Integral and Multidimensional Linear Distinguishers with Correlation Zero

Andrey Bogdanov[1⋆], Gregor Leander[2⋆], Kaisa Nyberg[3⋆], Meiqin Wang[4⋆]

[1] KU Leuven, ESAT/SCD/COSIC and IBBT, Belgium
[2] Technical University of Denmark, Denmark
[3] Aalto University, Finland
[4] Shandong University, Key Laboratory of Cryptologic Technology and Information Security, Ministry of Education, Shandong University, Jinan 250100, China

Abstract. Zero-correlation cryptanalysis uses linear approximations holding with probability exactly 1/2. In this paper, we reveal fundamental links of zero-correlation distinguishers to integral distinguishers and multidimensional linear distinguishers. We show that an integral implies zero-correlation linear approximations and that a zero-correlation linear distinguisher is actually a special case of multidimensional linear distinguishers. These observations provide new insight into zero-correlation cryptanalysis which is illustrated by attacking a Skipjack variant and round-reduced CAST-256 without weak key assumptions.
Keywords: zero-correlation cryptanalysis, integral distinguishers, multidimensional linear distinguishers, Skipjack, CAST-256

1 Introduction

1.1 Zero-correlation

Zero-correlation cryptanalysis [7, 8] is a novel promising attack technique for block ciphers. The distinguishing property used in zero-correlation cryptanalysis is the existence of *zero-correlation linear approximations* over (a part of) the cipher. Those are linear approximations that hold true with a probability p of exactly $1/2$, that is, strictly unbiased approximations having a *correlation* $c = 2p - 1$ equal to 0.

The original work [7] provides a simple and efficient technique to find zero-correlation approximation but the distignuisher was rather weak. Recently, the work [8] proposed a more powerful distinguisher was proposed by exploiting the fact that zero-correlation approximations are numerous in susceptible ciphers. Though working fine in practice and being useful in cryptanalysis, the distinguisher of [8] has some constraints that we would like to overcome: (1) If there are ℓ zero-correlation linear approximations for an n-bit block cipher, the distinguisher of [8] has to make $\mathcal{O}(2^n/\sqrt{\ell})$ queries. So the data complexity does not go down as fast as ℓ grows. (2) The distinguisher of [8] relies on the assumption that all linear approximations with correlation zero are independent.

⋆ All authors are corresponding authors.

In most cases, including the attacks of [8] in fact, this assumption is formally not met, since all classes of zero-correlation approximations known so far are actually truncated, building linear spaces of dimension $\log_2 \ell$. That is, almost all ℓ approximations used will be linearly dependent, formally jeopardizing the assumption and another theory is needed to support the zero-correlation.

1.2 Our contributions

Zero-correlation and integrals. Integral distinguishers were originally proposed by Knudsen as a dedicated attack against the Rijndael-predecessor Square [12]. Integral distinguishers [21] are also known as square distinguishers for this reason, especially when applied to Square-type ciphers such as AES. Variants of integral distinguishers include saturation [23] and multiset distinguishers [5]. Integral distinguishers mainly make use of the observation that it is possible to fix some parts of the plaintext such that specific parts of the ciphertext are balanced, i.e. each possible partial value occurs the exact same number of times in the output.

In this paper, we demonstrate that an integral implies zero-correlation linear approximations, see Fig. 1. In the other direction, a zero-correlation distinguisher implies an integral distinguisher only if input and output linear masks in zero-correlation approximations are independent of each other. Note that the condition for the input and output masks to be detached from each other implies that, for instance, the 5-round zero-correlation property of balanced Feistel ciphers of [7] is not directly described by an integral.

In this sense, the fact the integrals imply zero-correlation distinguishers is especially intriguing as not only the ways the distinguishers are constructed are different but also the ways the resulting attacks work seem inherently different. In particular, this link allows using ℓ input masks and one output mask with correlation zero in a distinguisher with a data complexity of $2^n/\ell$. Thus, in these settings the above outlined link allows to reduce the data complexity of zero-correlation distinguishers by a factor of $\sqrt{\ell}$ (at the price of transforming the attack into a chosen-plaintext attack) compared to previous works.

Zero-correlation and multidimensional linear distinguishers. The basic idea of multidimensional cryptanalysis [1, 4, 13, 15, 17, 18] is that, given correlations of all linear approximations with non-zero correlation on a linear space formed by some cipher data, the probability distribution of the cipher data can be determined. Then, instead of the statistical behavior of a large set of mutually dependent linear approximations, one can examine the data distribution. Indeed, statistical behavior of multiple linear approximations has been analyzed only under the assumption of statistical independence [4]. The main advantage of the multidimensional approach is that it allows rigorous statistical analysis of linear approximations without the independence assumption. In traditional linear cryptanalysis, the focus is on linear approximations with correlations of large magnitude. The larger are the magnitudes of correlations, the more non-uniform is the distribution of the cipher data under consideration. The linear

distinguisher is then based on distinguishing the nonuniform cipher data distribution from an uniform distribution. For a more comprehensive recent survey on multidimensional linear distinguishers, the reader is referred to e.g. [16].

In this paper, we consider linear spaces of cipher data where correlations of all linear approximations are equal to zero. Our starting observation here is that in fact, being truncated, *zero-correlation approximations constitute a special case of multidimensional linear approximations*. However, unlike traditional multidimensional linear distinguishers where the cipher data behaves non-uniformly, the cipher data for zero-correlation is uniformly distributed. This requires the development of a statistical theory to distinguish a sample of such cipher data from a sample of random data drawn from an uniform distribution.

In contrast to [8], the new distinguisher does not need the assumption of the statistical independence for multiple zero-correlation linear approximations. While still requiring about $\mathcal{O}(2^n/\sqrt{\ell})$ cipher queries, it allows taking full advantage of all zero-correlation linear approximations available, independent or not. The distribution of the cipher data is accurately modeled as sampling from a multivariate hypergeometric distribution, while the random data is drawn from a multinomial distribution. This establishes an inherent link of zero-correlation to multidimensional linear distinguishers. In their essence, zero-correlation distinguishers constitute a special case of multidimensional linear-correlation distinguishers, see Fig. 1. We expect this technique to be useful in the cryptanalysis of many ciphers.

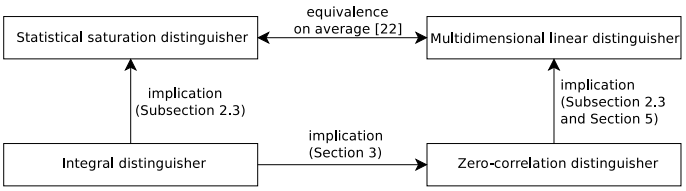

Fig. 1: Relations among distinguishers: zero-correlation, integral, statistical saturation, and multidimensional linear

Applications: Attacks on Skipjack variant and CAST-256. To emphasize the practical meaningfulness of our findings, we apply the new distinguishers to mount key recovery attacks on block ciphers.

Skipjack is the only block cipher known to be designed by NSA. It is a 32-round 4-line unbalanced Feistel-type network based on interleaving two types of round functions – Rule A and Rule B. The best known cryptanalytic result for Skipjack is the impossible differential cryptanalysis for 31 rounds given by Biham et al. [2] based on a 24-round impossible differential. We change the order of Rules A and B in Skipjack such that the longest impossible differential identified is over 21 rounds and show that it has a 30-round zero-correlation property. We can recover its key for 31 rounds with practical complexity using an integral zero-correlation attack.

CAST-256 was proposed as an AES candidate. It has 48 rounds. The best cryptanalysis so far in the classical single-key model without the weak-key assumption has been a linear attack on 24 rounds. We find 24-round zero-correlation linear approximations for CAST-256 and attack 28 rounds of CAST-256 using multidimensional zero-correlation cryptanalysis. At the same time, the longest impossible differential we are aware of is over 18 rounds (though there is an unspecified impossible differential for 20 rounds mentioned in the literature). Our multidimensional zero-correlation attack is the first attack on more than half of the full-round AES-candidate CAST-256 without the weak key assumption.

The remainder of the paper is organized as follows. In Section 2, we introduce some basic concepts and notions which will be useful throughout the paper. Section 3 establishes a strong link between the properties of integrals and zero-correlation approximations. Using an integral zero-correlation distinguisher, Section 4 cryptanalyzes a Skipjack variant resistant to impossible differential attack. Section 5 describes a link of zero-correlation approximations to multidimensional linear approximations and introduces a novel zero-correlation multidimensional linear distinguisher. Section 6 uses it to recover the key of 28 rounds of CAST-256. We conclude in Section 7.

2 Preliminaries

2.1 Linear approximations and balanced functions

$\mathbb{F}_2$ denotes the binary field of two elements and $\mathbb{F}_2^n$ is its extension of dimension n. Let x and $a \in \mathbb{F}_2^n$. Then $\langle a, x\rangle$ denotes their cannonical inner product on $\mathbb{F}_2^n$.

Given a function $H : \mathbb{F}_2^n \to \mathbb{F}_2^k$ the *correlation* c of the *linear approximation*

$$\langle b, H(x)\rangle + \langle a, x\rangle$$

for a k-bit output mask b and an n-bit input mask a is defined by

$$\Pr(\langle b, H(x)\rangle + \langle a, x\rangle = 0) = \frac{1+c}{2}$$

where the probability is taken over all choices of inputs x. A related measure for this correlation is *the Walsh- or Fourier-transformation*, defined as

$$\widehat{H}(a, b) = \sum_x (-1)^{\langle b, H(x)\rangle + \langle a, x\rangle}.$$

The fundamental relation between the Fourier transformation of H and the correlation of the linear approximation is given by

$$c = \frac{\widehat{H}(a, b)}{2^n}$$

and, thus, studying the correlation and studying the Fourier transformation are, up to scaling, equivalent.

We say a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^k$ is *balanced* if all preimages have identical size, i.e. if the size of the set

$$F^{-1}(y) := \{x \in \mathbb{F}_2^n \mid F(x) = y\}$$

is independent of y. Note that F being balanced implies $k \leq n$. We recall the following well-known characterization of balanced functions, see for example [10, Proposition 2]: A function $F : \mathbb{F}_2^n \to \mathbb{F}_2^k$ is balanced if and only if all its component functions are balanced, that is, if and only if for any non-zero $b \in \mathbb{F}_2^k$ it holds that $\widehat{F}(0, b) = 0$.

2.2 Decomposition of the target cipher

Assume that $H : \mathbb{F}_2^n \to \mathbb{F}_2^n$ is a (part of) cipher. To simplify notation and without loss of generality we split the inputs and outputs into two parts each.

$$H : \mathbb{F}_2^r \times \mathbb{F}_2^s \to \mathbb{F}_2^t \times F_2^u, H(x, y) = \begin{pmatrix} H_1(x, y) \\ H_2(x, y) \end{pmatrix}$$

Furthermore, the function T_λ defined by

$$T_\lambda : \mathbb{F}_2^s \to \mathbb{F}_2^t, T_\lambda(y) = H_1(\lambda, y)$$

will play a key role. The function T_λ is the function H when the first r bits of its input are fixed to λ and only the first t bits of the output are taken into account.

Table 1: Defining properties of some important distinguishers

Distinguisher	Defining property
multidimensional linear	$\sum_{a_1,b_1} \widehat{H}(a, b)^2$ non-random
statistical saturation	$\forall \lambda : \sum_{b_1} \widehat{T}_\lambda(0, b_1)^2$ non-random
integral	$\forall \lambda, b_1 : \widehat{T}_\lambda(0, b_1) = 0$
zero-correlation	$\forall a_1, b_1 : \widehat{H}(a, b) = 0$

2.3 Distinguishers and relations

Here we briefly outline the concepts behind four types of relevant distinguishers that we will be dealing with in this paper, which are also summarized in Table 1: *Zero-correlation distinguisher* uses the property that, for all input and output masks $a = (a_1, 0)$ and $b = (b_1, 0)$, the Fourier transformation of the cipher yields zero, $\widehat{H}(a, b) = 0$. *Integral distinguisher* is based on the property that, for all partial input fixations λ, the partial function of the cipher with this fixation is balanced in parts of its output. *Multidimensional linear distinguisher* relies upon the property that multiple Fourier coefficients of the cipher behave in a

non-random way, i.e. $\sum_{a_1,b_1} \widehat{H}(a,b)^2$ is non-random. *Statistical saturation distinguisher* builds upon the property that, for all partial input fixations λ, the partial function of the cipher with this fixation is non-random under Fourier transformation, i.e. $\sum_{b_1} \widehat{T}_\lambda(0,b_1)^2$ is non-random. While statistical saturation and multidimensional linear distinguishers concentrate on the cumulative properties holding for the partial Fourier spectra, integral and zero-correlation distinguishers deal with a set of individual properties of Fourier coefficients.

3 Zero-correlation and integral distinguishers

3.1 Conditional equivalence result

We start by stating the main result of this section, which is summarized in the following statement:

Proposition 1. *If the input and output linear masks a and b are independent, the approximation $\langle b, H(x)\rangle + \langle a, x\rangle$ has correlation zero for any $a = (a_1, 0)$ and any $b = (b_1, 0) \neq 0$ (zero-correlation) if and only if the function T_λ is balanced for any λ (integral).*

This basically means that, at least in terms of their defining properties, integral distinguishers imply zero-correlation distinguishers. The proof of Proposition 1 follows directly from the two lemmata below whose proofs are provided in the full version of this paper [6]. The tools used in the proofs mainly originate from results in the area of Boolean functions [22]. For instance, Lemma 2 is stated in different notation e.g. in [11, Proposition 9]).

The main technical tool is the next lemma linking the correlation of T_λ to the correlation of H.

Lemma 1. *With the notation from above, the following holds for any λ, b_1:*

$$2^s \widehat{T}_\lambda(0,b_1) = \sum_{a_1} (-1)^{\langle a_1, \lambda\rangle} \widehat{H}((a_1,0),(b_1,0)) \tag{1}$$

Lemma 1 already proves one direction of Proposition 1, namely, that zero-correlation approximations imply an integral under the condition that b_1 remains the same with the change of a_1. Lemma 1 is also especially useful for defining an integral distinguisher that is based on zero-correlation properties: Given a number of zero-correlation linear approximations (on the right-hand side of (1)), one checks if the corresponding partial function of the cipher is balanced (the left-hand side of (1)). This can be done for each partial input fixation λ separately.

The following direct corollary of Lemma 1 is even more telling and is the key in exhibiting the close link between zero-correlation distinguishers and integral distinguishers:

Lemma 2. *The following holds for any b_1:*

$$2^s \sum_{\lambda} \widehat{T}_\lambda(0,b_1)^2 = \sum_{a_1} \widehat{H}_1((a_1,0),b_1)^2$$

This lemma proves both directions of Proposition 1, including the fact that an integral implies zero-correlation distinguishers. In the sequel, we provide a more detailed description of the link and an example.

3.2 From zero-correlation to integral distinguishers (conditional)...

First, assume that $H : \mathbb{F}_2^n \to \mathbb{F}_2^n$ is a (part of) cipher vulnerable to zero-correlation attacks. More precisely, assume that for any $a = (a_1, 0)$ and any $b = (b_1, 0) \neq 0$ the relation $\langle b, H(x)\rangle + \langle a, x\rangle$ has correlation zero. We'd like to highlight two points here. The restriction to masks of the form $a = (a_1, 0)$ and $b = (b_1, 0)$, that is, to the masks where the last bits are fixed to zero, is solely for the simplicity of notations. However, the zero-correlation distinguishers considered here are of a special case: We assume not only that the used input and output masks form subspaces but also that this space of input and output masks is actually the direct product of the space of input masks and the space of output masks. Informally, the masks must not be coupled as they are for example in the attack on CAST-256 described in Section 6. We call such uncoupled input-output masks, for our equivalence result applies, *detached masks*.

Under those conditions, it follows from Lemma 2 above that $\widehat{T}_\lambda(0, b_1)$ equals zero for all $b_1 \neq 0$ and all λ. This yields that, for any λ the function T_λ mapping s bits to t bits is balanced. In other words, H exhibits the following integral distinguisher: Fixing the first s bits of H arbitrarily and encrypting all remaining 2^r possible plaintext, each possible t bits string occurs equally often in the first t bits of the output of H. In the particular case of $s = t$, the function T_λ is a permutation and, thus, each possible t-bit string should occur exactly once.

3.3 ...And back again (unconditional)

On the other hand, let us consider the case of a cipher that is vulnerable to an integral distinguisher in the following sense. Assume that, by fixing some (without loss of generality, the first s) bits in the input and encrypting all possible remaining plaintexts, one can identify a subset of t bits (again without loss of generality, the first t bits), each possible t-bit string occurs equally often. Then H is also vulnerable to a zero-correlation attack. More precisely, $\widehat{H}((a_1, 0), (b_1, 0)) = 0$ for all $a_1 \in \mathbb{F}_2^s$ and $b_1 \in \mathbb{F}_2^t$. Again, this follows directly from Lemma 2. In fact, an integral unconditionally implies zero-correlation.

3.4 Discussion of the link

As pointed out, this relation is intriguing as zero-correlation distinguishers and integral distinguishers are constructed quite differently. Moreover, not only the ways the distinguishers are constructed are different but also the ways the resulting attacks work seem inherently different.

The first difference is that zero-correlation attacks are usually known plaintext attacks (or using known distinct plaintexts, while integral attacks are usually chosen plaintext attacks. Moreover, for zero-correlation attacks, appending

rounds before the distinguisher normally does not increase the data complexity. On the other hand, appending rounds before an integral distinguisher often results in an increased data complexity as, for each (partial) key guess, one has to ensure that some values are fixed according to the distinguisher. Finally, integral distinguishers have the advantage that it is often possible to extend the distinguisher by relaxing the balanced property to a zero-sum property (or equivalently to the fact that a certain subfunction does not have maximal algebraic degree). For zero correlation attacks, such an extension is not known so far.

Thus, besides being interesting from a theoretical perspective, the above mentioned link clearly calls for further work on combining the specific advantages offered by both attacks.

Before discussing an application of this relation to mount an integral attack on a variant of Skipjack, we'd like to illustrate the above with AES as an example.

3.5 Example with AES

Fig.2 depicts the well-known 3-round integral distinguisher for AES. Starting with one active byte and fixing all other bytes results in all bytes being active after ShiftRows in the third round. In terms of zero-correlation distinguisher, the above discussion implies that for any non-zero input mask with (at least) one zero byte and any non-zero output mask which is zero in all but one byte the corresponding linear approximation is unbiased.

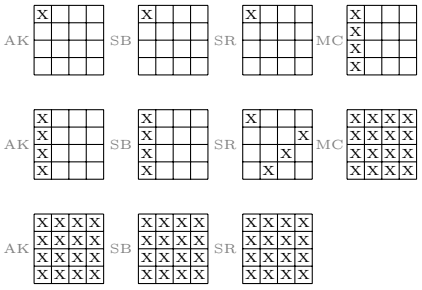

Fig. 2: The integral distinguisher on 3 rounds of AES. The X denotes an active byte.

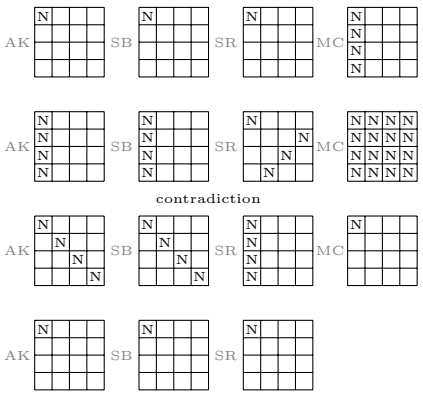

Fig. 3: Zero correlation distinguisher on 4 rounds of AES. The N denotes a non-zero byte in the mask.

Reciprocally, Fig.3 shows the 4 round zero-correlation distinguisher from [7]. For any non-zero mask which is zero in all-but-one bytes and any output mask with the same condition, the corresponding linear approximation is unbiased.

Now, again using the above discussion, this implies the following integral distinguisher on 4 rounds of AES. Fix any byte in the plaintext and encrypt all remaining 2^{120} possible plaintexts. Check if the output restricted to any byte results is a balanced function, that is, of each out of the possible 256 values is obtained exactly 2^{112} times. Note that this distinguisher was implicitly used for example in [14].

4 Integral zero-correlation for a Skipjack variant

4.1 Skipjack-BABABABA vs the original Skipjack-AABBAABB

Skipjack [25] is the only block cipher known to be designed by NSA. Skipjack is a 64-bit block cipher with an 80-bit key. It is an unbalanced Feistel network with 32 rounds of two types, called Rule A and Rule B. Each round is described in the form of a linear feedback shift register with additional non-linear keyed G permutation. Rule B is basically the inverse of Rule A with minor positioning differences. Skipjack applies eight rounds of Rule A, followed by eight rounds of Rule B, followed by another eight rounds of Rule A, followed by another eight rounds of Rule B. We refer to this original Skipjack algorithm as Skipjack-AABBAABB – A denoting four rounds of Rule A and B standing for four rounds of Rule B. The best known cryptanalytic result for the original Skipjack-AABBAABB is the impossible differential cryptanalysis for 31 rounds given by Biham et.al. [2] based on a 24-round impossible differential.

In Skipjack-BABABABA, four rounds of Rule B are applied first, followed by four rounds of Rule A, followed by another four rounds of Rule B, followed by another four rounds of Rule A. The rest of the cipher is exactly as in Skipjack-AABBAABB, amounting to 32 rounds in total. See the Fig.4a. Skipjack variants involving the change of order of Rules A and B were studied in [19, 20]. Though it was suggested that putting Rule B before Rule A might facilitate truncated differentials as a matter of principle, no attacks have been reported on Skipjack-BABABABA.

For Skipjack-BABABABA, the longest impossible differential we can find is over 21 rounds and covers less rounds than the 24-round impossible differential for the original Skipjack. However, in the following, we derive 30-round zero-correlation linear approximations for Skipjack-BABABABA.

4.2 Zero-correlation linear approximations for 30 rounds of Skipjack-BABABABA

Let the input masks for the first round be $(L_1, L_1, 0, 0)$ and the output mask for the last round be $(L_2, L_2, 0, 0)$ for any non-zero L_1 and L_2. Fig.4b depicts the evolution of both masks from the top and from the bottom towards the middle of the cipher. In the figure, M_i denotes an undetermined non-zero mask and R_i denotes an undetermined mask (zero or non-zero). From the input mask $(L_1, L_1, 0, 0)$ at the first round, the output mask of the 19-th round is (M_4, R_2, R_1, M_5). From the output mask $(L_2, L_2, 0, 0)$ at the 30-th round, the input mask of the 20-th

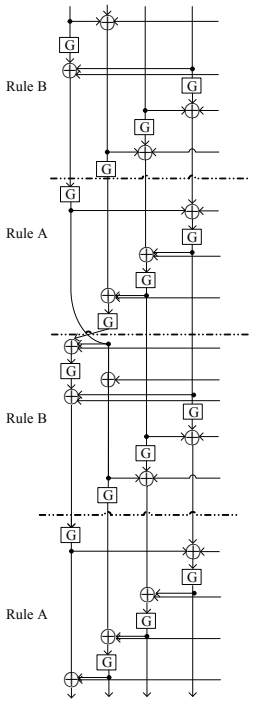

(a) Skipjack-BABABABA

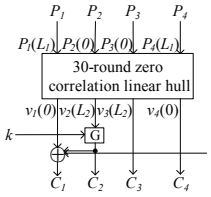

(c) Integral zero-correlation attack on 31-round Skipjack-BABABABA (values in brackets are masks, plain values are actual data processed)

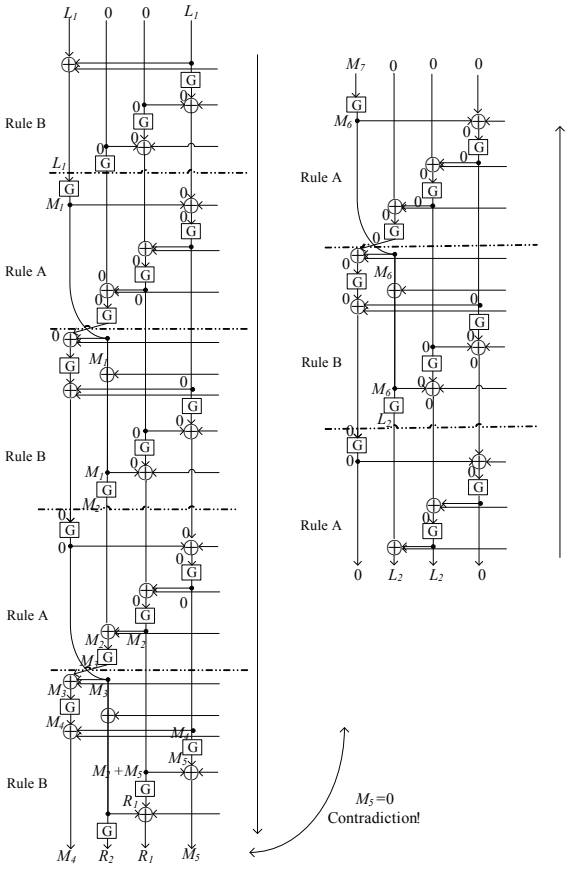

(b) 30-round zero-correlation linear approximations for Skipjack-BABABABA (M_i denotes an undetermined non-zero mask and R_i denotes an undetermined mask – zero or non-zero)

Fig. 4: Integral zero-correlation cryptanalysis of 31-round Skipjack-BABABABA

round is $(M_7, 0, 0, 0)$. Here we conclude that $(M_4, R_2, R_1, M_5) \neq (M_7, 0, 0, 0)$ as equality would imply that $M_5 = 0$ contradicting that $M_5 \neq 0$. Therefore, the linear hull of the 30-round linear approximation $(L_1, L_1, 0, 0) \rightarrow (L_2, L_2, 0, 0)$ does not contain linear trails of non-zero correlation contribution and, thus, has correlation zero.

Property 1. In Skipjack-BABABABA, each linear approximation of the form $(L_1, L_1, 0, 0) \rightarrow (L_2, L_2, 0, 0)$ for non-zero L_1 and L_2 over the 30 rounds $\mathrm{B}^3\mathrm{ABABABA}^3$ has zero correlation. Here $\mathrm{B}^3\mathrm{ABABABA}^3$ means that the 30 rounds start from three consecutive rounds of Rule B, followed by ABABAB and by three consecutive rounds of Rule A.

4.3 Zero-correlation integral attack on 31-round Skipjack-BABABABA

Here we describe how to use Proposition 1 to attack 31 rounds of Skipjack-B^3ABABABA using an integral distinguisher. Combining Proposition 1 with Property 1 leads to the following distinguisher.

Corollary 1. *With the notation of Fig.4c, for the 30-round Skipjack-$B^3ABABABA^3$, encrypting all 2^{48} plaintexts of the form $(P_1|P_2|P_3|P_1)$ each of the 2^{16} possible values of $v_2 \oplus v_3$ occurs exactly 2^{32} times.*

With the notation of Fig.4c, this distinguisher can now be used directly to mount a key-recovery attack on the 31 rounds of Skipjack-B^3ABABABA as follows.

- Initialize 2^{32} counters $V_1[C_2|C_3]$ to zero.
- Encrypt each of all 2^{48} plaintexts of the form $(P_1|P_2|P_3|P_1)$, and increase $V_1[C_2|C_3]$ by one.
- For each guess of the 2^{32} possible values for k:
  - Initialize 2^{16} counters $V_2[v]$ to zero.
  - Decrypt all 2^{16} values of C_2 to get $v_2|v_3$ and increase $V_2[v_2 \oplus v_3]$ by $V_1[C_2|C_3]$.
  - If one of the counters $V_2[v] \neq 2^{16}$, discard k as a wrong key-guess.

With high probability only the correct guess for k will not be discarded. As the key size for Skipjack is 80 bits, the remaining key bits can be brute-forced with a complexity of 2^{48}. The time complexity of this attack is roughly 2^{49} Skipjack encryptions and we have to store roughly 2^{32} counters. The data complexity is 2^{48} chosen plaintexts. Thus, this attack has practical complexities.

5 Zero-correlation and multidimensional linear distinguishers

5.1 Multidimensional linear setting

Given m linear approximations

$$\langle u_i, x\rangle + \langle w_i, y\rangle, \quad i = 1, \ldots, m,$$

where $x \in \mathbb{F}_2^n$ is plaintext and $y \in \mathbb{F}_2^t$ is some part of data in the encryption process, one obtains an m-tuple of bits by evaluating those for a plaintext-ciphertext pair. Instead of considering each such bit and its distribution independently as x varies, multidimensional linear cryptanalysis focuses on the analysis of the distribution of the m-tuples

$$z = (z_1, \ldots, z_m), \quad z_i = \langle u_i, x \rangle + \langle w_i, y \rangle.$$

Then we have the following relationship between the probability distribution of z and the correlations c_γ of all linear approximations $\gamma \in \mathbb{F}_2^m$:

$$\Pr[z] = 2^{-m} \sum_{\gamma \in \mathbb{F}_2^m} (-1)^{\langle \gamma, z \rangle} c_\gamma. \tag{2}$$

Note that this is actually the key in proving that for a balanced function all component functions have zero-correlation.

We denote by U and W the $m \times n$ and $m \times t$ matrices with rows u_i and w_i, respectively. Then we have $z = Ux + Wy$ and can write

$$\langle \gamma, z \rangle = \langle \gamma, Ux + Wy \rangle = \langle U^T \gamma, x \rangle + \langle W^T \gamma, y \rangle, \tag{3}$$

where $U^T\gamma$ and $W^T\gamma$ are linear combinations of the linear masks u_i and w_i, $i = 0, \ldots, m$, respectively.

5.2 How to make zero-correlation multidimensional

Now we are ready to formulate the zero-correlation distinguishing property as a special case of the multidimensional distinguishing property.

Zero-correlation distinguisher assumes that the correlations of all linear approximations $\langle u_i, x \rangle + \langle w_i, y \rangle$, $i = 1, \ldots, m$, and their nonzero linear combinations are equal to zero. (Note that this means, in particular, that these m linear approximations are statistically independent.) By (3), it follows that $c_\gamma = 0$, for all $\gamma \neq 0$. When substituting this information in the formula of $\Pr[z]$ in (2), we obtain that z has a uniform distribution in $\mathbb{F}_2^m$.

Let the adversary be given N distinct plaintexts for an n-bit block cipher and m linear approximations such that all their nonzero linear combinations have correlation zero. Then he can construct, as shown above, a function from $\mathbb{F}_2^n$ to $\mathbb{F}_2^m$ whose outputs z computed for all plaintexts are uniformly distributed m-tuples of bits in $\mathbb{F}_2^m$.

Such a completely uniform distribution is very unlikely to have been obtained from selecting the values at random in $\mathbb{F}_2^m$, even if the probability of each value is equal, spanning a linear space of $\ell = 2^m$ zero-correlation approximations of dimension m. But as we will see, it is possible to distinguish the non-random behavior of the cipher data already with much less data than the full codebook. The distribution of the cipher data follows *multivariate hypergeometric distribution*, while the data drawn at random from a uniform distribution on $\mathbb{F}_2^m$ follows *multinomial distribution*. These distributions have essentially different parameters for large sample sizes N and can be distinguished from each other. The distinguisher can be obtained as follows.

5.3 Multidimensional distinguisher for correlation zero

For each of the 2^m data values $z \in \mathbb{F}_2^m$, the attacker initializes a counter $V[z]$, $z = 0, 1, 2, \ldots, 2^m - 1$, to zero value. Then, for each distinct plaintext, the attacker computes the corresponding data value in $\mathbb{F}_2^m$ (by evaluating the m basis linear approximations) and increments the counter $V[z]$ of this data value by one. Then the attacker computes the statistic T for this distribution as

$$T = \sum_{i=0}^{2^m-1} \frac{(V[z] - N2^{-m})^2}{N2^{-m}(1 - 2^{-m})}. \tag{4}$$

The statistic T will have two distinct distributions for the cipher exhibiting zero-correlation and a randomly drawn permutation which is our wrong-key hypothesis assumption:

Proposition 2. *For sufficiently large sample size N and number ℓ of zero-correlation linear approximations given for the cipher, the statistic T follows a χ^2-distribution for the cipher approximately with mean and variance*

$$\mu_0 = \mathrm{Exp}(T_{cipher}) = (\ell - 1)\frac{2^n - N}{2^n - 1} \quad \textit{and} \quad \sigma_0^2 = \mathrm{Var}(T_{cipher}) = 2(\ell - 1)\left(\frac{2^n - N}{2^n - 1}\right)^2$$

and for a randomly drawn permutation with mean and variance

$$\mu_1 = \mathrm{Exp}(T_{random}) = \ell - 1 \quad \textit{and} \quad \sigma_1^2 = \mathrm{Var}(T_{random}) = 2(\ell - 1).$$

The proof of this proposition is available in the full version of this paper [6].

5.4 Distinguishing complexity

Applying the standard normal approximation of χ^2 to the two different distributions of the statistic T in Proposition 2, one can compute data complexities N of the distinguisher, given error probabilities. As a rule of thumb, we can conclude that it is sufficient to have $N \approx 2^{n+2-\frac{m}{2}}$ distinct plaintexts and their corresponding ciphertexts to distinguish the cipher distribution from randomly drawn permutation. A more precise distinguishing complexity is given by the following statement.

Corollary 2. *Under the assumptions of Proposition 2, for type-I error probability α_0 (the probability to wrongfully discard the cipher), type-II error probability α_1 (the probability to wrongfully accept a randomly chosen permutation as the cipher), for an n-bit block cipher exhibiting ℓ zero-correlation linear approximations forming an $\log_2 \ell$-dimensional linear space, the distinguishing complexity N can be approximated as*

$$N = \frac{2^n(q_{1-\alpha_0} + q_{1-\alpha_1})}{\sqrt{\ell/2} - q_{1-\alpha_1}},$$

where $q_{1-\alpha_0}$ and $q_{1-\alpha_1}$ are the respective quantiles of the standard normal distribution.

Note that this statistical test is based on the decision threshold of $\tau = \mu_0 + \sigma_0 q_{1-\alpha_0} = \mu_1 - \sigma_1 q_{1-\alpha_1}$: If the statistic $T \leq \tau$, the test outputs 'cipher'. Otherwise, if the statistic $T > \tau$, the test returns 'random'.

6 Multidimensional zero-correlation for 28-round CAST-256

6.1 Description of CAST-256

As a first-round AES candidate, CAST-256 is designed based on CAST-128. The block size is 128 bits, and the key size can be 128, 192 or 256 bits. CAST-256 has 48 rounds for all key sizes. The design of CAST-256 is a generalized Feistel network with 4 lines as illustrated in Fig.5a.

We denote the 128-bit block of CAST-256 as $\beta = (A|B|C|D)$, where A, B, C and D are 32 bits each. Two types of round function, the *forward quad-round* $Q(\cdot)$ and the *reverse quad-round* $\bar{Q}(\cdot)$ are used in CAST-256.

The forward quad-round $\beta := Q_i(\beta)$ is defined as consecutive application of 4 rounds as follows:

$$C = C \oplus F_1(D, {K_{R_1}}^{(i)}, {K_{M_1}}^{(i)}),\ \ B = B \oplus F_2(C, {K_{R_2}}^{(i)}, {K_{M_2}}^{(i)}),$$
$$A = A \oplus F_3(B, {K_{R_3}}^{(i)}, {K_{M_3}}^{(i)}),\ \ D = D \oplus F_1(A, {K_{R_4}}^{(i)}, {K_{M_4}}^{(i)}).$$

Similarly, the reverse quad-round $\beta := \bar{Q}_i(\beta)$ is defined as:

$$D = D \oplus F_1(A, {K_{R_4}}^{(i)}, {K_{M_4}}^{(i)}),\ \ A = A \oplus F_3(B, {K_{R_3}}^{(i)}, {K_{M_3}}^{(i)}),$$
$$B = B \oplus F_2(C, {K_{R_2}}^{(i)}, {K_{M_2}}^{(i)}),\ \ C = C \oplus F_1(D, {K_{R_1}}^{(i)}, {K_{M_1}}^{(i)}),$$

where ${K_R}^{(i)} = \{{K_{R_1}}^{(i)}, {K_{R_2}}^{(i)}, {K_{R_3}}^{(i)}, {K_{R_4}}^{(i)}\}$ is the set of rotation keys for the i-th quad-round, and ${K_M}^{(i)} = \{{K_{M_1}}^{(i)}, {K_{M_2}}^{(i)}, {K_{M_3}}^{(i)}, {K_{M_4}}^{(i)}\}$ is the set of masking keys for the i-th quad-round.

The encryption procedure for CAST-256 consists of 6 forward quad-rounds followed by 6 reverse quad-rounds, counting 48 rounds in total. Decryption is identical to encryption except that the sets of quad-round keys ${K_R}^{(i)}$ and ${K_M}^{(i)}$ are applied in the reverse order. The keys are obtained from an up to 256-bit master key by encrypting it with a CAST-256-type cipher (acting on on eight 32-bit words) with known constants as subkeys.

The functions F_1, F_2 and F_3 are exactly those of CAST-128. They use four 8x32-bit S-boxes based on bent functions, modular addition, modular subtraction, XOR and key-dependent rotation. See Fig. 5a.

6.2 24-Round zero-correlation linear approximations for CAST-256

Property 2. For 24-round CAST-256 (3 forward quad-rounds followed by 3 reverse quad-rounds, or rounds 13-36), if the input mask is $(0|0|0|L_1)$ and the output mask is $(0|0|0|L_2)$, the correlation of the linear approximation for the 24-round CAST-256 is zero, where $L_1 \neq L_2, L_1 \neq 0$, and $L_2 \neq 0$.

The proof of this property is available in the full version of this paper [6].

As compared to this 24-round property, the longest impossible differential for CAST-256 we are aware of covers 18 rounds [28]. The work [3] claims unspecified 20-round impossible differentials. Thus, the zero-correlation property for CAST-256 is at least 4 rounds longer than the one of impossible differential.

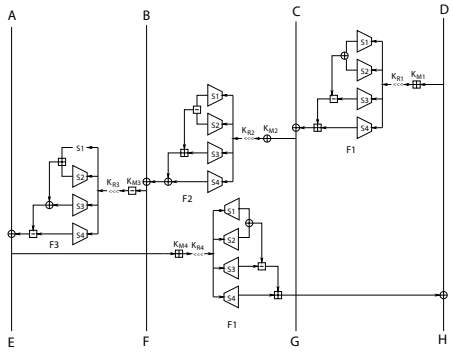

(a) Forward quad-round of CAST-256

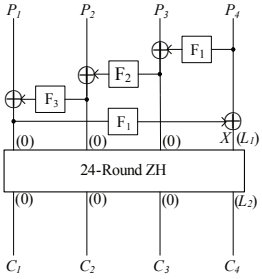

(c) Multidimensional zero-correlation cryptanalysis of 28-round CAST-256: 4 forward quad-rounds and 3 reverse quad-rounds (values in brackets denote masks, plain values are actual data processed)

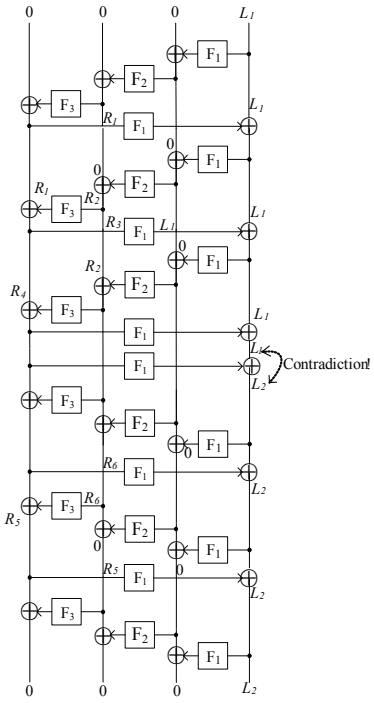

(b) 24-round zero-correlation linear approximations of CAST-256 (3 forward and 3 reverse quad-rounds)

Fig. 5: Multidimensional zero-correlation cryptanalysis of 28-round CAST-256

6.3 Key recovery for 28-round CAST-256

We use the 24-round zero-correlation linear approximations of Property 2 to attack 28 rounds of CAST-256. Fig. 5c illustrates the recovery of the subkey values from the first round to the fourth round. The attack works as follows.

For each possible 148-bit subkey value $\kappa = {K_R}^{(1)} | {K_M}^{(1)}$:

1. Allocate a 64-bit global counter $V[z]$ for each of 2^{64} possible values of the 64-bit vector z and set it to 0. $V[z]$ will contain the number of times the vector value z occurs for the current key guess κ. The vector z is the concatenation of evaluations of 64 basis zero-correlation masks.

2. For each of N distinct plaintext-ciphertext pairs:
   (a) Partially encrypt 4 rounds and get 64-bit value for $X|C_4$.
   (b) Evaluate all 64 basis zero-correlation masks on $X|C_4$ and put the evaluations to the vector z.
   (c) Increment $V[z]$.
3. Compute the χ^2 statistic $T = N2^{64} \sum_{z=0}^{2^{64}-1} \left(\frac{V[z]}{N} - \frac{1}{2^{64}} \right)^2$.
4. If $T < \tau$, then the subkey guess κ is a possible subkey candidate and all master keys it is compatible with are tested exhaustively against a maximum of 3 plaintext-ciphertext pairs.

Table 2: Summary of attacks on CAST-256: KP = Known Plaintexts, CP = Chosen Plaintexts.

Rounds	Key size	Attack	Data	Time	Memory (bytes)	Ratio of weak keys	Ref.
16	$128, 192, 256$	boomerang	$2^{49.3}$CP	–	–	1	[26]
24	192 or 256	linear	$2^{124.1}$KP	$2^{156.52}$	–	1	[27]
36	256	differential	2^{123}CP	2^{182}	–	2^{-35}	[24]
28	256	multidim. ZC	$2^{98.8}$KP	$2^{246.9}$	2^{68}	1	Here

In this attack, using Corollary 2, we set the type-I error probability (the probability to miss the right key) to $\alpha_0 = 2^{-2.7}$ and the type-II error probability (the probability to accept a wrong key) to $\alpha_1 = 2^{-14}$. Thus, we get $q_{1-\alpha_0} = 1$ and $q_{1-\alpha_1} = 3.84$. Here, $\tau = \sigma_1 \cdot q_{\alpha_1} + \mu_1 \approx 2^{64}$.

Corollary 2 suggests that the data complexity is $N = 2^{98.8}$ distinct plaintext-ciphertexts with those parameters. The success probability of the entire attack is $1 - \alpha_0 \approx 0.846$.

The time complexity is $2^{246.8}$ times of one-round encryption and $2^{246.8}$ memory accesses to a memory of size 2^{64}. Under the assumption that one memory access with size 2^{64} is equivalent to one 28-round CAST-256 encryption, the total time complexity would be about $2^{246.9}$ 28-round CAST-256 encryptions. Due to $\alpha_1 = 2^{-14}$ and the total number of recovered bits is 148, the number of the remaining subkey values is $2^{-14} \cdot 2^{148} = 2^{134}$. Then we exhaustively search other $256 - 148 = 108$ subkey bits, the time complexity will be $2^{134+108} = 2^{242}$ times of 28-round encryptions.

The memory requirements are 2^{64} 128-bit words needed for $V[z]$, or 2^{68} bytes.

In all, the data complexity is about $2^{98.8}$ known plaintexts, the time complexity is about $2^{246.9}$ 28-round CAST-256 encryptions and the memory requirements are 2^{64} blocks. This is the first attack on more than half of the full-round AES-candidate CAST-256 without the weak key assumption. See Table 2 for a summary and a comparison of attacks.

7 Conclusions

In this paper, we establish fundamental links between zero-correlation distinguishers on the one hand and integral and multidimensional linear distinguishers on the other. In particular, an integral implies a zero-correlation property and zero-correlation distinguishers can be seen as a special case of multidimensional linear distinguishers. These findings result in two novel distinguishers for zero-correlation based on integral and multidimensional linear distinguishers. To obtain the latter, we refine the theory of multidimensional linear distinguishers. We illustrate these new distinguishers by mounting attacks on a Skipjack variant and CAST-256.

Acknowledgements. Andrey Bogdanov is postdoctoral fellow of the Fund for Scientific Research - Flanders (FWO). This work has been supported in part by the IAP Programme P6/26 BCRYPT of the Belgian State, by the European Commission under contract number ICT-2007-216676 ECRYPT NoE phase II, by KU Leuven-BOF (OT/08/027), by the Research Council KU Leuven (GOA TENSE), by NSFC Projects (No.61133013 and No.61070244), by 973 project (2013CB834205) as well as Interdisciplinary Research Foundation of Shandong University (No.2012JC018).

References

1. T. Baigneres, P. Junod, S. Vaudenay. How Far Can We Go Beyond Linear Cryptanalysis? ASIACRYPT 2004, LNCS, vol. 3329, pp. 432–450, Springer-Verlag, 2004.
2. E. Biham, A. Biryukov, A. Shamir. Cryptanalysis of Skipjack reduced to 31 rounds using impossible differentials. Advances in Cryptology: EUROCRYPT'99, LNCS, vol. 1592, pp. 12–23, Springer-Verlag, 1999.
3. E. Biham, A. Biryukov, A. Shamir. Miss in the Middle Attacks on IDEA and Khufu. FSE'99, LNCS, vol. 1636, pp. 124–138, Springer-Verlag, 1999.
4. Biryukov, A., Cannière, C.D., Quisquater, M. On Multiple Linear Approximations. CRYPTO '04, LNCS, vol. 3152, pp. 1–22, Springer-Verlag, 2004.
5. A. Biryukov, A. Shamir. Structural Cryptanalysis of SASAS. EUROCRYPT 2001, LNCS, vol. 2045, pp. 394–405, Springer-Verlag, 2001.
6. A. Bogdanov, G. Leander, K. Nyberg, M Wang. Integral and Multidimensional Linear Distinguishers with Correlation Zero. IACR ePrint Archive report, 2012.
7. A. Bogdanov, V. Rijmen. Linear Hulls with Correlation Zero and Linear Cryptanalysis of Block Ciphers. Designs, Codes and Cryptography, Springer-Verlag, to appear, 2012. Preprint available as Cryptology ePrint Archive: Report 2011/123, http://eprint.iacr.org/2011/123.
8. A. Bogdanov, M. Wang. Zero Correlation Linear Cryptanalysis with Reduced Data Complexity. FSE'12, LNCS, Anne Canteaut (ed.), Springer-Verlag, to appear, 2012.
9. J. Borst, L. R. Knudsen, and V. Rijmen. Two Attacks on Reduced IDEA. In EUROCRYPT'97, LNCS, vol. 1233, pp. 1–13, Springer-Verlag, 1997.
10. C. Carlet. *Vectorial (multi-output) Boolean Functions for Cryptography.* Cambridge University Press, to appear.

11. C. Carlet. *Boolean Functions for Cryptography and Error Correcting Codes*. Cambridge University Press, to appear (preliminary version available online at www-rocq.inria.fr/codes/Claude.Carlet/chap-fcts-Bool.pdf.
12. J. Daemen, L.R. Knudsen, V. Rijmen. The Block Cipher Square. FSE'97, LNCS, vol. 1267, pp. 149–165, Springer-Verlag, 1997.
13. H. Englund, A. Maximov. Attack the Dragon. INDOCRYPT 2005. LNCS, vol. 3797, pp. 130–142, Springer-Verlag, 2005.
14. N. Ferguson, J. Kelsey, S. Lucks, B. Schneier, M. Stay, D. Wagner, D. Whiting. Improved cryptanalysis of Rijndael. FSE'00, LNCS, vol. 1978, pp. 213–230, Springer-Verlag, 2000.
15. M. Hermelin, J.Y. Cho, K. Nyberg. Multidimensional Extension of Matsui's Algorithm 2. FSE 2009, LNCS, vol. 5665, pp. 209–227, Springer-Verlag, 2009.
16. M. Hermelin, K. Nyberg. Linear cryptanalysis using multiple linear approximations. In: Junod, P., Canteaut, A. (eds.) Advanced Linear Cryptanalysis of Block and Stream Ciphers. IOS Press, 2011.
17. M. Hermelin, K. Nyberg, J.Y. Cho. Multidimensional Linear Cryptanalysis of Reduced Round Serpent. ACISP 2008, LNCS, vol. 5107, pp. 203–215, Springer-Verlag, 2008.
18. B.S. Kaliski, J., Robshaw, M.J.B. Linear Cryptanalysis Using Multiple Approximations. CRYPTO'94, LNCS, vol. 839, pp. 26–39, Springer-Verlag, 1994.
19. L. R. Knudsen, M. J. B. Robshaw, D. Wagner. Truncated Differentials and Skipjack. CRYPTO'99, LNCS, vol. 1666, pp. 165–180, Springer-Verlag, 1999.
20. L. R. Knudsen, D. Wagner. On the structure of Skipjack. Discrete Applied Mathematics 111(1-2): 103-116 (2001)
21. L. R. Knudsen, D. Wagner. Integral Cryptanalysis. FSE 2002, LNCS, vol. 2365, pp. 112–127, Springer-Verlag, 2002.
22. G. Leander. On Linear Hulls, Statistical Saturation Attacks, PRESENT and a Cryptanalysis of PUFFIN. EUROCRYPT'11, LNCS vol. 6632, pp. 303–322, Springer-Verlag, 2011.
23. S. Lucks. The Saturation Attack - A Bait for Twofish. FSE 2001, LNCS, vol. 2355, pp. 1–15, Springer-Verlag, 2002.
24. H. Seki and T. Kaneko. Differential Cryptanalysis of CAST-256 Reduced to Nine Quad-rounds. IEICE Transactions on Fundamentals of Electronics Communications and Computer Sciences, E84A(4), pp. 913– 918, 2001.
25. Skipjack and KEA Algorithm Specifications, Version 2.0, 29 May 1998. Available at the National Institute of Standards and Technology's web page, http://csrc.nist.gov/groups/ST/toolkit/documents/skipjack/skipjack.pdf
26. D. Wagner. The Boomerang Attack. FSE99, LNCS 1636, pp. 156170, Springer-Verlag, 1999.
27. M. Wang, X. Wang and C. Hu. New Linear Cryptanalytic Results of Reduced-Round of CAST-128 and CAST-256. SAC 2008, LNCS, vol. 5381, pp. 429–441, Springer-Verlag, 2009.
28. Q. Wang and J. Chen. 18-Round Impossible Differential for CAST-256, 2012.