

# PRINCE – A Low-Latency Block Cipher for Pervasive Computing Applications

## Extended Abstract\*

Julia Borghoff<sup>1,\*\*</sup>, Anne Canteaut<sup>1,2,\*\*\*</sup>, Tim Güneysu<sup>3</sup>, Elif Bilge Kavun<sup>3</sup>,  
Miroslav Knezevic<sup>4</sup>, Lars R. Knudsen<sup>1</sup>, Gregor Leander<sup>1,†</sup>, Ventsislav Nikov<sup>4</sup>,  
Christof Paar<sup>3</sup>, Christian Rechberger<sup>1</sup>, Peter Rombouts<sup>4</sup>, Søren S. Thomsen<sup>1</sup>,  
and Tolga Yalçın<sup>3</sup>

<sup>1</sup> Technical University of Denmark

<sup>2</sup> INRIA, Paris-Rocquencourt, France

<sup>3</sup> Ruhr-University Bochum, Germany

<sup>4</sup> NXP Semiconductors, Leuven, Belgium

**Abstract.** This paper presents a block cipher that is optimized with respect to latency when implemented in hardware. Such ciphers are desirable for many future pervasive applications with real-time security needs. Our cipher, named PRINCE, allows encryption of data within one clock cycle with a very competitive chip area compared to known solutions. The fully unrolled fashion in which such algorithms need to be implemented calls for innovative design choices. The number of rounds must be moderate and rounds must have short delays in hardware. At the same time, the traditional need that a cipher has to be iterative with very similar round functions disappears, an observation that increases the design space for the algorithm. An important further requirement is that realizing decryption and encryption results in minimum additional costs. PRINCE is designed in such a way that the overhead for decryption on top of encryption is negligible. More precisely for our cipher it holds that decryption for one key corresponds to encryption with a related key. This property we refer to as  $\alpha$ -reflection is of independent interest and we prove its soundness against generic attacks.

## 1 Introduction

The area of lightweight cryptography, i.e., ciphers with particularly low implementation costs, has drawn considerable attention over the last years. Among the

---

\* Due to page limitations, several details are omitted in this proceedings version. A full version is available at [10].

\*\* Supported by the Danish research council for Technology and Production Sciences, Grant No.10-093667.

\*\*\* Partially supported by DGA/DS under Contract 2011.60.055.

† Partially supported by the Danish-Chinese Center for Applications of Algebraic Geometry in Coding Theory and Cryptography (Danish National Research Foundation and the National Science Foundation of China, Grant No.11061130539).

best studied algorithms are the block ciphers CLEFIA, Hight, KATAN, KTANTAN, Klein, mCrypton, LED, Piccolo and PRESENT [33,23,15,20,29,21,32,9], as well as the stream ciphers Grain, Mickey, and Trivium [22,2,16]. Particular interest in lightweight symmetric ciphers is coming from industry, as becoming evident in the adoption of CLEFIA and PRESENT in the ISO/IEC Standard 29192-2. The dominant metric according to which the majority of lightweight ciphers have been optimized is chip area, typically measured in gate equivalences (GE), i.e., the cipher area normalized to the area of a 2-input NAND gate in a given standard cell library. This is certainly a valid optimization objective in cases where there are extremely tight power or cost constraints, in particular passive RFID tags. However, depending on the application, there are several other implementation parameters according to which a cipher should have lightweight characteristics. There are several important applications for which a low-latency encryption and instant response time is highly desirable, such as instant authentication or block-wise read/write access to memory devices, e.g., solid-state hard disks. There are also embedded applications where current block ciphers in multiple-clock architectures could be sufficiently fast, but the needed high clock rates are not supported by the system. For instance, in many FPGA designs clock rates above 200 MHz are often difficult to realize. It can also be anticipated that given the ongoing growth of pervasive computing, there will be many more future embedded systems that require low-latency encryption, especially applications with real-time requirements, e.g., in the automotive domain. Moreover, [27] as well as [24] show that low-latency goes hand in hand with energy efficiency, another crucial criterion in many (other) applications.

For all these cases, we like to have symmetric ciphers that can *instantaneously* encrypt a given plaintext, i.e., the entire encryption *and* decryption should take place within the shortest possible delay. This seemingly simple problem poses a considerable challenge with today's cryptosystems — in particular if encryption and decryption should both be available on a given platform. Software implementations of virtually all strong ciphers take hundreds or thousands of clock cycles, making them ill suited for a designer aiming for low-latency cryptography. In the case of stream ciphers implemented in hardware, the high number of clock cycles for the initialization phase makes them not suitable for this task, especially when secret keys need to be regularly changed. Moreover, if we want to encrypt small blocks selected at random (e.g., encryption of sectors on solid-state disks), stream ciphers are not suited<sup>1</sup>. This leaves block ciphers as the remaining viable solution. However, the round-based, i.e., iterative, nature of virtually all existing block ciphers, as shown for the case of AES, makes low-latency implementation a non-trivial task. A round-based hardware architecture of the AES-128 requires ten clock cycles to output a ciphertext which we do not consider instantaneous as it is still too long for some applications. As a remedy, the ten rounds can be loop-unrolled, i.e., the circuit that realizes the single round is repeated ten times. Now, the cipher returns a ciphertext within a single clock cycle — but at the cost of a very long critical path. This yields a very slow absolute response time

---

<sup>1</sup> A possible exception are random-access stream ciphers such as Salsa [5].

and clock frequencies, e.g., in the range of a few MHz. Furthermore, the unrolled architecture has a high gate count in the range of several tens of thousand GE, implying a high power consumption and costs. Both features are undesirable, especially if one considers that many of the applications for instantaneous ciphers are in the embedded domain. Following the same motivation and reasoning as above [27] compares several lightweight ciphers with respect to latency and as a conclusion calls for new designs that are optimized for low-latency.

**Our Contribution.** Based on the above discussion our goal is to design a new block cipher which is optimized with respect to the following criteria if implemented in hardware:

1. The cipher can perform instantaneous encryption, a ciphertext is computed within a single clock cycle. There is no warm-up phase.
2. If implemented in modern chip technology, low delays resulting in moderately high clock rates can be achieved.
3. The hardware costs are moderate (i.e., considerably lower than fully unrolled versions of AES or PRESENT).
4. Encryption and decryption should both be possible with low costs and overhead.

We would like to remark that existing lightweight ciphers such as PRESENT do *not* fulfill Criteria 2 and 3 (low delay, small area) due to their large number of rounds. In order to fulfill Criterion 4, one needs to design a cipher for which decryption and encryption use (almost) identical pieces of hardware. This is an important requirement since the unrolled nature of instantaneous ciphers leads to circuits which are large and it is thus clearly advantageous if large parts of the implementation can be used both for encryption and decryption.

Besides designing a new lightweight cipher that is for the first time optimized with respect to the goals above, PRINCE has several innovative features that we like to highlight.

First, a fully unrolled design increases the possible design choices enormously. With a fully unrolled cipher, the traditional need that a cipher has to be iterative with very similar round functions disappears. This in turn allows us to efficiently implement a cipher where decryption with one key corresponds to encryption with a related key. This property we refer to as  $\alpha$ -reflection is of independent interest and we prove its soundness against generic attacks. As a consequence, the overhead of implementing decryption over encryption becomes negligible. Note that previous approaches to minimizing the overhead of decryption over encryption, for example in the ciphers NOEKEON and ICEBERG usually require multiplexer in each round. While for a round-based implementation this does not make a difference, our approach is clearly preferable for a fully unrolled implementation, as we require multiplexer only once at the beginning of the circuit.

Another difference to known lightweight ciphers like PRESENT is that we balance the cost of an Sbox-layer and the linear layer. As it turns out optimizing

the cost of the Sbox chosen has a major influence on the overall cost of the cipher. As an Sbox that performs well in one technology does not necessarily perform well in another technology, we propose the PRINCE-family of ciphers that allows to freely choose the Sbox within a (large) set of Sboxes fulfilling certain criteria. Our choice for the linear layer can be seen as being inbetween a bit-permutation layer PRESENT (implemented with wires only) and AES (implemented with considerable combinatorial logic). With the expense of only 2 additional XOR-gates per bit over a simple bit-permutation layer, we achieve an almost-MDS property that helps to prove much better bounds against various classes of attacks and in turn allows to significantly reduce the number of rounds and hence latency.

As a result, PRINCE compares very favorable to existing ciphers. For the same time constraints and technologies, PRINCE uses 6–7 times less area than PRESENT-80 and 14–15 times less area than AES-128. In addition to this, our design uses about 4–5 times less area than other ciphers in the literature (see Section 5 and in particular Tables 1 and 2 for a detailed comparison and technology details). To facilitate further study and fairer comparisons, we also report synthesis results using the open-source standard-cell library NANGATE [30]. We also like to mention that, although this is not the main objective of the cipher, PRINCE compares reasonably well to other lightweight ciphers when implemented in a round-based fashion (see [10]).

We believe that our consideration can be of major value for industry and can at the same time stimulate the scientific community to pursue research on lightweight ciphers with different optimization goals.

**Organization of the Paper.** We introduce an instance of PRINCE-family of ciphers and state our security claims in Section 2. Design decisions are discussed in Section 3 where we also describe the entire PRINCE-family. We provide security proofs and evaluations considering cryptanalytical attacks in Section 4. In Section 5 we finally present implementation results and comparisons with other lightweight ciphers for a range of hardware technologies. For further details, including a detailed security analysis against standard attacks as well as test vectors, we refer to [10].

## 2 Cipher Description

PRINCE is a 64-bit block cipher with a 128-bit key. The key is split into two parts of 64 bits each,

$$k = k_0 || k_1$$

and extended to 192 bits by the mapping


$$(k_0 || k_1) \rightarrow (k_0 || k'_0 || k_1) := (k_0 || (k_0 \ggg 1) \oplus (k_0 \ggg 63) || k_1).$$

PRINCE is based on the so-called  $FX$  construction [7,25]: the first two subkeys  $k_0$  and  $k'_0$  are used as whitening keys, while the key  $k_1$  is the 64-bit key for a 12-round block cipher we refer to as  $\text{PRINCE}_{\text{core}}$ . We provide test vectors in the full version of the paper [10].



**Specification of  $\text{PRINCE}_{\text{core}}$ .**

The whole encryption process of  $\text{PRINCE}_{\text{core}}$  is depicted below.



Each round of  $\text{PRINCE}_{\text{core}}$  consist of a key addition, an Sbox-layer, a linear layer, and the addition of a round constant.

**$k_i$ -add.** Here the 64-bit state is xored with the 64-bit subkey.

**S-Layer.** The cipher uses one 4-bit Sbox. The action of the Sbox in hexadecimal notation is given by the following table.

$x$	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
$S[x]$	B	F	3	2	A	C	9	1	6	7	8	0	E	5	D	4

**The Matrices:  $M/M'$ -layer.** In the  $M$  and  $M'$ -layer the 64-bit state is multiplied with a  $64 \times 64$  matrix  $M$  (resp.  $M'$ ) defined in Section 3.3.

**$RC_i$ -add.** In the  $RC_i$ -add step a 64-bit round constant is xored with the state. We define the constants used below (in hex notation)

$RC_0$	0000000000000000
$RC_1$	13198a2e03707344
$RC_2$	a4093822299f31d0
$RC_3$	082efa98ec4e6c89
$RC_4$	452821e638d01377
$RC_5$	be5466cf34e90c6c
$RC_6$	7ef84f78fd955cb1
$RC_7$	85840851f1ac43aa
$RC_8$	c882d32f25323c54
$RC_9$	64a51195e0e3610d
$RC_{10}$	d3b5a399ca0c2399
$RC_{11}$	c0ac29b7c97c50dd

Note that, for all  $0 \leq i \leq 11$ ,  $RC_i \oplus RC_{11-i}$  is the constant  $\alpha = c0ac29b7c97c50dd$ ,  $RC_0 = 0$  and that  $RC_1, \dots, RC_5$  and  $\alpha$  are derived from the fraction part of  $\pi = 3.141\dots$

From the fact that the round constants satisfy  $RC_i \oplus RC_{11-i} = \alpha$  and that  $M'$  is an involution, we deduce that the core cipher is such that the inverse of  $\text{PRINCE}_{core}$  parametrized with  $k$  is equal to  $\text{PRINCE}_{core}$  parametrized with  $(k \oplus \alpha)$ . We call this property of  $\text{PRINCE}_{core}$  the  $\alpha$ -reflection property. It follows that, for any expanded key  $(k_0 || k'_0 || k_1)$ ,

$$D_{(k_0 || k'_0 || k_1)}(\cdot) = E_{(k'_0 || k_0 || k_1 \oplus \alpha)}(\cdot)$$

where  $\alpha$  is the 64-bit constant  $\alpha = c0ac29b7c97c50dd$ . Thus, for decryption one only has to do a very cheap change to the master key and afterwards reuse the exact same circuit.

**Security Claims.** For an adversary that is able to acquire  $2^n$  plaintext/ciphertext pairs in a model with a single fixed unknown key  $k$ , we claim that the effort to find the key is not significantly less expensive than  $2^{127-n}$  calls to the encryption or decryption function. In Section 4.1 we give a bound matching this claim in the ideal cipher model that does consider the special relation between the encryption and decryption operations. One way to interpret this is that any attack violating our security claim will have to use more properties of the cipher than the relation between the encryption and decryption operations.

We explicitly state that we do not have claims in related-key or known- and chosen-key models as we do not consider them to be relevant for the intended use cases. In particular, as for any cipher based on the  $FX$  construction or on the Even-Mansour scheme [18], there exists a trivial distinguisher for PRINCE in the related-key model: for any difference  $\Delta$ , the ciphertexts corresponding to  $m$  and  $(m \oplus \Delta)$  encrypted under keys  $(k_0 || k_1)$  and  $((k_0 \oplus \Delta) || k_1)$  respectively, differ from  $((\Delta \ggg 1) \oplus (\Delta \ggg 63))$  with probability 1.

**Reduced Versions.** Many classes of cryptanalytic attacks become more difficult with an increased number of rounds. In order to facilitate third-party

cryptanalysis and estimate the security margin, reduced-round variants need to be considered. We encourage to study round-reduced variants of PRINCE where the symmetry around the middle is kept, and rounds are added in an inside-out fashion, i.e. for every additional round  $\mathfrak{R}_i$  its inverse is also added. Another natural way to reduce PRINCE is to consider the cipher without the key whitening layer,  $\text{PRINCE}_{\text{core}}$ .

### 3 Design Decisions

In this section we explain our design decisions. First note that an SP-network is preferable over a Feistel-cipher, since a Feistel-cipher operates only on half the state resulting often in a higher number of rounds. In order to minimize the number of rounds and still achieve security against linear and differential attacks, we adopted the wide-trail strategy [11]. As not all round functions have to be identical for a cipher aiming for a fully unrolled implementation as PRINCE, it is very tempting to directly use the concept of code-concatenation [13] to achieve a high number of active Sboxes over 4 rounds of the cipher. However, not only a serial implementation benefits from similar round functions. It is also very helpful for ensuring a minimum number of active Sboxes. Assume that, using the code-concatenation approach, one can ensure that rounds  $R_i$  to  $R_{i+3}$  have at least 16 active Sboxes. While this is nice, the problem is that it does not ensure that rounds  $R_{i-1}$  to  $R_{i+2}$  or  $R_{i+1}$  to  $R_{i+4}$  have 16 active Sboxes as well if the individual rounds are very different in nature. We therefore decided to follow a design that on one hand allows to use the freedom given by a fully enrolled design and on the other hand still keeps the round functions similar enough to prove some bounds on the resistance against linear and differential attacks.

In this context, one of the main features of the design is that decryption can be implemented on top of encryption with a minimal overhead. This is achieved by designing a cipher which is symmetric around the middle round, a very simple key scheduling, and a special choice of round constants.

#### 3.1 Aligning Encryption with Decryption

The use of a core cipher having the  $\alpha$ -reflection property, with two additional whitening keys, offers a nice alternative to the usual design strategy which consists in using involucional components — Noekeon [12], Khazad [4], Anubis [3], Iceberg [35] or SEA [34] are some examples of such ciphers with involucional components. Actually, the general construction used in PRINCE has the following advantages:

- It allows a much larger choice of Sboxes, which may lead to a lower implementation cost, since the Sbox is not required to be an involution. It is worth noticing that the fact that both the Sbox and its inverse are involved in the encryption function does not affect the cost of the fully-unrolled implementations we consider;

- In ciphers with involutorial components, the overhead due to the implementation of the inverse key scheduling can be reduced by adding some symmetry in the subkey sequence. But this may introduce weak keys or potential slide attacks. The fact that all components are involutions may also introduce some regularities in the cyclic structure of the cipher which can be exploited in some attacks [6]. The resistance of PRINCE to this type of attacks will be extensively discussed in Section 4.2.
- It is an open problem to prove the security of ciphers with ideal, involutorial components against generic attacks. We show in Section 4.1 that ciphers with the  $\alpha$ -reflection property (for  $\alpha \neq 0$ ) has a proof of security similar to that of the  $FX$  construction.
- Previous approaches to minimizing the overhead of decryption over encryption usually require multiplexer in each round while our approach requires multiplexer only once at the beginning of the circuit.

### 3.2 The PRINCE-Family: Choosing the Sbox

As discussed in more detail in Section 5, the cost of the Sbox, i.e., its area and critical path, is a substantial part of the overall cost. Thus, choosing an Sbox which minimizes those costs is crucial for obtaining competitive results. As the cost of an Sbox depends on various parameters, such as the technology, the synthesis tool, and the library used, one cannot expect that there is one optimal Sbox for all environments. In fact, in order to achieve optimal results it is preferable to choose your favorite Sbox. In order to ensure the security of the resulting design, an Sbox  $S : \mathbb{F}_2^4 \rightarrow \mathbb{F}_2^4$  for the PRINCE-Family has to fulfill the following criteria.


1. The maximal probability of a differential is  $1/4$
2. There are exactly 15 differentials with probability  $1/4$ .
3. The maximal absolute bias of a linear approximation is  $1/4$ .
4. There are exactly 30 linear approximations with absolute bias  $1/4$ .
5. Each of the 15 non-zero component functions has algebraic degree 3.

As it can be deduced for example from [28] up to affine equivalence there are only 8 Sboxes fulfilling those criteria. Thus, another way of defining an Sbox for the PRINCE-Family is to say that it has to be affine equivalent to one of the eight Sboxes  $S_i$  given in the full version of this paper [10].

### 3.3 The Linear Layer

In the  $M$  and  $M'$ -layer the 64-bit state is multiplied with a  $64 \times 64$  matrix  $M$  (resp.  $M'$ ) defined below. We have different requirements for the two different linear layers. The  $M'$ -layer is only used in the middle round, thus  $M'$  has to be an involution to ensure the  $\alpha$ -reflection property. This requirement does not apply for the  $M$ -layer used in the round functions. Here we want to ensure full diffusion after two rounds. To achieve this we combine the  $M'$ -mapping with an application of matrix  $SR$  which behaves like the AES shift rows and permutes the 16 nibbles in the following way.





Thus  $M = SR \circ M'$ .

Additionally the implementation costs should be minimized, meaning that the number of ones in the matrices  $M'$  and  $M$  should be minimal, while at the same time it should be guaranteed that at least 16 Sboxes are active in 4 consecutive rounds (see full version [10] for details). Thus, trivially each output bit of an Sbox has to influence 3 Sboxes in the next round and therefore the minimum number of ones per row and column is 3. Thus we can use the following four  $4 \times 4$  matrices as building blocks for the  $M'$ -layer.

$$M_0 = \begin{pmatrix} 0000 \\ 0100 \\ 0010 \\ 0001 \end{pmatrix}, M_1 = \begin{pmatrix} 1000 \\ 0000 \\ 0010 \\ 0001 \end{pmatrix}, M_2 = \begin{pmatrix} 1000 \\ 0100 \\ 0000 \\ 0001 \end{pmatrix}, M_3 = \begin{pmatrix} 1000 \\ 0100 \\ 0010 \\ 0000 \end{pmatrix}$$

In the next step we generate a  $4 \times 4$  block matrix  $\hat{M}$  where each row and column is a permutation of the four  $4 \times 4$  matrices  $M_0, \dots, M_3$ . The row permutations are chosen such that we obtain a symmetric block matrix. The choice of the building blocks and the symmetric structure ensures that the resulting  $16 \times 16$  matrix is an involution. We define

$$\hat{M}^{(0)} = \begin{pmatrix} M_0 & M_1 & M_2 & M_3 \\ M_1 & M_2 & M_3 & M_0 \\ M_2 & M_3 & M_0 & M_1 \\ M_3 & M_0 & M_1 & M_2 \end{pmatrix} \quad \hat{M}^{(1)} = \begin{pmatrix} M_1 & M_2 & M_3 & M_0 \\ M_2 & M_3 & M_0 & M_1 \\ M_3 & M_0 & M_1 & M_2 \\ M_0 & M_1 & M_2 & M_3 \end{pmatrix}.$$

In order to obtain a permutation for the full 64-bit state we construct a  $64 \times 64$  block diagonal matrix  $M'$  with  $(\hat{M}^{(0)}, \hat{M}^{(1)}, \hat{M}^{(1)}, \hat{M}^{(0)})$  as diagonal blocks. The matrix  $M'$  is an involution with  $2^{32}$  fixed points, which is average for a randomly chosen involution [19, Page 596]. The linear layer  $M$  is not an involution anymore due to the composition of  $M'$  and shift rows, which is not an involution.

### 3.4 The Key Expansion

The 128-bit key  $(k_0 || k_1)$  is extended to a 192-bit key  $(k_0 || k'_0 || k_1)$  by a linear mapping of the form

$$(k_0 || k_1) \mapsto (k_0 || P(k_0) || k_1).$$

This expansion should be such that it makes peeling of rounds (both at the beginning and at the end) by partial key guessing difficult for the attacker. In particular, we would like that each pair of subkeys among  $k_1$  and the quantities  $(k_0 \oplus k_1)$  and  $(k'_0 \oplus k_1)$  takes all the  $2^{128}$  possible values when  $(k_0 || k_1)$  varies in the set of 128-bit words. In other words, the set of all triples  $(k_0 || P(k_0) || k_1)$  should correspond to an MDS code of length 3 and size  $2^{128}$  over  $\mathbf{F}_2^{64}$ . This equivalently means that both  $x \mapsto P(x)$  and  $x \mapsto x \oplus P(x)$  should be permutations of  $\mathbf{F}_2^{64}$ . Note that no bit-permutation  $P$  satisfies this condition. Indeed, both the all-zero vector and the all-one vector satisfy  $P(x) \oplus x = 0$ .

Thus, a hardware-optimal choice for  $P$  such that both  $P$  and  $P \oplus \text{Id}$  are permutations is

$$P(x) = (x \ggg 1) \oplus (x \ggg 63),$$

*i.e.*,  $P(x_{63}, \dots, x_0) = (x_0, x_{63}, \dots, x_2, x_1 \oplus x_{63})$ . Then, we can easily check that  $P(x) = 0$  (resp.  $P(x) = x$ ) has a unique solution.

## 4 Security Analysis

This section investigates the security of the general construction of PRINCE. In particular, we show that the  $\alpha$ -reflection property of the core cipher does not introduce any generic attack with complexity significantly lower than the known generic attacks against the  $FX$  construction. However, in the particular case of  $\text{PRINCE}_{\text{core}}$ , the  $\alpha$ -reflection property comes from some symmetries in the construction, including the use of an involution as middle round. Thus, we investigate in Section 4.2 whether weaknesses similar to those identified for involutorial ciphers could also appear in the case of PRINCE. An evaluation of the security of PRINCE regarding more classical attacks, including linear, differential and algebraic but also to the recently introduced biclique attacks is provided in the full version [10].

### 4.1 On Generic Attacks: Security Proof

The  $FX$  construction, introduced by Rivest for increasing the resistance of DES to exhaustive key-search [7], consists in deriving a block cipher  $E$  with  $(2n + \kappa)$ -bit key and  $n$ -bit block from a block cipher  $F$  with  $\kappa$ -bit key and  $n$ -bit block by xoring the input and output of  $F$  with a pre-whitening key and a post-whitening key:

$$E_{k_0, k_1, k_2}(x) = F_{k_1}(x \oplus k_0) \oplus k_2.$$

Kilian and Rogaway [25,26] proved that, if the core cipher  $F$  is ideal, then this construction achieves  $(\kappa + n - 1 - \log T)$ -bit security where  $T$  is the number of pairs of inputs and outputs for  $F$  known by the attacker. This result obviously does not apply in the case of PRINCE since the core cipher  $F$  in PRINCE can be easily distinguished from a family of random permutations due to the  $\alpha$ -reflection property, *i.e.*,  $F_k^{-1} = F_{k \oplus \alpha}$  for any  $k$ . Here, we want to quantify the impact of this property on the generic attacks against the  $FX$  construction. For instance, it appears that a decryption oracle also gives a related-key oracle with the fixed-key relation  $(k_0, k_2, k_1) \rightarrow (k_2, k_0, k_1 \oplus \alpha)$  and it is important to determine whether an adversary can profit from this relation.

A similar question was investigated by Kilian and Rogaway for showing that the complementation property of DES decreases the security level by a single bit [25, Section 4]. In the case of the  $\alpha$ -reflection property, we like to model the core cipher  $F$  as an ideal cipher, that is as a set of random permutations, with the (only!) additional relation that  $F_{k \oplus \alpha}(x) = F_k^{-1}(x)$ . Informally, this can be

seen as picking only half of the  $2^\kappa$  permutations independently at random, while the second half is defined by the encryption vs decryption relation above.

More precisely, we consider for  $F$  a keyed permutation with a  $(\kappa - 1)$ -bit key, operating on  $n$ -bit blocks. Let  $\alpha$  be a nonzero element in  $\mathbf{F}_2^\kappa$ . We decompose the set of  $\kappa$ -bit words into two subsets as  $\mathbf{F}_2^\kappa = H \cup (\alpha \oplus H)$  where  $H$  is some linear subspace of dimension  $(\kappa - 1)$  which does not contain  $\alpha$ , *e.g.*, if  $\text{lsb}(\alpha) = 1$ ,  $H$  is the set of all  $n$ -bit words  $x$  with  $\text{lsb}(x) = 0$ . In the following,  $H$  is identified with the set of  $(\kappa - 1)$ -bit words. It is worth noticing that such a decomposition does not exist when  $\alpha = 0$ , *i.e.*, when  $F$  is an involution. Therefore, the following construction is defined for  $\alpha \neq 0$  only. Now, we derive from  $F$  a block cipher with  $(2n + \kappa)$  key bits and  $n$ -bit blocks:

$$E_{k_0, k_1, k_2}(m) = \begin{cases} F_{k_1}(m \oplus k_0) \oplus k_2 & \text{if } k_1 \in H \\ F_{k_1 \oplus \alpha}^{-1}(m \oplus k_0) \oplus k_2 & \text{if } k_1 \in (\alpha \oplus H) \end{cases}$$

This construction, we refer to as  $\tilde{F}X$ -construction, corresponds to the  $FX$  construction applied to  $\tilde{F}$  where  $\tilde{F}$  is the family of  $2^\kappa$  permutations defined by

$$\tilde{F}_k(x) = \begin{cases} F_k(x) & \text{if } k \in H \\ F_{k \oplus \alpha}^{-1}(x) & \text{if } k \in (\alpha \oplus H) \end{cases}$$

The only difference with the construction considered in the case of the complementation property is that  $F$  is extended by using the inverse permutations  $F_k, k \in H$ , instead of the permutations themselves. But, we can obtain a similar result.

More precisely, when analyzing the original  $FX$  construction, Kilian and Rogaway [25] consider the following problem. Let  $\mathcal{A}$  be an adversary with access to three oracles:  $E, F$  and  $F^{-1}$ . During the game, the adversary may make queries to  $E$ , to  $F$  and  $F^{-1}$ . Any query to the  $F/F^{-1}$  oracle consists of a pair  $(k, x)$  in  $\mathbf{F}_2^\kappa \times \mathbf{F}_2^n$  and the oracle returns an element in  $\mathbf{F}_2^n$ . A query to the  $E$  oracle consists of an  $n$ -bit element, and an  $n$ -bit value is returned. The aim of this adversary is then to guess whether the  $E$  oracle computes  $FX_k$  for some random key  $k$ , or if it computes  $\pi$  for a random permutation of  $\mathbf{F}_2^n$ . Then, a game-hopping argument leads to the following upper-bound on the advantage of any such adversary.

**Theorem 1.** [25] *The advantage of any adversary who makes  $D$  queries to the  $E$  oracle and  $T$  queries to the  $F/F^{-1}$  oracle satisfies*

$$\text{Adv}_{FX}^{\text{CPA}}(\mathcal{A}) = \left| \Pr[k \stackrel{\$}{\leftarrow} \mathbf{F}_2^{\kappa+2n}, F \stackrel{\$}{\leftarrow} (\mathcal{P}_n)^{2^\kappa} : \mathcal{A}^{FX_k, F, F^{-1}} = 1] - \Pr[\pi \stackrel{\$}{\leftarrow} \mathcal{P}_n, F \stackrel{\$}{\leftarrow} (\mathcal{P}_n)^{2^\kappa} : \mathcal{A}^{\pi, F, F^{-1}} = 1] \right| \leq DT2^{-(n+\kappa-1)},$$

where  $x \stackrel{\$}{\leftarrow} S$  means that  $x$  is uniformly chosen at random from a set  $S$ ,  $\mathcal{P}_n$  denotes the set of permutations of  $\mathbf{F}_2^n$  and  $F \stackrel{\$}{\leftarrow} (\mathcal{P}_n)^{2^\kappa}$  means that  $F$  is a family of  $2^\kappa$  independently chosen random permutations.

We deduce a similar result for the  $\tilde{F}X$  construction.

**Corollary 1.** *The advantage of any adversary who makes  $D$  queries to the  $E$  oracle and  $T$  queries to the  $F/F^{-1}$  oracle satisfies*

$$\begin{aligned} \text{Adv}_{\tilde{F}X}^{\text{CPA}}(\mathcal{A}) &= \left| \Pr[k \xleftarrow{\$} \mathbf{F}_2^{\kappa+2n}, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{\tilde{F}X_k, F, F^{-1}} = 1] \right. \\ &\quad \left. - \Pr[\pi \xleftarrow{\$} \mathcal{P}_n, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{\pi, F, F^{-1}} = 1] \right| \leq DT2^{-(n+\kappa-2)} \end{aligned}$$

*Proof.* We decompose

$$\begin{aligned} P_c &= \Pr[k \xleftarrow{\$} \mathbf{F}_2^{\kappa+2n}, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{\tilde{F}X_k, F, F^{-1}} = 1] \\ &= \Pr[k_0, k_2 \xleftarrow{\$} \mathbf{F}_2^n, k_1 \xleftarrow{\$} H, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{\tilde{F}X_{k_0, k_1, k_2}, F, F^{-1}} = 1] \\ &\quad \times \Pr[k_1 \in H] \\ &\quad + \Pr[k_0, k_2 \xleftarrow{\$} \mathbf{F}_2^n, k_1 \xleftarrow{\$} \alpha \oplus H, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{\tilde{F}X_{k_0, k_1, k_2}, F, F^{-1}} = 1] \\ &\quad \times \Pr[k_1 \in \alpha \oplus H] \\ &= \frac{1}{2} \Pr[k_0, k_2 \xleftarrow{\$} \mathbf{F}_2^n, k_1 \xleftarrow{\$} H, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{FX_{k_0, k_1, k_2}, F, F^{-1}} = 1] \\ &\quad + \frac{1}{2} \Pr[k_0, k_2 \xleftarrow{\$} \mathbf{F}_2^n, k_1 \xleftarrow{\$} H, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{F^{-1}X_{k_0, k_1, k_2}, F, F^{-1}} = 1], \end{aligned}$$

since

$$\tilde{F}X_{k_0, k_1, k_2}(x) = \begin{cases} FX_{k_0, k_1, k_2}(x) & \text{if } k_1 \in H \\ F^{-1}X_{k_0, k_1 \oplus \alpha, k_2}(x) & \text{if } k_1 \in \alpha \oplus H. \end{cases}$$

Obviously,

$$\Pr[\mathcal{A}^{F^{-1}X_{k_0, k_1, k_2}, F, F^{-1}} = 1] = \Pr[\mathcal{A}^{FX_{k_0, k_1, k_2}, F, F^{-1}} = 1]$$

leading to

$$P_c = \Pr[k_0, k_2 \xleftarrow{\$} \mathbf{F}_2^n, k_1 \xleftarrow{\$} H, F \xleftarrow{\$} (\mathcal{P}_n)^{2^{\kappa-1}} : \mathcal{A}^{FX_{k_0, k_1, k_2}, F, F^{-1}} = 1].$$

It directly follows from Theorem 1 that

$$\text{Adv}_{\tilde{F}X}^{\text{CPA}}(\mathcal{A}) = \text{Adv}_{FX}^{\text{CPA}}(\mathcal{A}) \leq DT2^{-(n+\kappa-2)}.$$

□

As noticed in [25], this bound is still valid in a chosen-ciphertext scenario; it can also be extended to the case where the whitening keys are related, for instance if  $k_2 = k_0$  or  $k_2 = P(k_0)$  as in PRINCE. Both generalizations apply to the  $\tilde{F}X$  construction as well.

The bound obtained for the  $FX$  construction is achieved, for instance by the slide attack due to Biryukov and Wagner [8] and by its recent generalization named slidex [17]. A chosen-plaintext variant of this attack allows to exploit the  $\alpha$ -reflection property for reducing the security level by one bit, compared to the original  $FX$  construction. This attack, detailed in the full version, has an average time complexity corresponding to  $2^{\kappa+n-\log_2 D}$  computations of the core cipher  $F$  for any number  $D$  of pairs of chosen plaintexts-ciphertexts.

## 4.2 Impact of the Construction Implementing the $\alpha$ -Reflection Property

As mentioned earlier, one particular feature of PRINCE is the  $\alpha$ -reflection property of  $\text{PRINCE}_{\text{core}}$ . But, not surprisingly, the construction we used for obtaining this feature also has structural properties, including an involutorial middle round, and care has to be taken when designing a cipher with such a structure. In this section we analyse the influence of this construction on the security of the cipher. In particular, we are interested in the so-called profile of the core cipher, *i.e.*, in the sequence of the lengths of all cycles in the decomposition of  $\text{PRINCE}_{\text{core}}$ .

A first strategy for exploiting some information on the profile of the core cipher is the following. If the decomposition of the core cipher is independent from the key, then this decomposition can be used as a distinguishing property for recovering some information on the whitening keys. The simplest illustration of this type of attack is when the core cipher is an involution, *i.e.* when  $\alpha = 0$  which is the only case where Corollary 1 does not apply. Indeed, the attack presented by Dunkelman *et al.* [17, Section 5.2] allows to recover the sum of the two whitening keys ( $k_0 \oplus k_2$ ) in the  $FX$  construction when  $F$  is an involution. This attack uses the fact that for two plaintext-ciphertext pairs  $(m, c)$  and  $(m', c')$  related by  $m' = E_{k_0, k_1, k_2}^{-1}(m \oplus k_0 \oplus k_2)$  it holds that  $m \oplus c = m' \oplus c'$ . Indeed,

$$\begin{aligned} m' \oplus c' &= E_{k_0, k_1, k_2}^{-1}(m \oplus k_0 \oplus k_2) \oplus m \oplus k_0 \oplus k_2 \\ &= k_0 \oplus F_{k_1}^{-1}(m \oplus k_0) \oplus m \oplus k_0 \oplus k_2 = F_{k_1}(m \oplus k_0) \oplus m \oplus k_2 \\ &= m \oplus c \end{aligned}$$

where the last-but-one equality uses that  $F_{k_1}$  is an involution. Thus, plaintext-ciphertext pairs  $(m, c)$  and  $(m', c')$  such that  $c' = m \oplus k_0 \oplus k_2$  can be easily detected. Such a collision can be found if the attacker has access to  $2^{\frac{n+1}{2}}$  known plaintext-ciphertext pairs, and it provides the value of  $(k_0 \oplus k_2)$ . Moreover, in the particular case of PRINCE,  $k_2$  is related to  $k_0$  by  $k_2 = P(k_0)$  where  $x \mapsto x \oplus P(x)$  is a permutation (see Section 3.4). Therefore, the whitening key  $k_0$  can be deduced from  $(k_0 \oplus k_2)$  in this case. It follows that, when the core cipher is an involution, the whole key can then be recovered with time complexity  $2^\kappa$  (corresponding to an exhaustive search for  $k_1$ ) and data complexity  $2^{\frac{n+1}{2}}$ . This confirms that Corollary 1 does not hold for  $\alpha = 0$ .

This type of attack can be generalized to the case where the profile of the core cipher does not depend on  $k_1$ : since  $\text{PRINCE}_{\text{core}}$  has a reasonable block size, its cycle structure could be precomputed and then used as a distinguishing property for  $(k_0 \oplus k_2)$ . Indeed, the profile of  $E_{k_0, k_1, k_2} : m \mapsto k_2 \oplus F_{k_1}(m \oplus k_0)$  depends on  $(k_0 \oplus k_2)$  only. It follows that, for each  $n$ -bit word  $\delta$ , we could compute one or a few cycles of  $x \mapsto F_{k_1}(x \oplus k_0 \oplus k_2 \oplus \delta)$  in a chosen-plaintext scenario where the attacker knows a sequence of plaintext-ciphertext pairs  $(m_i, c_i)$  with  $m_{i+1} = c_i \oplus \delta$ . A valid candidate for  $(k_0 \oplus k_2)$  is a value  $\delta$  which leads to a cycle having a length which appears in the precomputed profile of  $F_{k_1}$ .

We checked whether the cycle structure of  $\text{PRINCE}_{\text{core}}$  has some peculiarities which do not depend on its key. Based on the technique used by Biryukov for analyzing involutory ciphers [6], we can observe the profile of the reduced version of  $\text{PRINCE}_{\text{core}}$  with 4 Sbox layers where we keep the symmetry around the middle does not depend on the key. Actually, this reduced version can be written as

$$G = (R_5^{-1} \circ \text{Add}_{k_1 \oplus \alpha}) \circ (S^{-1} \circ M' \circ S) \circ (\text{Add}_{k_1} \circ R_5)$$

where  $R_5$  corresponds to  $\mathfrak{R}_5$  without the key addition. Since  $S^{-1} \circ M' \circ S$  is an involution, the cycle structure of  $\text{Add}_{k_1 \oplus \alpha} \circ (S^{-1} \circ M' \circ S) \circ \text{Add}_{k_1}$  depends on  $\alpha$  only and not on  $k_1$ . Its profile then remains unchanged after a right composition with  $R_5$  and a left composition with its inverse. However, this property does not hold anymore when an additional round is included since the next key addition  $\text{Add}_{k_1 \oplus \alpha} \circ G \circ \text{Add}_{k_1}$  modifies the cycle structure of  $G$  in a way which depends on the values  $G$ , and not only on its profile. Therefore, it appears that the previously mentioned attack strategy does not apply if  $\text{PRINCE}_{\text{core}}$  contains more than 6 Sbox layers.

In the light of the previous analysis, a more relevant attack method consists in using the fact that the core cipher may have a peculiar cycle decomposition for some weak keys. For instance, if there exists some weak keys  $k_1$  for which  $\text{PRINCE}_{\text{core}}$  is an involution, then this class of keys can be detected from the knowledge of  $2^{\frac{n+1}{2}}$  pairs of plaintext-ciphertext by counting the number of collisions for  $m \oplus c$ . And the technique from [17] that we have previously described also recovers the whitening key. It is worth noticing that this attack applies to DESX and allows to detect the use of the four weak keys of DES [14] for which DES is an involution. A similar weakness would appear if, in  $\text{PRINCE}_{\text{core}}$ , we have used two subkeys  $k_1$  and  $k'_1$  in turn as round keys. Keeping the remaining structure of  $\text{PRINCE}_{\text{core}}$  results in the following relation

$$F_{(k_1 || k'_1)}^{-1} = F_{(k'_1 \oplus \alpha || k_1 \oplus \alpha)} .$$

However, this has serious – and interesting – consequences for the security of the resulting cipher. For the class of keys such that  $k'_1 = k_1 \oplus \alpha$ , it holds that

$$F_{(k_1 || k'_1)}^{-1} = F_{(k_1 || k'_1)},$$

that is, the core cipher is an involution. This class of weak keys can then be easily detected. It then appears that some particular related-key distinguishers for the core cipher may be exploited for detecting the corresponding class of keys. To be very clear, we do not consider related key-attacks here in the classical sense of enlarging the power of an adversary. But without a careful choice, the construction we used for implementing the  $\alpha$ -reflection property might result in key-recovery attacks for certain weak-key classes, as soon as the core cipher is vulnerable to related key-attacks.

## 5 Implementation

Besides the main target low-latency, low-cost hardware implementation is one of the design objectives of PRINCE. To achieve low-latency, a fully unrolled design should be considered for implementation. During the design process of PRINCE the cost of each function was investigated and each component was carefully designed in order to get the lowest possible gate count without compromising security. One of the most critical and expensive operations of the cipher is the substitution, where we use the same Sbox 16 times (rather than having 16 different Sboxes). Therefore, the implementation of PRINCE started with a search for the most suitable Sbox for the target design specifications. In order to achieve an implementation with low delay and gate count, we analyzed many Sbox instances to identify one with optimal combinational logic and propagation paths. Then, the targeted unrolled design was implemented with the resulting *optimal* Sbox.

In the implementation process, *Cadence NCVerilog 06.20-p001* is used for simulation and *Cadence Encounter RTL Compiler v10.1* for synthesis. Since gate count and delay parameters are heavily technology dependent, the implementations have been synthesized for three different technology libraries: 130 nm and 90 nm low-leakage Faraday libraries from UMC, and 45 nm generic NANGATE Open Cell Library. In all syntheses, typical operating conditions were assumed.

The unrolled version of PRINCE is a direct mapping to hardware of the cipher defined in Section 2. Multiplexers select encryption and decryption keys accordingly. The only costs associated with the key whitening stages are XOR gates and multiplexers used for whitening key selection. However, in practice, due to the unrolled nature of the implementation, these additions reduce to XOR operations with constants, which in turn reduce to inverters or no additional gates at all. Furthermore, these inverters are combined with the preceding or following matrix multiplications, which are implemented with cascaded XOR gates. In cases where an XOR is sourced by the output from an inverter, or is sourcing input of an inverter, it is simply replaced by an XNOR gate and the sourced/sourcing inverter is removed. Since both XOR and XNOR have the same gate count, the overall effect of the round constant addition on area reduces to zero.

The unrolled implementation of PRINCE results are listed in Table 1 for different technologies with respect to different timing constraints. In this table, a unit delay (UD) parameter is used to enable a fair comparison between different technologies. It is the average delay of a single inverter gate (with lowest drive - X1) within a ring oscillator under zero wireload conditions in the target technology (6.7 ps, 31.9 ps, and 43.6 ps for 45 nm, 90 nm, and 130 nm, respectively). We also implemented PRESENT-80, PRESENT-128, LED-128 and AES-128 and applied the same metrics to adequately evaluate the achievements of our new cipher (note that in some cases the key size – and also our security claim – is different: PRINCE does not claim to offer 128-bit security and security against related key-attacks). In order to achieve both encryption and decryption capability in PRESENT and LED, we had to implement both true and inverse Sboxes and select their output by a multiplexer, which doubled the Sbox area with respect to an encryption-only implementation. For AES, we just had to

implement the inverse affine transform since the finite field inversion module could be shared between encryption and decryption. In addition to this comparison, Table 2 shows the extrapolated results (which are calculated by removing register and control logic area from the total gate count, and multiplying the rest by the number of rounds) for other unfolded cipher instances obtained from round-based cipher implementations provided by previous works. Note that all ciphers in the table include encryption and decryption functionality with 128-bit key size, however the comparison is difficult as the block size is different in some cases (also note that the ciphers having 128-bit block size are obviously much bigger and more power consuming than a 64-bit block cipher).

We also measured maximum frequencies achievable by unrolled versions of PRINCE under two different conditions: The frequency where the area of synthesized design starts to deviate from the unconstrained area – 158.9, 38.4 and 35.5 MHz, and the frequency where the timing slack becomes zero – 212.8, 71.8 and 54.3 MHz. Both figures are given for 45 nm, 90 nm, and 130 nm, respectively.

**Table 1.** Area/power comparison of unrolled versions of PRINCE and other ciphers

	Tech.	Nangate 45nm Generic			UMC 90nm Faraday			UMC 130nm Faraday		
	Constr.(UD)	1000	3162	10000	1000	3162	10000	1000	3162	10000
PRINCE <sup>*</sup>	Area(GE)	8260	8263	8263	7996	7996	7996	8679	8679	8679
	Power(mW)	38.5	17.9	8.3	26.3	10.9	3.9	29.8	11.8	4.1
PRESENT-80	Area(GE)	63942	51631	50429	113062	49723	49698	119196	51790	51790
	Power(mW)	1304.6	320.9	98.0	1436.9	144.9	45.5	1578.4	134.9	42.7
PRESENT-128	Area(GE)	68908	56668	55467	120271	54576	54525	126351	56732	56722
	Power(mW)	1327.1	330.4	99.1	1491.1	149.9	47.8	1638.7	137.4	43.6
LED-128	Area(GE)	109811	109958	109697	281240	286779	98100	236770	235106	111496
	Power(mW)	2470.7	835.7	252.3	5405.0	1076.3	133.7	5274.8	1133.9	163.6
AES-128	Area (GE)	135051	135093	118440	421997	130835	118522	347860	141060	130764
	Power (mW)	3265.8	1165.7	301.6	8903.2	587.4	186.8	8911.2	876.8	229.1

**Table 2.** Extrapolated area of unrolled versions of other ciphers against PRINCE

Technology	Area* (GE)
CLEFIA-128 [1]	28035 (18 rounds unfolded, 130nm CMOS)
HIGHT-128 [23]	42688 (32 rounds unfolded, 250nm CMOS)
mCrypton-128 [29]	37635 (13 rounds unfolded, 130nm CMOS)
Piccolo-128 [32]	25668 (31 rounds unfolded, 130nm CMOS)

\* Area requirements extrapolated from round-based implementations.

## References

1. Akishita, T., Hiwatari, H.: Very Compact Hardware Implementations of the Block Cipher CLEFIA. In: Miri, A., Vaudenay, S. (eds.) SAC 2011. LNCS, vol. 7118, pp. 278–292. Springer, Heidelberg (2012)
2. Babbage, S., Dodd, M.: The MICKEY Stream Ciphers. In: Robshaw, Billet [31], pp. 191–209



3. Barreto, P.S.L.M., Rijmen, V.: The ANUBIS Block Cipher. Submission to the NESSIE project (2000), <http://www.larc.usp.br/~pbarreto/AnubisPage.html>
4. Barreto, P.S.L.M., Rijmen, V.: The Khazad Legacy-level Block Cipher. Submission to the NESSIE project (2000), <http://www.larc.usp.br/~pbarreto/KhazadPage.html>
5. Bernstein, D.J.: The Salsa20 Family of Stream Ciphers. In: Robshaw, Billet [31], pp. 84–97
6. Biryukov, A.: Analysis of Involutional Ciphers: Khazad and Anubis. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 45–53. Springer, Heidelberg (2003)
7. Biryukov, A.: DES-X (or DESX). In: Encyclopedia of Cryptography and Security, 2nd edn., p. 331. Springer (2011)
8. Biryukov, A., Wagner, D.: Advanced Slide Attacks. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 589–606. Springer, Heidelberg (2000)
9. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M., Seurin, Y., Vikkelsøe, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
10. Borghoff, J., Canteaut, A., Güneşyü, T., Kavun, E.B., Knežević, M., Knudsen, L.R., Leander, G., Nikov, V., Paar, C., Rechberger, C., Rombouts, P., Thomsen, S.S., Yalçın, T.: PRINCE – A Low-latency Block Cipher for Pervasive Computing Applications. IACR Cryptology ePrint Archive, 529 (2012)
11. Daemen, J.: Cipher and Hash Function Design, Strategies Based on Linear and Differential Cryptanalysis. PhD thesis, Katholieke Universiteit Leuven (1995)
12. Daemen, J., Peeters, M., Van Assche, G., Rijmen, V.: The NOEKEON Block Cipher. Submission to the NESSIE project (2000), <http://gro.noekeon.org/>
13. Daemen, J., Rijmen, V.: Codes and Provable Security of Ciphers. In: Enhancing Cryptographic Primitives with Techniques from Error Correcting Codes. NATO Science for Peace and Security Series D - Information and Communication Security 23, vol. 1807, pp. 60–80. IOS Press (2009)
14. Davies, D.W.: Some Regular Properties of the ‘Data Encryption Standard’ Algorithm. In: Advances in Cryptology, CRYPTO 1982, pp. 89–96. Plenum Press, New York (1982)
15. De Cannière, C., Dunkelman, O., Knežević, M.: KATAN and KTANTAN — A Family of Small and Efficient Hardware-Oriented Block Ciphers. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 272–288. Springer, Heidelberg (2009)
16. De Cannière, C., Preneel, B.: Trivium Specifications. eSTREAM, ECRYPT Stream Cipher Project (2006)
17. Dunkelman, O., Keller, N., Shamir, A.: Minimalism in Cryptography: The Even-Mansour Scheme Revisited. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 336–354. Springer, Heidelberg (2012)
18. Even, S., Mansour, Y.: A Construction of a Cipher From a Single Pseudorandom Permutation. In: Matsumoto, T., Imai, H., Rivest, R.L. (eds.) ASIACRYPT 1991. LNCS, vol. 739, pp. 210–224. Springer, Heidelberg (1993)
19. Flajolet, P., Sedgewick, R.: Analytic Combinatorics. Cambridge University Press (2009)
20. Gong, Z., Nikova, S., Law, Y.W.: KLEIN: A New Family of Lightweight Block Ciphers. In: Juels, A., Paar, C. (eds.) RFIDSec 2011. LNCS, vol. 7055, pp. 1–18. Springer, Heidelberg (2012)

21. Guo, J., Peyrin, T., Poschmann, A., Robshaw, M.: The LED Block Cipher. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 326–341. Springer, Heidelberg (2011)
22. Hell, M., Johansson, T., Meier, W.: Grain: A Stream Cipher for Constrained Environments. *International Journal of Wireless and Mobile Computing* 2(1), 86–93 (2007)
23. Hong, D., Sung, J., Hong, S., Lim, J., Lee, S., Koo, B.-S., Lee, C., Chang, D., Lee, J., Jeong, K., Kim, H., Kim, J.-S., Chee, S.: HIGHT: A New Block Cipher Suitable for Low-Resource Device. In: Goubin, L., Matsui, M. (eds.) CHES 2006. LNCS, vol. 4249, pp. 46–59. Springer, Heidelberg (2006)
24. Kerckhof, S., Durvaux, F., Hocquet, C., Bol, D., Standaert, F.-X.: Towards Green Cryptography: A Comparison of Lightweight Ciphers from the Energy Viewpoint. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 390–407. Springer, Heidelberg (2012)
25. Kilian, J., Rogaway, P.: How to Protect DES against Exhaustive Key Search. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 252–267. Springer, Heidelberg (1996)
26. Kilian, J., Rogaway, P.: How to Protect DES Against Exhaustive Key Search (An Analysis of DESX). *J. Cryptology* 14(1), 17–35 (2001)
27. Knežević, M., Nikov, V., Rombouts, P.: Low-Latency Encryption – Is “Lightweight = Light + Wait”? In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 426–446. Springer, Heidelberg (2012)
28. Leander, G., Poschmann, A.: On the Classification of 4 Bit S-Boxes. In: Carlet, C., Sunar, B. (eds.) WAIFI 2007. LNCS, vol. 4547, pp. 159–176. Springer, Heidelberg (2007)
29. Lim, C.H., Korkishko, T.: mCrypton – A Lightweight Block Cipher for Security of Low-Cost RFID Tags and Sensors. In: Song, J.-S., Kwon, T., Yung, M. (eds.) WISA 2005. LNCS, vol. 3786, pp. 243–258. Springer, Heidelberg (2006)
30. NANGATE. The NanGate 45nm Opencell Library, <http://www.nangate.com>
31. Robshaw, M., Billet, O. (eds.): New Stream Cipher Designs. LNCS, vol. 4986. Springer, Heidelberg (2008)
32. Shibutani, K., Isobe, T., Hiwatari, H., Mitsuda, A., Akishita, T., Shirai, T.: *Piccolo*: An Ultra-Lightweight Blockcipher. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 342–357. Springer, Heidelberg (2011)
33. Shirai, T., Shibutani, K., Akishita, T., Moriai, S., Iwata, T.: The 128-Bit Blockcipher CLEFIA (Extended Abstract). In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 181–195. Springer, Heidelberg (2007)
34. Standaert, F.-X., Piret, G., Gershenfeld, N., Quisquater, J.-J.: SEA: A Scalable Encryption Algorithm for Small Embedded Applications. In: Domingo-Ferrer, J., Posegga, J., Schreckling, D. (eds.) CARDIS 2006. LNCS, vol. 3928, pp. 222–236. Springer, Heidelberg (2006)
35. Standaert, F.-X., Piret, G., Rouvroy, G., Quisquater, J.-J., Legat, J.-D.: ICEBERG: An Involutive Cipher Efficient for Block Encryption in Reconfigurable Hardware. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 279–299. Springer, Heidelberg (2004)