Bison: Instantiating the Whitened Swap-Or-Not Construction

Anne Canteaut, Virginie Lallemand, Gregor Leander, Patrick Neumann,
Friedrich Wiemer

► To cite this version:

Anne Canteaut, Virginie Lallemand, Gregor Leander, Patrick Neumann, Friedrich Wiemer. Bison: Instantiating the Whitened Swap-Or-Not Construction. Eurocrypt 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, May 2019, Darmstadt, Germany. 10.1007/978-3-030-17659-4_20 . hal-02431714

HAL Id: hal-02431714

<https://hal.inria.fr/hal-02431714>

Submitted on 8 Jan 2020

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

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

BISON

Instantiating the Whitened Swap-Or-Not Construction

Anne Canteaut¹, Virginie Lallemand², Gregor Leander²,
Patrick Neumann², and Friedrich Wiemer²

¹ Inria, Paris, France

`anne.canteaut@inria.fr`

² Horst Görtz Institute for IT-Security, Ruhr University Bochum, Germany
`firstname.lastname@rub.de`

Abstract We give the first practical instance – BISON – of the Whitened Swap-Or-Not construction. After clarifying inherent limitations of the construction, we point out that this way of building block ciphers allows easy and very strong arguments against differential attacks.

Keywords Whitened Swap-Or-Not · Instantiating Provable Security · Block Cipher Design · Differential Cryptanalysis

1 Introduction

Block ciphers are among the most important cryptographic primitives as they are at the core responsible for a large fraction of all our data that is encrypted. Depending on the mode of operation (or used construction), a block cipher can be turned into an encryption function, a hash-function, a message authentication code or an authenticated encryption function.

Due to their importance, it is not surprising that block ciphers are also among the best understood primitives. In particular the Advanced Encryption Standard (AES) [2] has been scrutinized by cryptanalysts ever since its development in 1998 [19] without any significant security threat discovered for the full cipher (see e. g. [27,26,7,6,23,28,29]).

The overall structure of AES, being built on several (round)-permutations interleaved with a (binary) addition of round keys is often referred to as key-alternating cipher and is depicted in Fig. 1.

The first cipher following this approach was, to the best of our knowledge, the cipher MMB [17], while the name key-alternating cipher first appears in [20] and in the book describing the design of the AES [21]. Nowadays many secure ciphers follow this construction.

Interestingly, besides its overwhelming use in practice and the intense cryptanalytic efforts spent to understand its practical security, the generic (or idealized) security of key-alternating ciphers has not been investigated until 2012. Here, generic or idealized security refers to the setting where the round functions R_i

are modeled as random permutations. An (computational unbounded) attacker is given access to those round functions via oracle queries and additional oracle access to either the block cipher or a random permutation. The goal of the attacker is to tell apart those two cases. As any attack in this setting is obviously independent of any particular structure of the round function, those attacks are generic for all key-alternating ciphers. In this setting, the construction behind key-alternating ciphers is referred to as the iterated Even-Mansour construction. Indeed, the Even-Mansour cipher [25] can be seen as a one-round version of the key-alternating cipher where the round function is a random permutation.

The first result on the iterated Even-Mansour construction (basically focusing on the two-round version) was given in [10]. Since then, quite a lot of follow-up papers, e.g. [32,3,38,30], managed to improve and generalize this initial result significantly. In particular, [15] managed to give a tight security bound for any number of rounds. Informally, for breaking the r -round Even-Mansour construction, any attacker needs to make roughly $2^{\frac{r}{r+1}n}$ oracle queries.

While this bound can be proven tight for the iterated Even-Mansour construction, it is unsatisfactory for two reasons. First, one might hope to get better security bounds with different constructions and second one might hope to lower the requirement of relying on r random permutations.

Motivated by this theoretical defect and the importance of encrypting small domains with full security (see e.g. [42]), researchers started to investigate alternative ways to construct block ciphers with the highest possible security level under minimal assumptions in ideal models. The most interesting result along those lines is the construction by Tessaro [48]. His construction is based on the Swap-or-Not construction by [31], which was designed for the setting where the component functions are secret. Instead of being based on random permutations, this construction requires only a set of random (Boolean) functions. Tessaro’s construction, coined Whitened Swap-Or-Not (WSN for short), requires only two public random (Boolean) functions f_i with n -bit input, and can be proven to achieve full security, see Section 2 for more details.

However, and this is the main motivation for our work, *no instance of this construction is known*. This situation is in sharp contrast to the case of the iterated Even-Mansour construction, where many secure instances are known for a long time already, as discussed above.

Without such a concrete instance, the framework of [48] remains of no avail. As soon as one wants to use the framework in any way, one fundamentally has to

Figure 1: Key-alternating construction for r rounds, using unkeyed round permutations R_1 to R_r . In practical instantiations, the round keys k_i are typically derived from a master key by some key schedule.

instantiate the Boolean functions modeled as ideal functionalities by efficiently computable functions. Clearly, the above mentioned bound in the ideal model does not say anything about any concrete instance. Tessaro phrases this situation as follows:

Heuristically, however, one hopes for even more: Namely, that under a careful implementation of the underlying component, the construction retains the promised security level. [48]

There has actually been one instance of the previous construction [31], but this has been broken almost instantaneously and completely, as parts of the encryption function were actually linear, see [52]. This failure to securely instantiate the construction points to an important hurdle. Namely, proving the generic bounds and analyzing the security of an instance are technically very different tasks. The security of any block cipher is, with the current state of knowledge, always the security against known attacks. In particular, when designing any concrete block cipher, one has to argue why linear and differential attacks do not threaten the construction.

Our Contribution

Consequently, in this paper we investigate the important, but so far overlooked, aspect of instantiating the WSN construction with a practical secure instance. Practical secure meaning, just like in the case of AES, that the block cipher resists all known attacks. We denote this instance as **BISON** (for Bent whItened Swap Or Not). Our insights presented here are twofold.

First, we derive some inherent restrictions on the choice of the round function f_i . In a nutshell, we show that f_i has to be rather strong, in the sense that its output bit has to basically depend on all input bits. Moreover, we show that using less than n rounds will always result in an insecure construction. Those, from a cryptanalytic perspective rather obvious, results are presented in Section 3. Again, but from a different angle, this situation is in sharp contrast to key-alternating ciphers. In the case of key-alternating ciphers, even with a rather small number of rounds (e.g. ten in the case of AES-128) and rather weak round functions (in case of the AES round function any output bit depends on 32 input bits only and the whole round function decomposes into four parallel functions on 32 bits each) we get ciphers that provide, to the best of our knowledge today and after a significant amount of cryptanalysis, full security.

Second, despite those restrictions of the WSN construction, that have significant impact on the performance of any instance, there are very positive aspects of the WSN construction as well. In Section 4, we first define a family of WSN instances which fulfill our initial restrictions.

As we will show in detail, this allows to argue very convincingly that our instance is secure against differential attacks. Indeed, under standard assumptions, we can show that the probability of any (non-trivial) *differential* is upper bounded by 2^{-n+1} where n is the block size, a value that is close to the ideal case. This

significantly improves upon what is the state of the art for key-alternating ciphers. Deriving useful bounds on differentials is notoriously hard and normally one therefore has to restrict to bounding the probability of *differential characteristics* only. Our results for differential cryptanalysis can be of independent interest in the analysis of maximally unbalanced Feistel networks or nonlinear feedback shift registers.

We specify our concrete instance as a family of block ciphers for varying input length in Section 5. In our instance, we attach importance to simplicity and mathematical clarity. It is making use of bent functions, i. e. maximally non-linear Boolean functions, for instantiating f and linear feedback shift registers (LFSRs) for generating the round keys. Another advantage of BISON is that it defines a whole family of block ciphers, one for any odd block size. In particular it allows the straightforward definition of small scale variants to be used for experiments.

Finally we discuss various other attacks and argue why they do not pose a threat for BISON in Section 6. Particularly the discussion on algebraic attacks might be of independent interest. For this we analyse the growth of the algebraic degree over the rounds. In contrast to what we intuitively expect – an exponential growth (until a certain threshold) as in the case for SPNs [11] – the degree of the WSN construction grows linearly in the degree of the round function f_i . This result can also be applied in the analysis of maximally unbalanced Feistel networks or nonlinear feedback shift registers.

Related Work

The first cipher, a Feistel structure, that allowed similarly strong arguments against differential attacks was presented by Nyberg and Knudsen [45], see also [44] for a nice survey on the topic. This cipher was named CRADIC, as Cipher Resistant Against Differential Cryptanalysis but is often simply referenced as the KN cipher. However, the cipher has been broken quickly afterwards, with the invention of interpolation attacks [34]. Another, technically very different approach to get strong results on resistance against attacks we would like to mention is the decorrelation theory [51]. Interestingly, both previous approaches rely rather on one strong component, i. e. round function, to ensure security, while the WSN approach, and in particular BISON, gains its resistance against differential attacks step by step.

Regarding the analysis of differentials, extensive efforts have been expended to evaluate the MEDP/MELP of SPN ciphers, and in particular of the AES. Some remarkable results were published by [46] and then subsequently improved by [35] with a sophisticated pruning algorithm. Interestingly, further work by [22] and later by [13] revealed that such bounds are not invariant under affine transformations – an equivalence notion often exploited for classification of S-boxes when studying their strength against differential cryptanalysis. All these works stress out how difficult it is to evaluate the MEDP/MELP of SPNs, even for a small number of rounds. On the contrary, and as we are going to elaborate in the remaining of this paper, computing the MEDP of BISON is rather straightforward and independent of the exact details of the components. This can be compared to

Figure 2: Schematic view of the WSN construction.

the wide trail strategy that, making use of the branch number and the superbox argument, allows bounding the probability of any differential characteristic for a large class of SPNs. Our arguments allow to bound the differential probability for a large class of WSN instances.

2 Preliminaries

We briefly recall the Whitened Swap-or-Not construction, recapitulate properties of Boolean functions and shortly cover differential and linear cryptanalysis. We denote by $\mathbb{F}_2$ the finite field with two elements and by $\mathbb{F}_2^n$ the n -dimensional vector space over $\mathbb{F}_2$, i. e. the set of all n -bit vectors with a bitwise xor as the addition.

2.1 Whitened Swap-or-Not

The WSN is defined as follows. Given two round keys k_i, w_i , the i th round R_{k_i, w_i} computes

$$R_{k_i, w_i} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$$

$$R_{k_i, w_i}(x) := x + f_{b(i)}(w_i + \max\{x, x + k_i\}) \cdot k_i$$

where $f_{0,1} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ are modeled as two ideal random functions, the $\max$ function returns the lexicographic biggest value in the input set, and $+$ denotes the addition in $\mathbb{F}_2$ (the bitwise xor). The index $b(i)$ equals zero for the first half of the rounds and one for the second half (see Fig. 2 for a graphical overview of the encryption process).

In the remainder of the paper, we denote by $E_{k,w}^r(x)$ the application of r rounds of the construction to the input x with round keys k_i and w_i derived from the master key (k, w) . Every round is involutory, thus for decryption one only has to reverse the order of the round keys.

Note that the usage of the maximum function is not decisive but that it can be replaced by any function Φ_k that returns a unique representative of the set $\{x, x + k\}$, see [48]. In other words it can be replaced by any function such that $\Phi_k(x) = \Phi_k(y)$ if and only if $y \in \{x, x + k\}$.

The main result given by Tessaro on the security of the WSN is the following:

Proposition 1 (Security of the WSN (Informal) [48]). *The WSN construction is $(2^{n-\mathcal{O}(\log n)}, 2^{n-\mathcal{O}(1)})$ -secure for $\mathcal{O}(n)$ rounds.*

Thus, any adversary trying to distinguish the WSN construction from a random permutation and making at most $2^{n-\mathcal{O}(\log n)}$ queries to the block cipher and $2^{n-\mathcal{O}(1)}$ queries to the underlying function has negligible advantage. Here, the round keys are modeled as independent and uniformly distributed random variables.

2.2 Boolean Functions

A *Boolean function* is defined as a function f mapping n bits to one bit. Any Boolean function

$$f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$$

can be uniquely expressed by its algebraic normal form (ANF), i.e. as a (reduced) multivariate polynomial with n variables $x_0, \dots, x_{n-1}$. For $u \in \mathbb{F}_2^n$ we denote

$$x^u = \prod_{i=0}^{n-1} x_i^{u_i}.$$

The ANF of f can be expressed as

$$f(x) = \sum_{u \in \mathbb{F}_2^n} \lambda_u x^u$$

for suitable choices of $\lambda_u \in \mathbb{F}_2$. The degree of f , denoted by $\deg(f)$ is defined as the maximal weight of a monomial present in the ANF of f . That is

$$\deg(f) = \max\{\text{wt}(u) \mid u \in \mathbb{F}_2^n \text{ such that } \lambda_u \neq 0\}.$$

In the context of symmetric cryptography, the differential and linear behavior of a Boolean function play an important role.

The *derivative* of a function f in *direction* α is defined as $\Delta_\alpha(f)(x) := f(x) + f(x + \alpha)$. Informally, studying the behavior of this derivative is at the core of differential cryptanalysis. If we generalize to the derivative of a vectorial Boolean function $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$, we can additionally specify an output difference β . The differential distribution table (DDT) captures the distribution of all possible derivatives; its entries are

$$\text{DDT}_F[\alpha, \beta] := |\{x \in \mathbb{F}_2^n \mid \Delta_\alpha(F)(x) = \beta\}|,$$

where we leave out the subscript, if it is clear from the context. Note that α is usually referred to as the input difference and β as the output difference.

In a similar way, we can approach the linear behavior of a Boolean function, that is its similarity to any linear function. The *Fourier coefficient* of a function $f : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, which is defined as

$$\hat{f}(\alpha) := \sum_{x \in \mathbb{F}_2^n} (-1)^{\langle \alpha, x \rangle + f(x)},$$

is a very useful way to measure this similarity. Here, the notation $\langle a, b \rangle$ denotes the inner product, defined as $\langle a, b \rangle := \sum_{i=1}^n a_i b_i$. Recall that any affine Boolean function can be written as $x \mapsto \langle \alpha, x \rangle + c$ for some fixed $\alpha \in \mathbb{F}_2^n$ and a constant $c \in \mathbb{F}_2$. In particular, it follows that any such affine function has one Fourier coefficient equal to $\pm 2^n$. More generally, the *nonlinearity* of f , defined as $\text{NL}(f) := 2^n - \max_{\alpha} |\hat{f}(\alpha)|$, measures the minimal Hamming-distance of f to the set of all affine functions.

Analogously to the DDT, for a vectorial Boolean function $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$, we define

$$\hat{F}(\alpha, \beta) = \sum_{x \in \mathbb{F}_2^n} (-1)^{\langle \alpha, x \rangle + \langle \beta, F(x) \rangle},$$

and the linear approximation table (LAT) contains the Fourier coefficients

$$\text{LAT}_F[\alpha, \beta] := \hat{F}(\alpha, \beta).$$

Again we leave out the subscript, if it is clear from the context. Here α is usually referred to as the input mask and β as the output mask. Another representation that is sometimes preferred is the *correlation matrix* that in a similar way contains the correlation values for all possible linear approximations, see [18]. The correlation values are simply scaled versions of the Fourier coefficients, i. e.

$$\Pr[\langle \alpha, x \rangle + \langle \beta, F(x) \rangle = 0] = \frac{1}{2} + \frac{\text{cor}_F(\alpha, \beta)}{2} = \frac{1}{2} + \frac{\hat{F}(\alpha, \beta)}{2^{n+1}}.$$

The advantage of the correlation matrix notation is that the correlation matrix of a composition of functions is nothing but the product of the corresponding matrices. For the linear approximation table, additional scaling is required.

Bent Functions. As they will play an important role in our design of BISON, we recall the basic facts of bent functions. Boolean functions on an even number n of input bits that achieve the highest possible nonlinearity of $2^n - 2^{\frac{n}{2}}$ are called *bent*. Bent functions have been introduced by Rothaus [47] and studied ever since, see also [14, Section 8.6]. Even so bent functions achieve the highest possible nonlinearity, their direct use in symmetric cryptography is so far very limited. This is mainly due to the fact that bent functions are not balanced, i. e. the distribution of zeros and ones is (slightly) biased.

Using Parseval's equality, it is easy to see that a function is bent if and only if all its Fourier coefficients are $\pm 2^{\frac{n}{2}}$. Moreover, an alternative classification that will be of importance for BISON, is that a function is bent if and only if all (non-trivial) derivatives $\Delta_{\alpha}(f)$ are balanced Boolean functions [41].

While there are many different primary and secondary constructions³ of bent functions known, for simplicity and for the sake of ease of implementation, we decided to focus on the simplest known bent functions which we recall next, see also [14, Section 6.2].

³ Primary constructions give bent functions from scratch, while secondary constructions build new bent functions from previously defined ones.

Lemma 1 ([24]). *Let $n = 2m$ be an even integer. The function*

$$f : \mathbb{F}_2^m \times \mathbb{F}_2^m \rightarrow \mathbb{F}_2$$

$$f(x, y) := \langle x, y \rangle$$

is a quadratic bent function. Moreover, any quadratic bent function is affine equivalent to f .

2.3 Differential and Linear Cryptanalysis

The two most important attacks on symmetric primitives are *differential* and *linear* cryptanalysis, respectively developed by Biham and Shamir [5] and by Matsui [40] for the analysis of the Data Encryption Standard. The general idea for both is to find a non-random characteristic in the differential, resp. linear, behavior of the scheme under inspection. Such a property can then be used as a distinguisher between the cipher and a random permutation and in many cases leads to key-recovery attacks.

It is inherently hard to compute these properties for the whole function, thus one typically exploits the special structure of the cipher. For round-based block ciphers one usually makes use of linear and differential characteristics that specify not only the input and output masks (resp. differences) but also all intermediate masks after the single rounds.

In the case of differential cryptanalysis, an r -round characteristic δ is defined by $(r + 1)$ differences

$$\delta = (\delta_0, \dots, \delta_r) \in \mathbb{F}_2^{(r+1)n}.$$

For so-called Markov ciphers and assuming the *independence of round keys*, we can compute the probability of a characteristic *averaged over all round-key sequences*:

$$\text{EP}(\delta) = \prod_{i=0}^{r-1} \Pr[F(x) + F(x + \delta_i) = \delta_{i+1}] = \prod_{i=0}^{r-1} \frac{\text{DDT}_F[\delta_i, \delta_{i+1}]}{2^n},$$

where the encryption iterates the round function F for r rounds. Moreover we usually assume the *hypothesis of stochastic equivalence* introduced by Lai *et al.* [37], stating that the actual probability for any fixed round key equals the average.

In contrast to the normal *characteristic* that defines the exact differences before and after each round, a *differential* takes every possible intermediate differences into account and fixes only the overall input and output differences (which are the two values an attacker can typically control).

However, while bounding the average probability of a differential characteristic is easily possible for many ciphers (using in particular the wide-trail strategy introduced in [16]), bounding the average probability of a differential, which is denoted as the expected differential probability (EDP), is not. Nevertheless, some effort was spent to prove bounds on the maximum EDP (MEDP) for two rounds of some key-alternating ciphers [21,33,46,13].

Similarly, for linear cryptanalysis, an r -round characteristic (also called trail or path) for a round function F is defined by $(r + 1)$ masks

$$\theta = (\theta_0, \dots, \theta_r) \in \mathbb{F}_2^{(r+1)n}$$

and its correlation is defined as

$$\mathbf{cor}_F(\theta) := \prod_{i=0}^{r-1} \mathbf{cor}_F(\theta_i, \theta_{i+1}) = \prod_{i=0}^{r-1} \frac{\widehat{F}(\theta_i, \theta_{i+1})}{2^n}$$

and it can be shown that the correlation of a composition can be computed as the sum over the trail correlations. More precisely,

$$\mathbf{cor}_{E_k^r}(\alpha, \beta) = \sum_{\substack{\theta \\ \theta_0 = \alpha, \theta_r = \beta}} \mathbf{cor}_F(\theta), \quad (1)$$

where the encryption E_k^r iterates the round function F for r rounds.

This is referred to as the linear hull (see [43]). While not visible in order to simplify notation, the terms in Eq. (1) are actually key dependent and thus for some keys they either could cancel out or amplify the overall correlation. For more background, we refer to e.g. [9] and [36]. For a key-alternating cipher with independent round keys, the average over all round-key sequences of the correlation $\mathbf{cor}_{E_k^r}(\alpha, \beta)$ is zero for any pair of nonzero masks (α, β) (see e.g. [21, Section 7.9]). Then, the most relevant parameter of the distribution is its variance, which corresponds to the average square correlation, and is called the *expected linear potential*. Again, bounding the ELP is out of reach for virtually any practical cipher, while for bounding the correlation of a single trail, one can again use the wide-trail strategy mentioned above. Upper bounds for the MELP of two rounds of AES are also given in [33,46,13].

Finally we would like to note that the round keys in an actual block cipher instance are basically never independent and identically distributed over the full key space, but instead derived from a key schedule, rendering the above assumption plain wrong. While the influence of key schedules is a crucially understudied topic and for specific instances strange effects can occur, see [1,36], the above assumption are seen as valid heuristics for most block ciphers.

3 Inherent Restrictions

In this section we point out two inherent restrictions on any practical secure instance, i.e. generic for the WSN construction. Those restrictions result in general conditions on both the minimal number of rounds to be used and general properties of the round functions $f_{b(i)}$. In particular, those insights are taken into account for BISON. While these restrictions are rather obvious from a cryptanalytic point of view, they have a severe impact on the performance of any concrete instance. We discuss performance in more detail in the full version [12, Section 7].

3.1 Number of Rounds

As in every round of the cipher, we simply add (or not) the current round key k_i , the ciphertext can always be expressed as the addition of the plaintext and a (message dependent) linear combination of all round keys k_i . The simple but important observation to be made here is that, as long as the round keys do not span the full space, the block cipher is easily attackable.

Phrased in terms of linear cryptanalysis we start with the following lemma.

Lemma 2. *For any number of rounds $r < n$ there exists an element $u \in \mathbb{F}_2^n \setminus \{0\}$ such that*

$$\widehat{E_{k,w}^r}(u, u) = 2^n,$$

that is the equation

$$\langle u, x \rangle = \langle u, E_{k,w}^r(x) \rangle$$

holds for all $x \in \mathbb{F}_2^n$.

Proof. Let $k_1, \dots, k_r$ be the round keys derived from k and denote by

$$U = \text{span} \{k_1, \dots, k_r\}^\perp$$

the dual space of the space spanned by the round keys, i. e.

$$\forall u \in U, \forall 1 \leq i \leq r \text{ it holds that } \langle u, k_i \rangle = 0.$$

As $r < n$ by assumption, the dimension of $\text{span} \{k_1, \dots, k_r\}$ is smaller than n and thus $U \neq \{0\}$. Therefore, U contains a non-zero element

$$u \in \text{span} \{k_1, \dots, k_r\}^\perp$$

and it holds that

$$\langle u, E_{k,w}^r(x) \rangle = \langle u, x + \sum_{i=1}^r \lambda_i k_i \rangle = \langle u, x \rangle + \langle u, \sum_{i=1}^r \lambda_i k_i \rangle = \langle u, x \rangle. \quad (2)$$

Even more importantly, this observation leads directly to a known plaintext attack with very low data-complexity. Given a set of t plaintext/ciphertext (p_i, c_i) pairs, an attacker simply computes

$$V = \text{span} \{p_i + c_i \mid 1 \leq i \leq t\} \subseteq \text{span} \{k_j \mid 1 \leq j \leq r\}.$$

Given $t > r$ slightly more pairs than rounds, and assuming that $p_i + c_i$ is uniformly distributed in $\text{span} \{k_j\}$ (otherwise the attack only gets even stronger)⁴ implies that

$$V = \text{span} \{k_j\}$$

⁴ E. g. if, with high probability, the $p_i + c_i$ do not depend on one or more k_j 's, the described attack can be extended to one or more rounds with high probability.

with high probability and V can be efficiently computed. Furthermore, as above $\dim(\text{span}\{k_j\})$ is at most r , we have $V^\perp \neq \{0\}$. Given any $u \neq 0$ in $V^\perp$ allows to compute one bit of information on the plaintext given only the ciphertext and particularly distinguish the cipher from a random permutation in a chosen-plaintext setting efficiently.

A similar argument shows the following:

Lemma 3. *For any number of rounds r smaller than $2n - 3$ there exist nonzero α and β , such that*

$$\widehat{E_{k,w}^r}(\alpha, \beta) = 0$$

Proof. We restrict to the case $r \geq n$ as otherwise the statement follows directly from the lemma above. Indeed, from Parseval equality, the fact that $\widehat{E_{k,w}^r}(\alpha, \alpha) = 2^n$ implies that $\widehat{E_{k,w}^r}(\alpha, \beta) = 0$ for all $\beta \neq \alpha$. Let $k_1, \dots, k_r$ be the round keys derived from k and choose non-zero elements $\alpha \neq \beta$ such that

$$\alpha \in \text{span}\{k_1, \dots, k_{n-2}\}^\perp \quad \text{and} \quad \beta \in \text{span}\{k_{n-1}, \dots, k_r\}^\perp.$$

Note that, as $r \leq 2n - 3$ by assumption such elements always exist. Next, we split the encryption function in two parts, the first $n - 2$ rounds E_1 and the remaining $r - (n - 2) < n$ rounds E_2 , i.e.

$$E_{k,w}^r = E_2 \circ E_1.$$

We can compute the Fourier coefficient of $E_{k,w}^r$ as

$$\widehat{E_{k,w}^r}(\alpha, \beta) = \sum_{\gamma \in \mathbb{F}_2^n} \frac{\widehat{E_1}(\alpha, \gamma)}{2^n} \cdot \frac{\widehat{E_2}(\gamma, \beta)}{2^n}.$$

Now, the above lemma and the choices of α and β imply that $\widehat{E_1}(\alpha, \gamma) = 0$ for $\gamma \neq \alpha$ and $\widehat{E_2}(\gamma, \beta) = 0$ for $\gamma \neq \beta$. Recalling that $\alpha \neq \beta$ by construction concludes the proof.

However, as the masks α and β depend on the key, and unlike above there does not seem to be an efficient way to compute those, we do not see a direct way to use this observation for an attack.

Summarizing the observations above, we get the following conclusion:

Rationale 1. *Any practical instance must iterate at least n rounds. Furthermore, it is beneficial if any set of n consecutive round keys are linearly independent.*⁵

After having derived basic bounds on the number of rounds for any secure instance, we move on to criteria on the round function itself.

⁵ If (some) round keys are linearly dependent, Lemma 3 can easily be extended to more rounds.

3.2 Round Function

Here, we investigate a very basic criterion on the round function, namely dependency on all input bits. Given the Boolean functions $f_{b(i)} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ used in the round function of $E_{k,w}^r$, an important question is, if it is necessary that the output bit of $f_{b(i)}$ has to depend on all input bits. It turns out that this is indeed strictly necessary for any secure instance, as summarized in the next rationale.

Rationale 2. *For a practical instance, the functions $f_{b(i)}$ has to depend on all bits. Even more, for any $\delta \in \mathbb{F}_2^n$ the probability of*

$$f_{b(i)}(x) = f_{b(i)}(x + \delta)$$

should be close to $\frac{1}{2}$.

Due to page constraints, we refer to [12, Lemma 4] for more details. It is worth noticing that the analysis leading to Rationale 2 applies to the original round function. However, as pointed out in [49, Section 3.1], in the definition of the round function, we can replace the function

$$x \mapsto \max \{x, x + k\}$$

by any function Φ_k such that $\Phi_k(x) = \Phi_k(x + k)$ for all x . While the following sections will focus on the case when Φ_k is linear, we will prove that Rationale 2 is also valid in this other setting.

Again, this should be compared to key-alternating ciphers, where usually not all output bits of a single round function depend on all input bits. For example for AES any output bit after one round depends only on 32 input bits and for PRESENT any output bit only depends on 4 input bits. However, while for key-alternating ciphers this does not seem to be problematic, and indeed allows rather weak round functions to result in a secure scheme, for the WSN construction the situation is very different.

Before specifying our exact instance, we want to discuss differential cryptanalysis of a broader family of instances.

4 Differential Cryptanalysis of BISON-like Instances

We coin an instance of the WSN construction “BISON-like”, if it iterates at least n rounds with linearly independent round keys $k_1, \dots, k_n$ and applies Boolean functions $f_{b(i)}$. As explained in [49, Section 3.1], in order to enable decryption it is required that the Boolean functions $f_{b(i)}$ return the same result for both x and $x + k$. In the original proposition by Tessaro, this is achieved by using the max function in the definition of the round function. Using this technique reduces the number of possible inputs for the $f_{b(i)}$ to 2^{n-1} . To simplify the analysis and to ease notation, we replace the max function with a *linear function* $\Phi_k : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{n-1}$ with $\ker(\Phi_k) = \{0, k\}$. From now on, we assume that any BISON-like instance

uses such a Φ_k instead of the max function. The corresponding round function has then the following form

$$R_{k_i, w_i}(x) := x + f_{b(i)}(w_i + \Phi_{k_i}(x))k_i. \quad (3)$$

With the above conditions, any BISON-like instance of the WSN construction is resistant to differential cryptanalysis, as we show in the remainder of this section.

For our analysis, we make two standard assumptions in symmetric cryptanalysis as mentioned above: the *independence of whitening round keys* w_i and the *hypothesis of stochastic equivalence* with respect to these round keys. That is, we assume round keys w_i to be independently uniformly drawn and the resulting EDP to equal the differential probabilities averaged over all w . Thus, the keys w_i act very much like the round key for a key-alternating cipher with respect to the probabilities of characteristics. We further back up this intuition by practical experiments (see Section 6.3 and [12, Appendix B]). For the round keys k_i we do not have to make such assumptions.

We first discuss the simple case of differential behaviour for one round only and then move up to an arbitrary number of rounds and devise the number of possible output differences and their probabilities.

4.1 From One-Round Differential Characteristics

Looking only at one round, we can compute the DDT explicitly:

Proposition 2. *Let $R_{k_i, w_i} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ be the WSN round function as in Eq. (3). Then its DDT consists of the entries*

$$\text{DDT}_R[\alpha, \beta] = \begin{cases} 2^{n-1} + \widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) & \text{if } \beta = \alpha \\ 2^{n-1} - \widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) & \text{if } \beta = \alpha + k \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Most notably, if f is bent, we have

$$\text{DDT}_R[\alpha, \beta] = \begin{cases} 2^n & \text{if } \alpha = \beta = k \text{ or } \alpha = \beta = 0 \\ 2^{n-1} & \text{if } \beta \in \{\alpha, \alpha + k\} \text{ and } \alpha \notin \{0, k\} \\ 0 & \text{otherwise.} \end{cases}$$

Proof. We have to count the number of solutions of $R(x) + R(x + \alpha) = \beta$:

$$\begin{aligned} \text{DDT}_R[\alpha, \beta] &= |\{x \in \mathbb{F}_2^n \mid R(x) + R(x + \alpha) = \beta\}| \\ &= |\{x \in \mathbb{F}_2^n \mid \alpha + [f(w + \Phi_k(x)) + f(w + \Phi_k(x + \alpha))] \cdot k = \beta\}| \end{aligned}$$

Since f takes its values in $\mathbb{F}_2$, the only output differences β such that $\text{DDT}_R[\alpha, \beta]$ may differ from 0 are $\beta = \alpha$ and $\beta = \alpha + k$. When $\beta = \alpha$, we have

$$\begin{aligned}\text{DDT}_R[\alpha, \alpha] &= |\{x \in \mathbb{F}_2^n \mid f(w + \Phi_k(x)) + f(w + \Phi_k(x + \alpha)) = 0\}| \\ &= |\{x \in \mathbb{F}_2^n \mid f(w + \Phi_k(x)) + f(w + \Phi_k(x) + \Phi_k(\alpha)) = 0\}| \\ &= 2 \cdot |\{x' \in \mathbb{F}_2^{n-1} \mid f(x') + f(x' + \Phi_k(\alpha)) = 0\}| \\ &= 2 \left(2^{n-2} + \frac{1}{2} \widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) \right).\end{aligned}$$

Similarly,

$$\begin{aligned}\text{DDT}_R[\alpha, \alpha + k] &= |\{x \in \mathbb{F}_2^n \mid f(w + \Phi_k(x)) + f(w + \Phi_k(x + \alpha)) = 1\}| \\ &= 2 \left(2^{n-2} - \frac{1}{2} \widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) \right).\end{aligned}$$

Most notably, when $\alpha \in \{0, k\}$, $\widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) = 2^{n-1}$. Moreover, when f is bent, $\widehat{\Delta_{\Phi_k(\alpha)}(f)}(0) = 2^{n-2}$ for all other values of α .

4.2 To Differentials over more Rounds

As previously explained, it is possible to estimate the probability of a differential characteristic over several rounds, averaged over the round keys, when the cipher is a Markov cipher. We now show that this assumption holds for any BISON-like instance of the WSN construction.

Lemma 4. *Let $R_{k,w} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ be the WSN round function as in Eq. (3). For any fixed $k \in \mathbb{F}_2^n$ and any differential $(\alpha, \beta) \in \mathbb{F}_2^n \times \mathbb{F}_2^n$, we have that*

$$\Pr_w [R_{k,w}(x + \alpha) + R_{k,w}(x) = \beta]$$

is independent of x . More precisely

$$\Pr_w [R_{k,w}(x + \alpha) + R_{k,w}(x) = \beta] = \Pr_x [R_{k,w}(x + \alpha) + R_{k,w}(x) = \beta] .$$

Proof. We have

$$\begin{aligned}& \{w \in \mathbb{F}_2^{n-1} \mid \Delta_\alpha(R_{k,w})(x) = \beta\} \\ &= \{w \in \mathbb{F}_2^{n-1} \mid (\Delta_{\Phi_k(\alpha)}(f)(w + \Phi_k(x))) \cdot k = \alpha + \beta\} \\ &= \begin{cases} \emptyset & \text{if } \beta \notin \{\alpha, \alpha + k\} \\ \Phi_k(x) + \text{Supp}(\Delta_{\Phi_k(\alpha)}(f)) & \text{if } \beta = \alpha + k \\ \Phi_k(x) + (\mathbb{F}_2^{n-1} \setminus \text{Supp}(\Delta_{\Phi_k(\alpha)}(f))) & \text{if } \beta = \alpha, \end{cases}\end{aligned}$$

where $\text{Supp}(g)$ denotes the support of a Boolean function g , i.e., the values x for which $g(x) = 1$. Clearly, the cardinality of this set does not depend on x . Moreover, this cardinality, divided by 2^{n-1} , corresponds to the value of

$$\Pr_x [R_{k,w}(x + \alpha) + R_{k,w}(x) = \beta]$$

computed in the previous proposition.

By induction on the number of rounds, we then directly deduce that any BISON-like instance of the WSN construction is a Markov cipher in the sense of the following corollary.

Corollary 1. *Let $E_{k,w}^i$ denote i rounds of a BISON-like instance of the WSN construction with round function R_{k_i,w_i} . For any number of rounds r and any round keys $(k_1, \dots, k_r)$, the probability of an r -round characteristic δ satisfies*

$$\Pr_w [E_{k,w}^i(x) + E_{k,w}^i(x + \delta_0) = \delta_i, \forall 1 \leq i \leq r] = \prod_{i=1}^r \Pr_x [R_{k_i,w_i}(x) + R_{k_i,w_i}(x + \delta_{i-1}) = \delta_i].$$

For many ciphers several differential characteristics can cluster in a *differential* over more rounds. This is the main reason why bounding the probability of differentials is usually very difficult if possible at all. For BISON-like instances the situation is much nicer; we can actually compute the EDP, i. e., the probabilities of the differentials averaged over all whitening key sequences $(w_1, \dots, w_r)$. This comes from the fact that any differential for less than n rounds contains at most one differential characteristic with non-zero probability. To understand this behavior, let us start by analyzing the EDP (averaged over the w_i) and by determining the number of possible output differences.

In the following, we assume that the input difference α is fixed, and we calculate the number of possible output differences. We show that this quantity depends on the relation between α and the k_i .

Lemma 5. *Let us consider r rounds of a BISON-like instance of the WSN construction with round function involving Boolean functions $f_{b(i)}$ having no (non-trivial) constant derivative. Assume that the first n round keys $k_1, \dots, k_n$ are linearly independent, and that $k_{n+1} = k_1 + \sum_{i=2}^n \gamma_i k_i$ for $\gamma_i \in \mathbb{F}_2$. For any non-zero input difference α , the number of possible output differences β such that*

$$\Pr_{w,x} [E_{k,w}^r(x + \alpha) + E_{k,w}^r(x) = \beta] \neq 0$$

is

$$\begin{cases} 2^r & \text{if } \alpha \notin \text{span}\{k_i\} \text{ and } r < n, \\ 2^r - 2^{r-\ell} & \text{if } \alpha = k_\ell + \sum_{i=1}^{\ell-1} \lambda_i^\alpha k_i \text{ and } r \leq n, \\ 2^n - 1 & \text{if } r > n. \end{cases}$$

Proof. By combining Corollary 1 and Proposition 2, we obtain that the average probability of a characteristic $(\delta_0, \delta_1, \dots, \delta_{r-1}, \delta_r)$ can be non-zero only if $\delta_i \in \{\delta_{i-1}, \delta_{i-1} + k_i\}$ for all $1 \leq i \leq r$. Therefore, the output difference δ_r must be of the form $\delta_r = \delta_0 + \sum_{i=1}^r \lambda_i k_i$ with $\lambda_i \in \mathbb{F}_2$. Moreover, for those characteristics, the average probability is non-zero unless there exists $1 \leq i < r$ such that $|\Delta_{\Phi_{k_i}(\delta_i)}(f_{b(i)})(0)| = 2^{n-1}$, i. e. $\Delta_{\Phi_{k_i}(\delta_i)}(f_{b(i)})$ is constant. By hypothesis, this only occurs when $\delta_i \in \{0, k_i\}$, and the impossible characteristics correspond to

the case when either $\delta_i = 0$ or $\delta_{i+1} = 0$. It follows that the valid characteristics are exactly the characteristics of the form

$$\delta_i = \delta_0 + \sum_{j=1}^i \lambda_j k_j$$

where none of the δ_i vanishes.

- When the input difference $\alpha \notin \text{span}\{k_i\}$, for any given output difference $\beta = \alpha + \sum_{i=1}^r \lambda_i k_i$, the r -round characteristic

$$(\alpha, \alpha + \lambda_1 k_1, \alpha + \lambda_1 k_1 + \lambda_2 k_2, \dots, \alpha + \sum_{i=1}^r \lambda_i k_i)$$

is valid since none of the intermediate differences vanishes.

- When $r \leq n$ and $\alpha = k_\ell + \sum_{i=1}^{\ell-1} \lambda_i^\alpha k_i$, the only possible characteristic from α to $\beta = \alpha + \sum_{i=1}^r \lambda_i k_i$ satisfies

$$\delta_j = \begin{cases} \sum_{i=1}^j (\lambda_i + \lambda_i^\alpha) k_i + \sum_{i=j+1}^\ell \lambda_i^\alpha k_i & \text{if } j \leq \ell \\ \sum_{i=1}^\ell (\lambda_i + \lambda_i^\alpha) k_i + \sum_{i=\ell+1}^j \lambda_i k_i & \text{if } j > \ell. \end{cases}$$

Since the involved round keys are linearly independent, we deduce that $\delta_j = 0$ only when $j = \ell$ and $\lambda_i = \lambda_i^\alpha$ for all $1 \leq i \leq \ell$. It then follows that there exists a valid characteristic from α to β unless $\lambda_i = \lambda_i^\alpha$ for all $1 \leq i \leq \ell$. The number of possible outputs β is then

$$(2^\ell - 1)2^{r-\ell} = 2^r - 2^{r-\ell}.$$

- If we increase the number of rounds to more than n , we have $\alpha = k_\ell + \sum_{i=1}^{\ell-1} \lambda_i^\alpha k_i$ for some $\ell \leq n$. If $\beta = \alpha + \sum_{i=1}^n \lambda_i k_i$ with $\sum_{i=1}^\ell \lambda_i k_i \neq \alpha$, then we can obviously extend the previous n -round characteristic to

$$(\alpha, \alpha + \lambda_1 k_1, \dots, \alpha + \sum_{i=1}^{n-1} \lambda_i k_i, \beta, \beta, \dots, \beta).$$

If $\sum_{i=1}^\ell \lambda_i k_i = \alpha$, β cannot be the output difference of an n -round characteristic. However, the following $(n+1)$ -round characteristic starting from $\delta_0 = \alpha$ is valid:

$$\delta_j = \begin{cases} k_1 + \sum_{i=2}^j \gamma_i k_i + \sum_{i=j+1}^\ell \lambda_i^\alpha k_i & \text{if } j \leq \ell \\ k_1 + \sum_{i=2}^j \gamma_i k_i + \sum_{i=\ell+1}^j \lambda_i k_i & \text{if } \ell < j \leq n \\ \beta & \text{if } j = n+1 \end{cases}$$

Indeed, $\delta_n = \beta + k_n$ implying that the last transition is valid. Moreover, it can be easily checked that none of these δ_j vanishes, unless $\beta = 0$. This implies that all non-zero output differences β are valid.

The last case in the above lemma is remarkable, as it states any output difference is possible after $n + 1$ rounds. To highlight this, we restate it as the following corollary.

Corollary 2. *For BISON-like instances with more than n rounds whose round keys $k_1, \dots, k_{n+1}$ satisfy the hypothesis of Lemma 5, and for any non-zero input difference, every non-zero output difference is possible.*

We now focus on a reduced version of the cipher limited to exactly n rounds and look at the probabilities for every possible output difference. Most notably, we exhibit in the following lemma an upper-bound on the MEDP which is minimized when n is odd and the involved Boolean functions $f_{b(i)}$ are bent. In other words, Rationale 2 which was initially motivated by the analysis in Section 3 for the original round function based on $x \mapsto \max(x, x + k)$ [48] is also valid when a linear function Φ_k is used.

Lemma 6. *Let us consider n rounds of a BISON-like instance of the WSN construction with round function involving Boolean functions $f_{b(i)}$. Let $k_1, \dots, k_n$ be any linearly independent round keys. Then, for any input difference $\alpha \neq 0$ and any β , we have*

$$\begin{aligned} \text{EDP}(\alpha, \beta) &= \Pr_{w,x} [E_{k,w}(x + \alpha) + E_{k,w}(x) = \beta] \\ &\leq \left(\frac{1}{2} + 2^{-n} \max_{1 \leq i \leq n} \max_{\delta \neq 0} \left| \widehat{\Delta_\delta(f_{b(i)})(0)} \right| \right)^{n-1}. \end{aligned}$$

More precisely, if all $f_{b(i)}$ are bent,

$$\text{EDP}(\alpha, \beta) = \begin{cases} 0 & \text{if } \beta = \sum_{i=\ell+1}^n \lambda_i k_i, & (5) \\ 2^{-n+1} & \text{if } \beta = k_\ell + \sum_{i=\ell+1}^n \lambda_i k_i, & (6) \\ 2^{-n} & \text{otherwise,} & (7) \end{cases}$$

where ℓ denotes as previously the latest round key that appears in the decomposition of α into the basis $(k_1, \dots, k_n)$, that is $\alpha = k_\ell + \sum_{i=1}^{\ell-1} \lambda_i k_i$.

The case of bent functions is visualized in Fig. 3, where we give an example of the three possibilities for three rounds.

Proof. As proved in Lemma 5, (α, β) is an impossible differential if and only if $\beta = \sum_{i=\ell+1}^n \lambda_i k_i$. For all other values of $\beta = \alpha + \sum_{i=1}^n \lambda_i k_i$, we have

$$\text{EDP}(\alpha, \beta) = \prod_{i=1}^n \left(\frac{1}{2} + (-1)^{\lambda_i} 2^{-n} \widehat{\Delta_{\Phi_{k_i}(\delta_i)}(f_{b(i)})(0)} \right)$$

Figure 3: Probabilities of output differences for three rounds and the cases of the input difference $\alpha = k_1 + k_2$, thus $\ell = 2$. Dotted transitions are impossible.

where $\delta_i = \alpha + \sum_{j=1}^i \lambda_j k_j$. The i -th term in the product is upper-bounded by

$$\frac{1}{2} + 2^{-n} \max_{1 \leq i \leq n} \max_{\delta \neq 0} \left| \widehat{\Delta_\delta(f_{b(i)})(0)} \right|$$

except if $\Phi_{k_i}(\delta_i) = 0$, i. e., $\delta_i \in \{0, k_i\}$. As seen in Lemma 5, the case $\delta_i = 0$ cannot occur in a valid characteristic. The case $\delta_i = k_i$ occurs if and only if $i = \ell$ and $\beta = k_\ell + \sum_{j=\ell+1}^n \lambda_j k_j$. In this situation, the ℓ -th term in the product equals 1. In the tree of differences this is visible as the collapsing of the two branches from two possible succeeding differences into only one, which then of course occurs with probability one, see upper branch of Fig. 3.

Most notably, all $f_{b(i)}$ are bent if and only if

$$\max_{1 \leq i \leq n} \max_{\delta \neq 0} \left| \widehat{\Delta_\delta(f_{b(i)})(0)} \right| = 0 ,$$

leading to the result.

This can be seen on Fig. 3: the $2^{n-\ell}$ possible differences coming from the collapsed branch have a transition of probability one in that round, resulting in an overall probability of 2^{-n+1} , see Eq. (6). For the lower part of Fig. 3, all the other differences are not affected by this effect and have a probability of 2^{-n} , see Eq. (7).

Because they allow us to minimize the MEDP, we now concentrate on the case of bent functions for the sake of simplicity, which implies that the block size is odd. However, if an even block size is more appropriate for implementation reasons, we could also define BISON-like instances based on maximally nonlinear functions.

It would be convenient to assume in differential cryptanalysis that the EDP of a differential does not increase when adding more rounds, while this does not

hold in general. However, this argument can easily be justified for BISON-like instances using bent functions, when averaging over the whitening keys w .

Proposition 3. *Let us consider $r \geq n$ rounds of a BISON-like instance of the WSN construction with bent functions $f_{b(i)}$. Let $k_1, \dots, k_n$ be any linearly independent round keys. Then the probability of any non-trivial differential, averaged over all whitening key sequences w is upper bounded by 2^{-n+1} .*

In other words, the MEDP of BISON-like instances with bent $f_{b(i)}$ for $r \geq n$ rounds is 2^{-n+1} .

Proof. By induction over r . The base case for $r = n$ rounds comes from Lemma 6. In the induction step, we first consider the case when the output difference β after r rounds differs from k_r . Then the output difference $\delta_r = \beta$ can be reached if and only if the output difference after $(r-1)$ rounds δ_{r-1} belongs to $\{\beta, \beta + k_r\}$. Then,

$$\begin{aligned} \text{EDP}^r(\alpha, \beta) &= \Pr_{w_r} [R_{k_r, w_r}(x_r) + R_{k_r, w_r}(x_r + \beta) = \beta] \text{EDP}^{r-1}(\alpha, \beta) \\ &\quad + \Pr_{w_r} [R_{k_r, w_r}(x_r) + R_{k_r, w_r}(x_r + \beta + k_r) = \beta] \text{EDP}^{r-1}(\alpha, \beta + k_r) \\ &= \frac{1}{2} (\text{EDP}^{r-1}(\alpha, \beta) + \text{EDP}^{r-1}(\alpha, \beta + k_r)) \leq 2^{-n+1}. \end{aligned}$$

When the output difference β after r rounds equals k_r , it results from $\delta_{r-1} = k_r$ with probability 1. In this case

$$\text{EDP}^r(\alpha, \beta) = \text{EDP}^{r-1}(\alpha, \beta) \leq 2^{-n+1}.$$

This bound is close to the ideal case, in which each differential has probability $1/(2^n - 1)$.

We now give a detailed description of our instance BISON.

5 Specification of BISON

As BISON-like instances should obviously generalise BISON, this concrete instance inherits the already specified parts. Thus BISON uses two bent functions $f_{b(i)}$, replaces the max function by Φ_k , and uses a key schedule that generates round keys, where all n consecutive round keys are linearly independent. The resulting instance for n bits iterates the WSN round function as defined below over $3 \cdot n$ rounds. The chosen number of rounds mainly stems from the analysis of the algebraic degree that we discuss in Section 6.

Security Claim. *We claim n -bit security for BISON in the single-key model. We emphasize that we do not claim any security in the related-key, chosen-key or known-key model.*

5.1 Round Function

For any nonzero round key k , we define $\Phi_k : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{n-1}$ as

$$\Phi_k(x) := (x_{i(k)} \cdot k + x)[1, \dots, i(k) - 1, i(k) + 1, \dots, n], \quad (8)$$

where $i(k)$ denotes the index of the lowest bit set to 1 in k , and the notation $x[1, \dots, j-1, j+1, \dots, n]$ returns the $(n-1)$ -bit vector, consisting of the bits of x except the j th bit.

Lemma 7. *The function $\Phi_k : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^{n-1}$ is linear and satisfies*

$$\ker(\Phi_k) = \{0, k\}.$$

The proof can be done by simply computing both outputs for x and $x + k$. For the preimage of $y \in \mathbb{F}_2^{n-1}$ and $j = i(k)$ we have

$$\Phi_k^{-1}(y) \in \left\{ \begin{array}{l} (y[1:j-1], 0, y[j:n-1]) + k[1:n], \\ (y[1:j-1], 0, y[j:n-1]) \end{array} \right\}. \quad (9)$$

Due to the requirement for the $f_{b(i)}$ being bent, we are limited to functions taking an even number of bits as input. The simplest example of a bent function is the inner product.

Eventually we end up with the following instance of the WSN round.

BISON's Round Function

For round keys $k_i \in \mathbb{F}_2^n$ and $w_i \in \mathbb{F}_2^{n-1}$ the round function computes

$$R_{k_i, w_i}(x) := x + f_{b(i)}(w_i + \Phi_{k_i}(x))k_i. \quad (10)$$

where

- Φ_{k_i} is defined as in Eq. (8),
- $f_{b(i)}$ is defined as

$$\begin{aligned} f_{b(i)} : \mathbb{F}_2^{n-1} &\rightarrow \mathbb{F}_2 \\ f_{b(i)}(x) &:= \langle x[1 : (n-1)/2], x[(n+1)/2 : n] \rangle + b(i), \end{aligned}$$

- and $b(i)$ is 0 if $i \leq \frac{r}{2}$ and 1 otherwise.

5.2 Key Schedule

In the i th round, the key schedule has to compute two round keys: $k_i \in \mathbb{F}_2^n$ and $w_i \in \mathbb{F}_2^{n-1}$. We compute those round keys as the states of LFSRs after i clocks, where the initial states are given by a master key K . The master key consists of two parts of n and $n-1$ bits, i.e.

$$K = (k, w) \in \mathbb{F}_2^n \times \mathbb{F}_2^{n-1}.$$

As the all-zero state is a fixed point for any LFSR, we *exclude the zero key* for both k and w . In particular $k = 0$ is obviously a weak key that would result in a ciphertext equal to the plaintext $p = E_{0,w}^r(p)$ for all p , independently of w or of the number of rounds r .

It is well-known that choosing a feedback polynomial of an LFSR to be primitive results in an LFSR of maximal period. Clocking the LFSR then corresponds to multiplication of its state with the companion matrix of this polynomial. Interpreted as elements from the finite field, this is the same as multiplying with a primitive element.

In order to avoid structural attacks, e.g. invariant attacks [39,50,28], as well as the propagation of low-weight inputs, we add round constants c_i to the round key w_i .

These round constants are also derived from the state of an LFSR with the same feedback polynomial as the w_i LFSR, initialized to the unit vector with the least significant bit set. To avoid synchronization with the w_i LFSR, the c_i LFSR clocks backwards.

BISON's Key Schedule

For two primitive polynomials $p_w(x), p_k(x) \in \mathbb{F}_2[x]$ with degrees $\deg(p_w) = n - 1$ and $\deg(p_k) = n$ and the master key $K = (k, w) \in \mathbb{F}_2^n \times \mathbb{F}_2^{n-1}$, $k, w \neq 0$ the key schedule computes the i th round keys as

$$\text{KS}_i : \mathbb{F}_2^n \times \mathbb{F}_2^{n-1} \rightarrow \mathbb{F}_2^n \times \mathbb{F}_2^{n-1}$$

$$\text{KS}_i(k, w) := (C(p_k)^i k, C(p_w)^{-i} e_1 + C(p_w)^i w) = (k_i, c_i + w_i)$$

where $C(\cdot)$ is the companion matrix of the corresponding polynomial, and $0 \leq i < r$.

In [12, Appendix A] we give concrete polynomials for $5 \leq n \leq 129$ -bit block sizes.

As discussed above, this key schedule has the following property, see also Rationale 1.

Lemma 8. *For BISON's key schedule, the following property holds: Any set of n consecutive round keys k_i are linearly independent. Moreover there exist coefficients λ_i such that*

$$k_{n+i} = k_i + \sum_{j=i+1}^{n+i-1} \lambda_j k_j.$$

Proof. To prove this, we start by showing that the above holds for the first n round keys, the general case then follows from a similar argumentation. We need to show that there exists no non-trivial $c_i \in \mathbb{F}_2$ so that $\sum_{i=1}^n c_i C(p_k)^i k = 0$, which is equivalent to showing that there exists no non-trivial $c_i \in \mathbb{F}_2$ so that $\sum_{i=0}^{n-1} c_i C(p_k)^i k = 0$. In this regard, we recall the notion of *minimal polynomial*

of k with respect to $C(p_k)$, defined as the monic polynomial of smallest degree $Q_L(k)(x) = \sum_{i=0}^d q_i x^i \in \mathbb{F}_2[x]$ such that $\sum_{i=0}^d q_i C(p_k)^i k = 0$. Referring to a discussion that has been done for instance in [4], we know that the minimal polynomial of k is a divisor of the minimal polynomial of $C(p_k)$. Since in our case our construction has been made so that this later is equal to p_k which is a primitive polynomial, we deduce that the minimal polynomial of $k \neq 0$ is p_k itself. Since the degree of p_k is equal to n , this prove that the first n keys are linearly independent.

The equation holds, since $p_k(0) = 1$.

6 Security Analysis

As we have already seen, BISON is resistant to differential cryptanalysis. In this section, we argue why BISON is also resistant to other known attacks.

6.1 Linear Cryptanalysis

For linear cryptanalysis, given the fact that BISON is based on a bent function, i. e. a maximally non-linear function, arguing that no linear characteristic with high correlation exist is rather easy. Again, we start by looking at the Fourier coefficients for one round.

From one Round Using the properties of f being bent, we get the following.

Proposition 4. *Let $R_{k,w} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ be the round function as defined in Eq. (10). Then, its LAT consists of the entries*

$$\widehat{R_{k,w}}(\alpha, \beta) = \begin{cases} 2^n & \text{if } \alpha = \beta \text{ and } \langle \beta, k \rangle = 0 \\ \pm 2^{\frac{n+1}{2}} & \text{if } \langle \alpha, k \rangle = 1 \text{ and } \langle \beta, k \rangle = 1 \\ 0 & \text{if } \langle \alpha + \beta, k \rangle = 1 \text{ or } (\alpha \neq \beta \text{ and } \langle \beta, k \rangle = 0) \end{cases} . \quad (11)$$

We prove the proposition in [12, Section 6.1.1, Proposition 4].

To more Rounds When we look at more than one round, we try to approximate the linear hull by looking at the strongest linear trail. As already discussed in Lemma 2, for $r < n$ there are trails with probability one. We now show that any trail's correlation for $r \geq n$ rounds is actually upper bounded by $2^{-\frac{n+1}{2}}$:

Proposition 5. *For $r \geq n$ rounds, the correlation of any non-trivial linear trail for BISON is upper bounded by $2^{-\frac{n+1}{2}}$.*

Proof. It is enough to show the above for any n -round trail. By contradiction, assume there exists a non-trivial trail $\theta = (\theta_0, \dots, \theta_n)$ with correlation one. Following Proposition 4, for every round i the intermediate mask θ_i has to fulfill $\langle \theta_i, k_i \rangle = 0$. Further $\theta_i = \theta_{i+1}$ for $0 \leq i < n$. Because all n round keys are linearly independent, this implies that $\theta_i = 0$, which contradicts our assumption. Thus, in at least one round the second or third case of Eq. (11) has to apply.

It would be nice to be able to say more about the linear hull, analogously to the differential case. However, for the linear cryptanalysis this looks much harder, due to the denser LAT. In our opinion developing a framework where bounding linear hulls is similarly easy as it is for BISON with respect to differentials is a fruitful future research topic.

6.2 Higher-Order Differentials and Algebraic Attacks.

High-order differential attacks, cube attacks, algebraic attacks and integral attacks all make use of non-random behaviour of the ANF of parts of the encryption function. In all these attacks the algebraic degree of (parts of) the encryption function is of particular interest. In this section, we argue that those attacks do not pose a threat to BISON.

We next elaborate in more detail on the algebraic degree of the WSN construction. In particular, we are going to show that the algebraic degree increases at most linearly with the number of rounds. More precisely, if the round function is of degree d , the algebraic degree after r rounds is upper bounded by $r(d - 1) + 1$.

Actually, we are going to consider a slight generalization of the WSN construction and prove the above statement for this generalization.

General Setting Consider an initial state of n bits given as $x = (x_0, \dots, x_{n-1})$ and a sequence of Boolean functions

$$f_i : \mathbb{F}_2^{n+i} \rightarrow \mathbb{F}_2$$

for $0 \leq i < r$. We define a sequence of values y_i by setting $y_0 = f_0(x)$ and

$$y_i = f_i(x_0, \dots, x_{n-1}, y_0, \dots, y_{i-1}),$$

for $1 \leq i < r$. Independently of the exact choice of f_i the degree of any y_ℓ , as a function of x can be upper bounded as stated in the next proposition.

Proposition 6. *Let f_i be a sequence of functions as defined above, such that $\deg(f_i) \leq d$. The degree of y_ℓ at step ℓ seen as a function of the bits of the initial state $x_0, \dots, x_{n-1}$ satisfies*

$$\deg(y_\ell) \leq (d - 1)(\ell + 1) + 1.$$

Moreover, for any $I \subseteq \{0, \dots, \ell\}$,

$$\deg\left(\prod_{i \in I} y_i\right) \leq (d - 1)(\ell + 1) + |I|.$$

Proof. The first assertion is of course a special case of the second one, but we add it for the sake of clarity. We prove the second, more general, statement by induction on ℓ .

Starting with $\ell = 0$, we have to prove that $\deg(y_0) \leq d$, which is obvious, as

$$y_0 = f_0(x_0, \dots, x_{n-1})$$

and $\deg(f_0) \leq d$.

Now, we consider some $I \subseteq \{0, \dots, \ell\}$ and show that

$$\deg\left(\prod_{i \in I} y_i\right) \leq (d-1)(\ell+1) + |I|.$$

We assume that $\ell \in I$, otherwise the result directly follows the induction hypothesis.

Since f_ℓ depends both on $y_0, \dots, y_{\ell-1}$ and x , we decompose it as follows:

$$y_\ell = f_\ell(y_0, \dots, y_{\ell-1}, x) = \sum_{\substack{J \subseteq \{0, \dots, \ell-1\} \\ 0 \leq |J| \leq \min(d, \ell)}} \left(\prod_{j \in J} y_j \right) g_J(x)$$

with $\deg(g_J) \leq d - |J|$ for all J since $\deg(f_\ell) \leq d$.

Then, for $I = \{\ell\} \cup I'$, we look at

$$y_\ell \left(\prod_{i \in I'} y_i \right) = \sum_{\substack{J \subseteq \{0, \dots, \ell-1\} \\ 0 \leq |J| \leq \min(d, \ell)}} \left(\prod_{j \in J \cup I'} y_j \right) g_J(x).$$

From the induction hypothesis, the term of index J in the sum has degree at most

$$\begin{aligned} (d-1)\ell + |J \cup I'| + \deg(g_J) &= (d-1)\ell + |J \cup I'| + d - |J| \\ &\leq (d-1)(\ell+1) + |J \cup I'| - |J| + 1 \\ &\leq (d-1)(\ell+1) + |J| + |I'| - |J| + 1 \\ &\leq (d-1)(\ell+1) + |I|. \end{aligned}$$

Special Case of BISON In the case of BISON, we make use of quadratic functions, and thus Proposition 6 implies that after r rounds the degree is upper bounded by $r+1$ only. Thus, it will take at least $n-2$ rounds before the degree reaches the maximal possible degree of $n-1$. Moreover, due to the construction of WSN, if all component functions of $E_{k,w}^r$ are of degree at most d , there will be at least one component function of $E_{k,w}^{r+n-1}$ of degree at most d . That is, there exist a vector $\beta \in \mathbb{F}_2^n$ such that

$$\langle \beta, E_{k,w}^{r+n-1}(x) \rangle$$

has degree at most d . Namely, for

$$\beta \in \text{span}\{k_r, \dots, k_{r+s}\}^\perp$$

it holds that

$$\deg\left(\langle \beta, E_{k,w}^{r+s}(x) \rangle\right) = \deg\left(\langle \beta, E_{k,w}^r(x) \rangle + \sum_{i=r}^{r+s} \lambda_i \langle \beta, k_i \rangle\right) = \deg(\langle \beta, E_{k,w}^r(x) \rangle).$$

We conclude there exists a component function of $E_{k,w}^{r+s}$ of non-maximal degree, as long as $0 \leq r \leq n - 2$ and $0 \leq s \leq n - 1$. Thus for BISON there will be at least one component function of degree less than $n - 1$ for any number of rounds $0 \leq r \leq 2n - 3$. However, similarly to the case of zero-correlation properties as described in Lemma 3, the vector β is key dependent and thus this property does not directly lead to an attack.

Finally, so far we only discussed upper bounds on the degree, while for arguing security, lower bounds on the degree are more relevant. As it seems very hard (just like for any cipher) to prove such lower bounds, we investigated experimentally how the degree increases in concrete cases. As can be seen in [12, Figure 4] the maximum degree is reached for almost any instance for $n + 5$ rounds. Most importantly, the fraction of instances where it takes more than $n + 2$ rounds decreases with increasing block length n . Moreover, the round function in BISON experimentally behaves with this respect as a random function, as can be seen in [12, Figure 5]. Thus, as the number of rounds is $3n$, we are confident that attacks exploiting the algebraic degree do not pose a threat for BISON.

Besides the WSN construction, a special case of the above proposition worth mentioning is a non linear feedback generator (NLFSR).

Degree of NLFSRs One well-known special case of the above general setting is an NLFSR or, equivalently a maximally unbalanced Feistel cipher, depicted below.

Proposition 6 implies that the degree of any NLFSR increases linearly with the number of rounds. To the best of our knowledge, this is the first time this has been observed in this generality. We like to add that this is in sharp contrast to how the degree increases for SPN ciphers. For SPN ciphers the degree usually increases exponentially until a certain threshold is reached [11].

6.3 Other attacks

We briefly discuss other cryptanalytic attacks.

Impossible Differentials In Lemma 5 and Corollary 2, we discuss that every output difference is possible after more than n rounds. Consequently, there are no impossible differentials for BISON.

Truncated Differentials Due to our strong bounds on differentials it seems very unlikely that any strong truncated differential exists.

Zero Correlation Linear Cryptanalysis In Lemma 3 we already discussed generic zero correlation linear hulls for the WSN construction. Depending on the actual key used, this technique may be used to construct a one-round-longer zero-correlation trail. For this, we need two *distinct* elements $\alpha \in \langle k_1, \dots, k_{n-1} \rangle^\perp$, $\beta \in \langle k_n, \dots, k_{2n-2} \rangle^\perp$, and construct the trail analogously to Lemma 3 (which may not exist, due to the key dependency).

Invariant Attacks For an invariant attack, we need a Boolean function g , s. t. $g(x) + g(E_{k,w}^r(x))$ is constant *for all* x and some *weak keys* (k, w) . As the encryption of any message is basically this message with some of the round keys added, key addition is the only operation which is performed. It has been shown in [4, Proposition 1] that any g which is invariant for a linear layer followed by the addition of the round key k_i as well as for the same up to addition of a different k_j , has a linear space containing $k_i + k_j$. In the case of the linear layer being the identity, the linear space actually contains also the k_i and k_j (by definition).

Thus, the linear space of any invariant for our construction has to contain $\text{span}\{k_1, \dots, k_{3n}\}$ which is obviously the full space $\mathbb{F}_2^n$. Following the results of [4], there are thus no invariant subspace or nonlinear invariant attack on BISON.

Related-Key Attacks In generic related-key attacks, the attacker is also allowed to exploit encryptions under a related, that is $k' = f(k)$, key – in the following, we restrict our analysis to the case where f is the addition with a constant. That is, the attacker cannot only request $E_{k,w}(x)$, and $E_{k,w}(x + \alpha)$, but also $E_{k+\beta, w+\beta'}(x)$ or $E_{k+\beta, w+\beta'}(x + \alpha)$, for α (difference in the input x), β (difference in the key k) and β' (difference in the key w) of her choice. As $\beta = \beta' = 0$ would result in the standard differential scenario, we exclude it for the remainder of this discussion. Similar, $\beta = k$ results in $\Phi_{k+\beta} = \Phi_0$, which we did not define, thus we also skip this case and refer to the fact that if an attacker chooses $\beta = k$, she basically already has guessed the secret key correctly.

For BISON, the following proposition holds.

Proposition 7. *For r rounds, the probability of any related-key differential characteristic for BISON, averaged over all whiting key sequences $(w_1, \dots, w_r)$, is upper bounded by $(\frac{3}{4})^r$.*

For more details and a proof of the proposition, see [12, Section 6.3.5, Proposition 7].

Further Observations During the design process, we observed the following interesting point: For sparse master keys k and w and message m , e. g. $k = w = m = 1$, in the first few rounds, nothing happens. This is mainly due to the choice of sparse key schedule polynomials p_w and p_k and the fact that f_0 outputs 0 if only one bit in its input is set (as $\langle 0, x \rangle = 0$ for any x).

To the best of our knowledge, this observation cannot be exploited in an attack.

Experimental Results We conducted experiments on small-scale versions of BISON with $n = 5$. The DDTs and LATs, depicted using the “Jackson Pollock representation” [8], for one to ten rounds are listed in [12, Appendix B]. In [12, Appendix B.1] one can see that the two cases of averaging over all possible w_i and choosing a fixed w_i results in very similar differential behaviors. Additionally, after $5 = n$ rounds, the plots do not change much.

The results in the linear case, see [12, Appendix B.2], are quite similar. The major difference here, is the comparable bigger entries for a fixed w_i . Nonetheless, most important is that there are no high entries in the average LAT which would imply a strong linear approximation for many keys. Additionally one also expects for a random permutation not too small LAT entries. Note that one can well observe the probability-one approximation for $4 = n - 1$ rounds (lower right corner of the corresponding plot).

7 Conclusion

Efficiency of symmetric ciphers have been significantly improved further and further, in particular within the trend of lightweight cryptography. However, when it comes to arguing about the security of ciphers, the progress is rather limited and the arguments basically did not get easier nor stronger since the development of the AES. In our opinion it might be worth shifting the focus to improving security arguments for new designs rather than (incrementally) improving efficiency. We see BISON as a first step in this direction.

With our instance for the WSN construction and its strong resistance to differential cryptanalysis, this framework emerges as an interesting possibility to design block ciphers. Unfortunately, we are not able to give better than normal arguments for the resistance to linear cryptanalysis. It is thus an interesting question, if one can find a similar instance of the WSN construction for which comparable strong arguments for the later type of cryptanalysis exist.

Alternative designs might also be worth looking at. For example many constructions for bent functions are known and could thus be examined as alternatives for the scalar product used in BISON. One might also look for a less algebraic design – but we do not yet see how this would improve or ease the analysis or implementation of an instance.

Finally, for an initial discussion of implementation figures, see [12, Section 7]. Another line of future work in this direction is the in-depth analysis of implementation optimizations and side channel-resistance of BISON.

Acknowledgments

We would like to thank the anonymous reviewers and Christof Beierle for their helpful comments, and Lucas Hartmann for the artistic design of BISON.

References

1. Abdelraheem, M.A., Ågren, M., Beelen, P., Leander, G.: On the distribution of linear biases: Three instructive examples. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 50–67. Springer, Heidelberg (Aug 2012)
2. Advanced Encryption Standard (AES) (Nov 2001)
3. Andreeva, E., Bogdanov, A., Dodis, Y., Mennink, B., Steinberger, J.P.: On the indifferentiability of key-alternating ciphers. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 531–550. Springer, Heidelberg (Aug 2013)
4. Beierle, C., Canteaut, A., Leander, G., Rotella, Y.: Proving resistance against invariant attacks: How to choose the round constants. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 647–678. Springer, Heidelberg (Aug 2017)
5. Biham, E., Shamir, A.: Differential cryptanalysis of DES-like cryptosystems. In: Menezes, A.J., Vanstone, S.A. (eds.) CRYPTO’90. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (Aug 1991)
6. Biryukov, A., Khovratovich, D.: Related-key cryptanalysis of the full AES-192 and AES-256. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 1–18. Springer, Heidelberg (Dec 2009)
7. Biryukov, A., Khovratovich, D., Nikolic, I.: Distinguisher and related-key attack on the full AES-256. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 231–249. Springer, Heidelberg (Aug 2009)
8. Biryukov, A., Perrin, L.: On reverse-engineering S-boxes with hidden design criteria or structure. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 116–140. Springer, Heidelberg (Aug 2015)
9. Blondeau, C., Nyberg, K.: Improved parameter estimates for correlation and capacity deviates in linear cryptanalysis. IACR Trans. Symm. Cryptol. 2016(2), 162–191 (2016), <http://tosc.iacr.org/index.php/ToSC/article/view/570>
10. Bogdanov, A., Knudsen, L.R., Leander, G., Standaert, F.X., Steinberger, J.P., Tischhauser, E.: Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 45–62. Springer, Heidelberg (Apr 2012)
11. Boura, C., Canteaut, A., De Cannière, C.: Higher-order differential properties of Keccak and Luffa. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 252–269. Springer, Heidelberg (Feb 2011)
12. Canteaut, A., Lallemand, V., Leander, G., Neumann, P., Wiemer, F.: BISON – instantiating the whitened swap-or-not construction. Cryptology ePrint Archive, Report 2018/1011 (2018)
13. Canteaut, A., Roué, J.: On the behaviors of affine equivalent sboxes regarding differential and linear attacks. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 45–74. Springer, Heidelberg (Apr 2015)
14. Carlet, C.: Boolean Functions for Cryptography and Error Correcting Codes. In: Crama, Y., Hammer, P. (eds.) Boolean Methods and Models. Cambridge University Press (2007)
15. Chen, S., Steinberger, J.P.: Tight security bounds for key-alternating ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 327–350. Springer, Heidelberg (May 2014)

16. Daemen, J.: Cipher and hash function design, strategies based on linear and differential cryptanalysis, PhD Thesis. K.U.Leuven (1995), <http://jda.noekeon.org/>
17. Daemen, J., Govaerts, R., Vandewalle, J.: Block ciphers based on modular arithmetic. In: Wolfowicz, W. (ed.) *State and Progress in the Research of Cryptography*, pp. 80–89. Fondazione Ugo Bordoni (1993)
18. Daemen, J., Govaerts, R., Vandewalle, J.: Correlation matrices. In: Preneel, B. (ed.) *FSE'94*. LNCS, vol. 1008, pp. 275–285. Springer, Heidelberg (Dec 1995)
19. Daemen, J., Rijmen, V.: The block cipher rijndael. In: *CARDIS'98*. LNCS, vol. 1820, pp. 277–284. Springer (1998)
20. Daemen, J., Rijmen, V.: The wide trail design strategy. In: Honary, B. (ed.) *8th IMA International Conference on Cryptography and Coding*. LNCS, vol. 2260, pp. 222–238. Springer, Heidelberg (Dec 2001)
21. Daemen, J., Rijmen, V.: *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography, Springer (2002)
22. Daemen, J., Rijmen, V.: Understanding two-round differentials in AES. In: Prisco, R.D., Yung, M. (eds.) *SCN 06*. LNCS, vol. 4116, pp. 78–94. Springer, Heidelberg (Sep 2006)
23. Derbez, P., Fouque, P.A., Jean, J.: Improved key recovery attacks on reduced-round AES in the single-key setting. In: Johansson, T., Nguyen, P.Q. (eds.) *EUROCRYPT 2013*. LNCS, vol. 7881, pp. 371–387. Springer, Heidelberg (May 2013)
24. Dillon, J.F.: A survey of bent functions. *The NSA technical journal* 191, 215 (1972)
25. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology* 10(3), 151–162 (Jun 1997)
26. Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D., Whiting, D.: Improved cryptanalysis of Rijndael. In: Schneier, B. (ed.) *FSE 2000*. LNCS, vol. 1978, pp. 213–230. Springer, Heidelberg (Apr 2001)
27. Gilbert, H., Minier, M.: A collision attack on 7 rounds of rijndael. In: *AES Candidate Conference*. vol. 230, p. 241 (2000)
28. Grassi, L., Rechberger, C., Rønjom, S.: Subspace trail cryptanalysis and its applications to AES. *IACR Trans. Symm. Cryptol.* 2016(2), 192–225 (2016), <http://tosc.iacr.org/index.php/ToSC/article/view/571>
29. Grassi, L., Rechberger, C., Rønjom, S.: A new structural-differential property of 5-round AES. In: Coron, J., Nielsen, J.B. (eds.) *EUROCRYPT 2017, Part II*. LNCS, vol. 10211, pp. 289–317. Springer, Heidelberg (Apr / May 2017)
30. Guo, C., Lin, D.: On the indistinguishability of key-alternating Feistel ciphers with no key derivation. In: Dodis, Y., Nielsen, J.B. (eds.) *TCC 2015, Part I*. LNCS, vol. 9014, pp. 110–133. Springer, Heidelberg (Mar 2015)
31. Hoang, V.T., Morris, B., Rogaway, P.: An enciphering scheme based on a card shuffle. In: Safavi-Naini, R., Canetti, R. (eds.) *CRYPTO 2012*. LNCS, vol. 7417, pp. 1–13. Springer, Heidelberg (Aug 2012)
32. Hoang, V.T., Tessaro, S.: Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In: Robshaw, M., Katz, J. (eds.) *CRYPTO 2016, Part I*. LNCS, vol. 9814, pp. 3–32. Springer, Heidelberg (Aug 2016)
33. Hong, S., Lee, S., Lim, J., Sung, J., Cheon, D.H., Cho, I.: Provable security against differential and linear cryptanalysis for the SPN structure. In: Schneier, B. (ed.) *FSE 2000*. LNCS, vol. 1978, pp. 273–283. Springer, Heidelberg (Apr 2001)
34. Jakobsen, T., Knudsen, L.R.: The interpolation attack on block ciphers. In: Biham, E. (ed.) *FSE'97*. LNCS, vol. 1267, pp. 28–40. Springer, Heidelberg (Jan 1997)

35. Keliher, L., Sui, J.: Exact maximum expected differential and linear probability for two-round advanced encryption standard. *IET Information Security* 1(2), 53–57 (2007)
36. Kranz, T., Leander, G., Wiemer, F.: Linear cryptanalysis: Key schedules and tweakable block ciphers. *IACR Trans. Symm. Cryptol.* 2017(1), 474–505 (2017)
37. Lai, X., Massey, J.L., Murphy, S.: Markov ciphers and differential cryptanalysis. In: Davies, D.W. (ed.) *EUROCRYPT’91*. LNCS, vol. 547, pp. 17–38. Springer, Heidelberg (Apr 1991)
38. Lampe, R., Seurin, Y.: Security analysis of key-alternating Feistel ciphers. In: Cid, C., Rechberger, C. (eds.) *FSE 2014*. LNCS, vol. 8540, pp. 243–264. Springer, Heidelberg (Mar 2015)
39. Leander, G., Abdelraheem, M.A., AlKhazimi, H., Zenner, E.: A cryptanalysis of PRINTcipher: The invariant subspace attack. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 206–221. Springer, Heidelberg (Aug 2011)
40. Matsui, M.: On correlation between the order of S-boxes and the strength of DES. In: Santis, A.D. (ed.) *EUROCRYPT’94*. LNCS, vol. 950, pp. 366–375. Springer, Heidelberg (May 1995)
41. Meier, W., Staffelbach, O.: Nonlinearity criteria for cryptographic functions. In: Quisquater, J.J., Vandewalle, J. (eds.) *EUROCRYPT’89*. LNCS, vol. 434, pp. 549–562. Springer, Heidelberg (Apr 1990)
42. Miracle, S., Yilek, S.: Cycle slicer: An algorithm for building permutations on special domains. *Cryptology ePrint Archive*, Report 2017/873 (2017), <http://eprint.iacr.org/2017/873>
43. Nyberg, K.: Linear approximation of block ciphers (rump session). In: Santis, A.D. (ed.) *EUROCRYPT’94*. LNCS, vol. 950, pp. 439–444. Springer, Heidelberg (May 1995)
44. Nyberg, K.: “provable” security against differential and linear cryptanalysis (invited talk). In: Canteaut, A. (ed.) *FSE 2012*. LNCS, vol. 7549, pp. 1–8. Springer, Heidelberg (Mar 2012)
45. Nyberg, K., Knudsen, L.R.: Provable security against a differential attack. *Journal of Cryptology* 8(1), 27–37 (Dec 1995)
46. Park, S., Sung, S.H., Lee, S., Lim, J.: Improving the upper bound on the maximum differential and the maximum linear Hull probability for SPN structures and AES. In: Johansson, T. (ed.) *FSE 2003*. LNCS, vol. 2887, pp. 247–260. Springer, Heidelberg (Feb 2003)
47. Rothaus, O.S.: On ‘bent’ functions. *Journal of Combinatorial Theory, Series A* 20(3), 300–305 (1976)
48. Tessaro, S.: Optimally secure block ciphers from ideal primitives. In: Iwata, T., Cheon, J.H. (eds.) *ASIACRYPT 2015, Part II*. LNCS, vol. 9453, pp. 437–462. Springer, Heidelberg (Nov / Dec 2015)
49. Tessaro, S.: Optimally secure block ciphers from ideal primitives. *Cryptology ePrint Archive*, Report 2015/868 (2015), <http://eprint.iacr.org/2015/868>
50. Todo, Y., Leander, G., Sasaki, Y.: Nonlinear invariant attack - practical attack on full SCREAM, iSCREAM, and Midori64. In: Cheon, J.H., Takagi, T. (eds.) *ASIACRYPT 2016, Part II*. LNCS, vol. 10032, pp. 3–33. Springer, Heidelberg (Dec 2016)
51. Vaudenay, S.: Provable security for block ciphers by decorrelation. In: *STACS’98*. LNCS, vol. 1373, pp. 249–275. Springer (1998)
52. Vaudenay, S.: The end of encryption based on card shuffling. *CRYPTO 2012 Rump Session* (2012), crypto.2012.rump.cr.yp.to PDF Link