{f} \alpha\right) \tag{31}$$

$$= \left(\Pr\left(\beta \xrightarrow{f} \alpha\right)\right)^2 \tag{32}$$

which concludes the proof. □

This lemma could be described as stating that Assumption 1 holds on average over the possible differences in the middle of the boomerang. It should be stressed though that the actual differences used in the middle layer of a boomerang attack are fixed; the actual probability of a boomerang can thus deviate strongly from this average.

A direct consequence of Lemma 1 is the following theorem, which also appears as Proposition 1 in [Nyb19].

Theorem 4. *Let a boomerang be given as defined in Theorem 2. The average probability of the return of the boomerang taken over the ciphertext differences δ is equal to the probability that two randomly chosen plaintext pairs with difference α have equal ciphertext differences.*

The probability of two random pairs of texts, both having the same difference α, to be mapped to same differences is closely related to the non-uniformity of the probability distributions of the ciphertext differences resulting from the plaintext difference α. The extreme examples are the APN functions which have the most uniform possible distribution of ciphertext differences with half of the probabilities $\Pr\left(\alpha \xrightarrow{\text{Enc}} w\right)$ equal to 2^{1-n} and the other half equal to zero. It is easy to see that for APN functions the average probability of a boomerang to return is equal to 2^{1-n}. Due to a lack of suitable APN bijections, practical constructions use differentially 4-uniform S-boxes instead. Interestingly, they have two-valued probabilities for propagation of 3-differences as shown in the following lemma, which also follows from the proof of Proposition 4 in [BC18]:

Lemma 2. *Let $f : \mathbb{F}_2^n \to \mathbb{F}_2^n$ be a differentially 4-uniform bijection. Let $\alpha \xrightarrow{f} \beta$ be a differential of probability p over f and let $A = \left\{(\alpha, \eta, \alpha \oplus \eta) \in \mathbb{F}_2^{3n} \;\middle|\; \eta \in \mathbb{F}_2^n\right\}$. Then the probability that $(\beta, \gamma, \beta \oplus \gamma) \xrightarrow[y]{f^{-1}} A$ is either p or 0 for any $\gamma, y \in \mathbb{F}_2^n$.*

Proof. As f is differentially 4-uniform, we know that there are at most four different values of y such that $\beta \xrightarrow[y]{f^{-1}} \alpha$ holds. Let y_0, y_1, y_2, and y_3 be such four values and let us assume without loss of generality that $y_0 \oplus y_1 = y_2 \oplus y_3 = \beta$. Now for some fixed y, let us look at the probability that $(\beta, \gamma, \beta \oplus \gamma) \xrightarrow[y]{f^{-1}} A$ holds. This is then equivalent to both $\beta \xrightarrow[y]{f^{-1}} \alpha$ and $\beta \xrightarrow[y \oplus \gamma]{f^{-1}} \alpha$ holding simultaneously. Now let us suppose that γ is one of the values 0, β, $y_0 \oplus y_2$, or $y_0 \oplus y_2 \oplus \beta$. Then clearly both differentials hold if and only if y is one of y_0, y_1, y_2, and y_3. On the converse if γ is not equal to any of the above values, then both differentials can never hold simultaneously. Thus the probability that both hold is either 0 or p. □

In words, the probability that two randomly chosen data pairs of the same difference follow the same differential strongly depends on the difference between the pairs. We now consider SPNs, that is, a cipher where the round function consists of a non-linear substitution layer comprised of S-boxes applied in parallel, a linear permutation layer, and a key addition. Lemma 2 then allows us to make the following statement:

Theorem 5. *Let* Enc *be a 2-round SPN using* 4*-uniform S-boxes in the substitution layer. Let* $\text{Enc} = E_1 \circ E_0$ *be such that E_0 and E_1 correspond to one round of the cipher each. Let further $\alpha \xrightarrow{E_0} \beta$ and $\gamma \xrightarrow{E_1} \delta$ be two differentials with probabilities p and q respectively. Then the probability of the boomerang constructed from these differentials is either 0 or pq.*

Proof. We determine the probability using Eq. (26). Let us evaluate the first factor. As E_0 consist only of an S-box layer and an affine layer, it is straightforward to see that $(\beta, \gamma, \beta \oplus \gamma) \xrightarrow[y]{f^{-1}} A$ holds if and only if the respective 3-difference transitions over all single S-boxes hold. But as the S-boxes are 4-uniform, we know from Lemma 2 that the probabilities of the transitions over the S-boxes are either 0 or correspond to the probability of a single differential over the S-box. This property is thus lifted to the complete round and we thus know that the first factor in Eq. (26) thus evaluates either to 0 or p. The argument for the second factor is analogous. □

We should note that it is enough for the boomerang to have a probability-zero transition in one S-box to set the total probability to zero. As a consequence any randomly chosen simple boomerang over such a cipher has a high probability of having a probability of zero. We give examples of this behavior on some well-known boomerang attacks in Section 7.

6 Comparison to the boomerang connectivity table (BCT)

The boomerang connectivity table (BCT) [CHP+18] is a tool that has been applied successfully in recent years. We would thus like to give a quick comparison between the techniques used in this paper and the BCT.

A BCT allows us to determine the probability that two trails connect successfully over an S-box when this S-box corresponds to the middle layer of a sandwich attack. The BCT for a given input difference α and output difference δ is defined as the number of right quartets over one n-bit S-box S:

$$\mathrm{BCT}(\alpha, \delta) := \left|\left\{x \in \mathbb{F}_2^n \;\middle|\; S^{-1}(S(x) \oplus \delta) \oplus S^{-1}(S(x \oplus \alpha) \oplus \delta) = \alpha\right\}\right|$$

Using Theorem 2, we can formulate this with our framework.

Theorem 6. *Given an S-box S, an input difference α and output difference δ, we can state the* BCT entry as

$$\mathrm{BCT}(\alpha, \delta) = \sum_{\substack{\eta \in \mathbb{F}_2^n \\ \varepsilon \in \mathbb{F}_2^n}} (\alpha, \eta, \alpha \oplus \eta) \xrightarrow{S} (\varepsilon, \delta, \varepsilon \oplus \delta)$$

We would like to point out that in contrast to Lemma 2, we need to sum over both η and ε. Thus the probabilities in the BCT are not limited by the highest entries in the difference distribution table (DDT). Thus ensuring that the trails do not connect directly in the middle but leave a middle round as in the sandwich attack can have a positive effect on the probability of the attack.

7 Boomerang attacks on Serpent

In this section, we will take a closer look at two of the most well-known applications of boomerang attacks, namely the amplified boomerang attack [KKS00] and the rectangle attack [BDK01] on the block cipher Serpent.

As an SPN with differentially 4-uniform S-boxes, Serpent is a prime test candidate for Theorem 5. Before we go into more detail into the specific boomerang attacks, let us first make one observation. Any simple boomerang constructed from a differential trail for the top part and another trail for the bottom part contains a 2-round boomerang at its core. This makes it necessary to take Theorem 5 into consideration also when the boomerang covers more than two rounds.

7.1 Short overview on Serpent

The block cipher Serpent [BAK98] was an AES candidate and ranked second in the final evaluation. Serpent is constructed as a 32-round substitution-permutation network (SPN) and has been designed to offer a very effective bit-sliced implementation. One round of Serpent consists of an S-box layer in which the same four-bit S-box is applied to all four-bit nibbles of the state followed by an affine layer and a round key addition. We provide a full description of Serpent in Appendix A to make the paper self-contained.

7.2 Amplified boomerang attack [KKS00]

In the amplified boomerang attack on Serpent [KKS00], a simple boomerang distinguisher on seven rounds (rounds one to seven) with a single differential trail for the top and a single differential trail for the bottom is used. The top part (E_0 in the paper) consists of rounds one to four while the bottom part (E_1 in the paper) consists of rounds five to

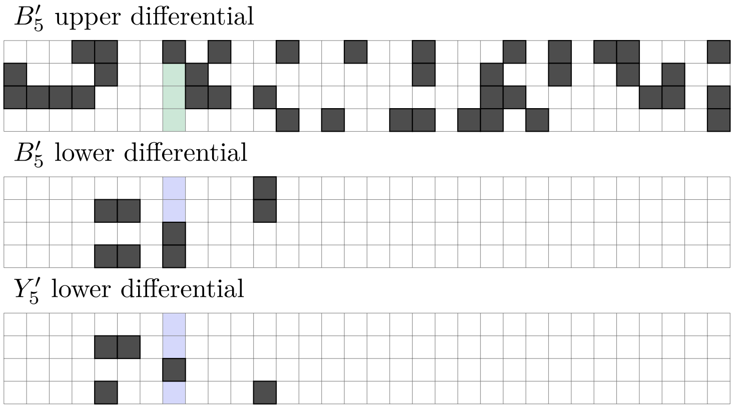

B_5' upper differential

B_5' lower differential

Y_5' lower differential

Figure 3: Differences involved in our analysis of the amplified boomerang attack on Serpent [KKS00]. Black bits are active, and the green and blue column marks S-box 24 in the upper and lower differential respectively.

seven. When taking a closer look at the probability of the inner two-round boomerang over rounds four and five, it is straightforward to see that there are several probability-zero transitions over the S-boxes.

Our approach here is to look at the first S-box transition of the lower differential, $B_5' \to Y_5'$ (see Fig. 3). We then look at the pairs that follow the differences for each S-box and check if two pairs have the difference from the upper B_5'. If there are no pairs with this requirement then the transition for that S-box is impossible and therefore the boomerang has probability 0. A similar analysis can be done using the BCT by checking the entry for upper B_5' to lower Y_5'. In this case the lower B_5' is disregarded as the boomerang is turned into a sandwich.

As an example consider S-box 24 in round 5 for which the attack requires two pairs such that $S(x) \oplus S(x \oplus c) = 4$ with a distance of 1. The only values for x are $\{4, 7, 8, b\}$, but since none of these have a distance of 1, the transition for this S-box is impossible. Any one of these is sufficient to give the inner and thus the complete boomerang a probability of zero. We can also see this from the BCT (Table 2) as BCT(1,4) = 0. This renders the attack invalid.

7.3 Rectangle attack [BDK01]

Let us now have a look at a more refined boomerang attack on Serpent, namely the rectangle attack [BDK01]. In this attack, a boomerang distinguisher is used on rounds one to eight and these rounds are split into two parts of four rounds each. Instead of using only one differential trail for each part, a set of different trails for both parts are used that share the same input or output difference (see Eq. (2) for the classical estimate). For the upper part $2^{13.4}$ trails are used. While the number of trails for the lower part is unspecified in the paper, the probability estimate from these is given. By considering the best $2^{40.0}$ trails of the trail type used in the paper, we get a classical estimate that is slightly higher than the one stated in the paper and we thus assume that these trails are a superset of the trails used in the original paper.[2] In total this leaves us with $2^{53.3}$ combinations of trails for the lower and upper part, giving us a classical estimate of $2^{-119.3}$ for the boomerang probability (this is slightly higher than the original estimate of $2^{-120.6}$ in [BDK01]).

For all of these trail combinations, we calculated the accurate transition probabilities

[2] For this estimate we consider all trails from round five to eight of the type specified in the Appendix of [BDK01] which activate at most 12 S-boxes in round five.

of the two-round inner boomerangs (rounds four and five). For the vast majority of trail combinations, the inner boomerang has a probability of zero, such that only 972 of the $2^{53.3}$ trail combinations are left with a non-zero probability. While this would leave the attack completely invalid if we stuck with the classical estimate (an estimate of $2^{-164.3}$ to be precise), we can apply the results of Theorem 5 positively on these remaining 972 trail combinations: as their inner boomerang has a non-zero probability, it has a much higher probability than the classical estimate would suggest. Interestingly combining the accurate probability for the inner boomerang with the classical estimates for the outer rounds of the boomerang, leaves us not only with a probability estimate that is close to the classical one, it even gives us a slightly higher one, namely $2^{-118.6}$.

Unfortunately when we try to correctly determine the probability of a boomerang that exceeds two rounds, we are faced with the problem of the exponential growth in the number of trails that need to be considered. This is particularly true for the rectangle attack on Serpent where we were not able to apply our methodology to more than the inner two rounds. However, Lemma 1 gives us reason to assume that it is justified to apply the classical estimate for the outer rounds, in particular when the diffusion of the cipher is prohibitively large for more accurate methods. As we are in that case considering a large number of distances between the text pairs, it seems acceptable to assume that they reasonably closely estimate the average probability. And this average (as in Eq. (27)) corresponds exactly to the classical estimate.

To improve the quality of the boomerang, we also evaluated the effect of considering more trails for the lower part. We found that when we allowed all trails of the type used in the original attack that activate at most 15 S-boxes in round five, the probability of the boomerang distinguisher improved to $2^{-116.3}$ (as opposed to $2^{-119.0}$ for the classical estimate with the same number of trails). This improves the estimate of the boomerang probability in comparison to the original estimate by a factor of $2^{4.3}$.

8 Summary and conclusion

In this paper, we took a close look at boomerang attacks and the classical estimate of their probability. We explicitly stated the assumption underlying the classical estimates of boomerang probabilities and showed that an inherent contradiction arises when we take this assumption for granted. Using the notion of d-differences and their transitions, a generalization of differential cryptanalysis, we were able to express the probability of boomerang distinguishers precisely in a model that only relies on the independence of rounds instead of the independence of differentials. We then used this formalization to prove a number of results.

One of the most important results is that we could rigorously prove that two-round boomerangs on SPN ciphers with differentially 4-uniform S-boxes—including ciphers such as AES, Serpent, or PRESENT—deviate strongly from their expected classical probability estimate. This results in a very high likelihood for boomerangs that only make use of two differentials to have probability zero, even when covering more than two rounds. On the other hand, this also allows cleverly constructed boomerangs to beat the classical estimate by a large margin.

As an application of these results, we took a closer look at two classical applications of boomerangs on Serpent. For the first attack [KKS00], we found that the boomerang, as constructed from two differentials, has in fact probability zero. For the second attack [BDK05], we found that although only a fraction of $2^{-43.4}$ of all possible considered trail combinations had a non-zero probability, the total estimate of the boomerang was hardly altered as the remaining trails showed a much higher probability than the classical estimate would suggest. As a matter of fact by including some more trails, we were able to improve the classical estimate of the boomerang by a factor of $2^{4.3}$.

How come that the probability estimate for the rectangle attack was so little influenced by the vast amount of probability-zero trail combinations? The explanation lies in the very large number of trails with comparable probabilities used in the attack. This allowed the classical estimate which only holds on average (as proven in Lemma 1) to describe the probability quite accurately.

From all this we conclude the following: probability estimates in boomerang attacks must be handled with the care. Simply combining two differentials to construct a boomerang can easily lead to the boomerang probability being zero, rendering the attack invalid. Detailed arguments and computer validations of the probability, where possible, should be a minimal requirement for all future boomerang attacks.

Acknowledgments

The authors would like to thank the anonymous reviewers and in particular Kaisa Nyberg for excellent feedback and comments that considerably improved the quality of the paper.

References

[BAK98] Eli Biham, Ross J. Anderson, and Lars R. Knudsen. Serpent: A new block cipher proposal. In Serge Vaudenay, editor, *Fast Software Encryption, 5th International Workshop, FSE '98, Paris, France, March 23-25, 1998, Proceedings*, volume 1372 of *Lecture Notes in Computer Science*, pages 222–238. Springer, 1998.

[BC18] Christina Boura and Anne Canteaut. On the boomerang uniformity of cryptographic sboxes. *IACR Trans. Symmetric Cryptol.*, 2018(3):290–310, 2018.

[BDK01] Eli Biham, Orr Dunkelman, and Nathan Keller. The rectangle attack - rectangling the serpent. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 340–357. Springer, 2001.

[BDK05] Eli Biham, Orr Dunkelman, and Nathan Keller. Related-key boomerang and rectangle attacks. In Ronald Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 507–525. Springer, 2005.

[BHL$^+$20] Hamid Boukerrou, Paul Huynh, Virginie Lallemand, Bimal Mandal, and Marine Minier. On the Feistel counterpart of the boomerang connectivity table introduction and analysis of the FBCT. *IACR Trans. Symmetric Cryptol.*, 2020(1):331–362, 2020.

[BK09] Alex Biryukov and Dmitry Khovratovich. Related-key cryptanalysis of the full AES-192 and AES-256. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2009.

[BS91] Eli Biham and Adi Shamir. Differential cryptanalysis of DES-like cryptosystems. *J. Cryptol.*, 4(1):3–72, 1991.

[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, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 683–714. Springer, 2018.

[DDV20] Stéphanie Delaune, Patrick Derbez, and Mathieu Vavrille. Catching the fastest boomerangs application to SKINNY. *IACR Trans. Symmetric Cryptol.*, 2020(4):104–129, 2020.

[DKRS20] Orr Dunkelman, Nathan Keller, Eyal Ronen, and Adi Shamir. The retracing boomerang attack. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 280–309. Springer, 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, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 393–410. Springer, 2010.

[DKS14] Orr Dunkelman, Nathan Keller, and Adi Shamir. A practical-time related-key attack on the KASUMI cryptosystem used in GSM and 3g telephony. *J. Cryptol.*, 27(4):824–849, 2014.

[HBS21] Hosein Hadipour, Nasour Bagheri, and Ling Song. Improved rectangle attacks on SKINNY and CRAFT. *IACR Trans. Symmetric Cryptol.*, 2021(2):140–198, 2021.

[JK97] Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In Eli Biham, editor, *Fast Software Encryption, 4th International Workshop, FSE '97, Haifa, Israel, January 20-22, 1997, Proceedings*, volume 1267 of *Lecture Notes in Computer Science*, pages 28–40. Springer, 1997.

[KKS00] John Kelsey, Tadayoshi Kohno, and Bruce Schneier. Amplified boomerang attacks against reduced-round MARS and Serpent. In Bruce Schneier, editor, *Fast Software Encryption, 7th International Workshop, FSE 2000, New York, NY, USA, April 10-12, 2000, Proceedings*, volume 1978 of *Lecture Notes in Computer Science*, pages 75–93. Springer, 2000.

[LGW12] Shusheng Liu, Zheng Gong, and Libin Wang. Improved related-key differential attacks on reduced-round LBlock. In Tat Wing Chim and Tsz Hon Yuen, editors, *Information and Communications Security - 14th International Conference, ICICS 2012, Hong Kong, China, October 29-31, 2012. Proceedings*, volume 7618 of *Lecture Notes in Computer Science*, pages 58–69. Springer, 2012.

[LMM91] Xuejia Lai, James L. Massey, and Sean Murphy. Markov ciphers and differential cryptanalysis. In Donald W. Davies, editor, *Advances in Cryptology - EUROCRYPT '91, Workshop on the Theory and Application of of Cryptographic Techniques, Brighton, UK, April 8-11, 1991, Proceedings*, volume 547 of *Lecture Notes in Computer Science*, pages 17–38. Springer, 1991.

[Mur11] Sean Murphy. The return of the cryptographic boomerang. *IEEE Transactions on Information Theory*, 57(4):2517–2521, 2011.

[NK95] Kaisa Nyberg and Lars R. Knudsen. Provable security against a differential attack. *J. Cryptol.*, 8(1):27–37, 1995.

[Nyb19] Kaisa Nyberg. The extended autocorrelation and boomerang tables and links between nonlinearity properties of vectorial Boolean functions. *IACR Cryptol. ePrint Arch.*, page 1381, 2019.

[SQH19] Ling Song, Xianrui Qin, and Lei Hu. Boomerang connectivity table revisited. application to SKINNY and AES. *IACR Trans. Symmetric Cryptol.*, 2019(1):118–141, 2019.

[Tie16] Tyge Tiessen. Polytopic cryptanalysis. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 214–239. Springer, 2016.

[Vau98] Serge Vaudenay. Provable security for block ciphers by decorrelation. In Michel Morvan, Christoph Meinel, and Daniel Krob, editors, *STACS 98, 15th Annual Symposium on Theoretical Aspects of Computer Science, Paris, France, February 25-27, 1998, Proceedings*, volume 1373 of *Lecture Notes in Computer Science*, pages 249–275. Springer, 1998.

[Wag99] David A. Wagner. The boomerang attack. In Lars R. Knudsen, editor, *Fast Software Encryption, 6th International Workshop, FSE '99, Rome, Italy, March 24-26, 1999, Proceedings*, volume 1636 of *Lecture Notes in Computer Science*, pages 156–170. Springer, 1999.

[WP19] Haoyang Wang and Thomas Peyrin. Boomerang switch in multiple rounds. application to AES variants and Deoxys. *IACR Trans. Symmetric Cryptol.*, 2019(1):142–169, 2019.

A Description of Serpent

Here we give a description of the parts of Serpent needed to follow the analysis in Section 7. Serpent has a state of 128 bits split into 4 words of 32 bits to allow for an efficient bit-sliced implementation. The convention for Serpent is to use B_i as the state before round i, so B_0 is the plaintext and B_{32} is the ciphertext. Encryption then proceeds as follows:

$$\begin{aligned} Y_i &= S_i(B_i \oplus K_i) & & \\ B_{i+1} &= L(Y_i) & & i = 0, \dots, 30 \\ B_{i+1} &= Y_i \oplus K_{i+1} & & i = 31 \end{aligned}$$

Here L is the linear layer which in the bit sliced version is described as follows. If we call the state words for X_0, X_1, X_2, X_3 after the key addition and the S-box layer then the linear layer can be described as:

$$\begin{aligned}
X_0, X_1, X_2, X_3 &= S_i(B_i \oplus K_i) \\
X_0 &= X_0 <<< 13 \\
X_2 &= X_2 <<< 3 \\
X_1 &= X_1 \oplus X_0 \oplus X_2 \\
X_3 &= X_3 \oplus X_2 \oplus (X_0 << 3) \\
X_1 &= X_1 <<< 1 \\
X_3 &= X_3 <<< 7 \\
X_0 &= X_0 \oplus X_1 \oplus X_3 \\
X_2 &= X_2 \oplus X_3 \oplus (X_1 << 7) \\
X_0 &= X_0 <<< 5 \\
X_2 &= X_2 <<< 22 \\
B_{i+1} &= X_0, X_1, X_2, X_3
\end{aligned}$$

Here $<<<$ is a left rotation and $<<$ is a left shift.

Serpent uses 8 different S-boxes such that round i uses S-box $i \mod 8$. The S-boxes are applied to 1 bit from each word and the same S-box is used for all bits. The bit from X_0 is the least significant and X_3 is the most significant bit. This allows them to be applied in a bit-sliced fashion. The S-boxes are provided in Table 1 and Table 2 gives the BCT of S5 which is used in out analysis in Section 7.

The keyschedule is not relevant for the analysis so we will leave out the description.

Table 1: S-boxes used in Serpent

S0:	3	8	15	1	10	6	5	11	14	13	4	2	7	0	9	12
S1:	15	12	2	7	9	0	5	10	1	11	14	8	6	13	3	4
S2:	8	6	7	9	3	12	10	15	13	1	14	4	0	11	5	2
S3:	0	15	11	8	12	9	6	3	13	1	2	4	10	7	5	14
S4:	1	15	8	3	12	0	11	6	2	5	4	10	9	14	7	13
S5:	15	5	2	11	4	10	9	12	0	3	14	8	13	6	7	1
S6:	7	2	12	5	8	4	6	11	14	9	1	15	13	3	10	0
S7:	1	13	15	0	14	8	2	11	7	4	12	10	9	3	5	6

Table 2: BCT of S5

	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
0	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16	16
1	16	00	00	02	00	02	04	00	00	06	02	02	00	00	02	04
2	16	00	00	00	00	04	06	02	00	00	02	10	00	04	04	00
3	16	02	00	02	06	00	00	02	04	00	00	00	06	02	00	00
4	16	00	00	00	00	02	08	02	00	04	00	04	00	06	00	06
5	16	06	02	00	00	02	06	00	00	04	00	00	00	00	04	08
6	16	02	02	02	00	00	04	02	00	02	00	04	00	04	02	00
7	16	02	00	02	06	02	00	00	04	00	00	00	06	00	00	02
8	16	00	00	02	04	00	02	00	04	02	00	00	04	02	02	02
9	16	00	02	00	00	04	00	02	02	00	02	06	02	00	04	00
a	16	02	02	02	02	00	00	00	02	00	02	02	00	02	00	00
b	16	04	00	00	00	06	00	02	00	00	00	04	00	02	08	06
c	16	00	02	02	04	02	00	02	04	02	02	00	04	00	00	00
d	16	02	02	02	00	00	00	02	02	02	02	00	02	00	00	00
e	16	04	02	00	02	00	00	00	02	02	02	00	00	02	04	04
f	16	00	02	00	08	00	02	00	08	00	02	00	08	00	02	00