

Block Ciphers – Focus on the Linear Layer (feat. PRIDE)*

Martin R. Albrecht^{1,**}, Benedikt Driessen^{2,***}, Elif Bilge Kavun^{3,†},
Gregor Leander^{3,‡}, Christof Paar³, and Tolga Yalçın^{4,***}

¹ Information Security Group, Royal Holloway, University of London, UK

² Infineon AG, Neubiberg, Germany

³ Horst Görtz Institute for IT Security, Ruhr-Universität Bochum, Germany

⁴ University of Information Science and Technology, Ohrid, Macedonia

Abstract. The linear layer is a core component in any substitution-permutation network block cipher. Its design significantly influences both the security and the efficiency of the resulting block cipher. Surprisingly, not many general constructions are known that allow to choose trade-offs between security and efficiency. Especially, when compared to Sboxes, it seems that the linear layer is crucially understudied. In this paper, we propose a general methodology to construct good, sometimes optimal, linear layers allowing for a large variety of trade-offs. We give several instances of our construction and on top underline its value by presenting a new block cipher. PRIDE is optimized for 8-bit micro-controllers and significantly outperforms all academic solutions both in terms of code size and cycle count.

Keywords: block cipher, linear layer, wide-trail, embedded processors.

1 Introduction

Block ciphers are one of the most prominently used cryptographic primitives and probably account for the largest portion of data encrypted today. This was facilitated by the introduction of Rijndael as the Advanced Encryption Standard (AES) [1], which was a major step forward in the field of block cipher design. Not only does AES offer strong security, but its structure also inspired many cipher designs ever since. One of the merits of AES (and its predecessor SQUARE [20]) was demonstrating that a well-chosen linear layer is not only crucial for the

* Due to page limitations, several details are omitted in this proceedings version. A full version is available at [2].

** Most of this work was done while the author was at the Technical University of Denmark

*** Most of this work was done while the authors were at Ruhr-Universität Bochum.

† The research was supported in part by the DFG Research Training Group GRK 1817/1.

‡ The research was supported in part by the BMBF Project UNIKOPS (01BY1040).

security (and efficiency) of a block cipher, but also allows to argue in a simple and thereby convincing way about its security.

There are two main design strategies that can be identified for block ciphers: Sbox-based constructions and constructions without Sboxes, most prominently those using addition, rotation, and XORs (ARX designs). Furthermore, Sbox-based designs can be split into *Feistel-ciphers and substitution-permutation networks (SPN)*. Both concepts have been successfully used in practice, the most prominent example of an SPN cipher being AES and the most prominent Feistel-cipher being the former Data Encryption Standard (DES) [22].

It is also worth mentioning that the concept of SPN has not only been used in the design of block ciphers but also for designing cryptographic permutations, most prominently for the design of several sponge-based hash functions including SHA-3 [11]. In SP networks, the round function consists of a non-linear layer composed of small Sboxes working in parallel on small chunks of the state and a linear layer that mixes those chunks. Thus, designing an SPN block cipher essentially reduces to choosing one (or several) Sboxes and a linear layer.

A lot of research has been devoted to the study of Sboxes. All Sboxes of size up to 4 bits have been classified (indeed, more than once – cf. [14,36,46]). Moreover, Sboxes with optimal resistance against differential and linear attacks have been classified up to dimension 5 [17]. In general, several constructions are known for good and optimal Sboxes in arbitrary dimensions. Starting with the work of Nyberg [43], this has evolved into its own field of research in which those functions are studied in great detail. A nice survey of the main results of this line of study is provided by Carlet [18].

The situation for the other main design part, the linear layer, is less clear.

1.1 The Linear Layer

For the design of the linear layer, two general approaches can be identified. One still widely-used method is to design the linear layer in a rather ad-hoc fashion, without following general design guidelines. While this might lead to very secure and efficient algorithms (cf. Serpent [3] and SHA-3 as prominent examples), it is not very satisfactory from a scientific point-of-view. The second general design strategy is the wide-trail strategy introduced by Daemen in [19] (see also [21]). Especially for the security against linear [41] and differential [12] attacks, the wide-trail strategy usually results in simple and strong security arguments. It is therefore not surprising that this concept has found its way in many recent designs (e.g. Khazad [9], Anubis [8], Grøstl [25], PHOTON [29], LED [30], PRINCE [16], mCrypton [39] to name but a few). In a nutshell, the main idea of the wide-trail strategy is to link the number of active Sboxes for linear and differential cryptanalysis to the minimal distance of a certain linear code associated with the linear layer of the cipher. In turn, choosing a good code (with some additional constraints) results in a large number of active Sboxes.

While the wide-trail strategy does provide a powerful tool for arguing about the security of a cipher, it does not help in actually designing an efficient linear layer (or the corresponding linear code) with a suitable number of active Sboxes.

Here, with the exception of early designs in [19] and later PRINCE and mCrypton, most ciphers following the wide-trail strategy simply choose an MDS matrix as the core component. This might guarantee an (partially) optimal number of active Sboxes, but usually comes at the price of a less efficient implementation. The only exception here is that, in the case of MDS matrices, the authors of PHOTON and LED made the observation that implementing such matrices in a serialized fashion improves hardware-efficiency. This idea was further generalized in [47,53], and more recently in [5].

It is our belief that, in many cases, it is advantageous to use a near-MDS matrix (or in general a matrix with a sub-optimal branch number) for the overall design. Furthermore, it is, in our opinion, utmost surprising that there are virtually no general constructions or guidelines that would allow an SPN design to benefit from security vs. efficiency trade-offs. This is in particular important when it comes to ciphers where specific performance characteristics are crucial, e.g. in lightweight cryptography.

1.2 The Current State of Lightweight Cryptography

In recent years, the field of lightweight cryptography has attracted a lot of attention from the cryptographic community. In particular, designing lightweight block ciphers has been a very active field for several years now. The dominant metric according to which the vast majority of lightweight ciphers have been optimized was and still is the chip area. While this is certainly a valid optimization objective, its relevance to real-world applications is limited. Nowadays, there are several interesting and strong proposals available that feature a very small area but simultaneously neglect other, important real-world constraints. Moreover, recent proposals achieve the goal of a small chip area by sacrificing execution speed to such an extent that even in applications where speed is supposedly uncritical, the ciphers are getting too slow¹.

Note that software solutions, i.e. low-end embedded processors, actually dominate the world of embedded systems and dedicated hardware is a comparably small fraction. Considering this fact, it is quite puzzling that efficiency on low-cost processors was disregarded for so long. Certainly, there were a few exceptions: Several theoretical and practical studies have already been done in this field. Practical examples include several proposals for instruction set extensions [38,42,48,37]. Among these, the Intel AES instruction set [31] is the most well-known and practically relevant one. There have also been attempts to come up with ciphers that are (partially) tailored for low-cost processors [51,50,54,26,10,32]. Of these, execution times of both SEA and ITUbee are rather high, mostly due to the high number of rounds. Furthermore, ITUbee uses 8-bit Sboxes, which occupy a vast amount of program memory storage. SPECK, on the other hand, seems to be an excellent *lightweight software cipher* in terms of both speed and program memory.

¹ See also [35] asking “Is lightweight = light + wait?”.

It is obvious that there are quite some challenges to be overcome in this relatively untouched area of lightweight software cryptography. The software cipher for embedded devices of the future should not only be compact in terms of program memory, but also be relatively fast in execution time. It should clearly be secure and, preferably, its security should be easily analysed and verified. The latter can possibly be achieved by building on conservative structures, which are conventionally costly in software implementation, thereby posing even harder challenges.

One major component influencing all or at least most of those criteria outlined above is the linear layer. Thus, it is important to have general constructions for linear layers that allow to explore and make optimal use of the possible trade-offs.

1.3 Our Contribution

In this paper, we take steps towards a better understanding of possible trade-offs for linear layers. After introducing necessary concepts and notation in Section 2, we give a general construction that allows to combine several strong linear mappings on a few number of bits into a strong linear layer for a larger number of bits (cf. Section 3). From a coding theory perspective, this construction corresponds to a construction known as block-interleaving (see [40], pages 131-132). While this idea is rather simple, its applicability is powerful. Implicitly, a specific instance of our construction is already implemented in AES. Furthermore, special instances of this construction are recently used in [7] and [28].

We illustrate our approach by providing several linear layers with an optimal or almost optimal trade-off between hardware-efficiency and number of active Sboxes in Section 4. Along similar lines, we present a classification of all linear layers fulfilling the criteria of the block cipher PRINCE in [2], Appendix C. Those examples show in particular that the construction given in Section 3 allows the construction of non-equivalent codes even when starting from equivalent ones. Secondly, we show that our construction also leads to very strong linear layers with respect to efficiency on embedded 8-bit micro-controllers. For this, we adopt a search strategy from [52] to find the *most efficient linear layer possible* within our constraints. We implemented this search on an FPGA platform to overcome the big computational effort involved and to have the advantage of reconfigurability. Details are described in Section 5.1.

With this, and as a second main contribution of our paper, we make use of our construction to design a new block cipher named PRIDE that significantly outperforms all existing block ciphers of similar key-sizes, with the exception of SIMON and SPECK [10]. One of the key-points here is that our construction of strong linear layers is nicely in line with a bit-sliced implementation of the Sbox layer. Our cipher is comparable, both in speed and memory size, to the new NSA block ciphers SIMON and SPECK, dedicated for the same platform. We conclude the paper in Section 6 with some open problems and pressing topics for further investigation. Finally, we note that while in this paper we focus on SPN ciphers, most of the results translate to the design of Feistel ciphers as well.

2 Notation and Preliminaries

In this section, we fix the basic notation and furthermore recall the ideas of the wide-trail strategy.

We deal with SPN block ciphers where the Sbox layer consist of n Sboxes of size b each. Thus the block size of the cipher is $n \times b$. The linear layer will be implemented by applying k binary matrices in parallel.

We denote by \mathbb{F}_2 the field with two elements and by \mathbb{F}_2^n the n -dimensional vector space over \mathbb{F}_2 . Note that any finite extension field \mathbb{F}_{2^b} over \mathbb{F}_2 can be viewed as the vector space \mathbb{F}_2^b of dimension b . Along these lines, the vector space $(\mathbb{F}_{2^b})^n$ can be viewed as the (nested) vector space $(\mathbb{F}_2^b)^n$.

Given a vector $x = (x_1, \dots, x_n) \in (\mathbb{F}_2^b)^n$ where each $x_i \in \mathbb{F}_2^b$ we define its weight² as

$$\text{wt}_b(x) = |\{1 \leq i \leq n \mid x_i \neq 0\}|.$$

Following [21], given a linear mapping $L : (\mathbb{F}_2^b)^n \rightarrow (\mathbb{F}_2^b)^n$ its differential branch number is defined as

$$\mathcal{B}_d(L) := \min\{\text{wt}_b(x) + \text{wt}_b(L(x)) \mid x \in (\mathbb{F}_2^b)^n, x \neq 0\}.$$

The cryptographic significance of the branch number is that the branch number corresponds to the minimal number of active Sboxes in any two consecutive rounds. Here an Sbox is called active if it gets a non-zero input difference in its input.

Given an upper bound p on the differential probability for a single Sbox along with a lower bound of active Sboxes immediately allows to deduce an upper bound for any differential characteristic³ using

$$\text{average probability for any non-trivial characteristic} \leq p^{\#\text{active Sboxes}}.$$

For linear cryptanalysis, the linear branch number is defined as

$$\mathcal{B}_l(L) := \min\{\text{wt}_b(x) + \text{wt}_b(L^*(x)) \mid x \in (\mathbb{F}_2^b)^n, x \neq 0\}$$

where L^* is the adjoint linear mapping. That is, with respect to the standard inner product, L^* corresponds to the transposed matrix of L .

In terms of correlation (cf., for example, [19]), an upper bound c on the absolute value of the correlation for a single Sbox results in a bound for any linear trail (or linear characteristic, linear path) via

$$\text{absolute correlation for a trail} \leq c^{\#\text{active Sboxes}}.$$

The differential branch number corresponds to the minimal distance of the \mathbb{F}_2 -linear code C over \mathbb{F}_2^b with generator matrix

$$G = [I \mid L^T]$$

² Of course $(\mathbb{F}_2^b)^n$ is isomorphic to \mathbb{F}_2^{nb} , but the weight is defined differently on each.

³ Averaging over all keys, assuming independent round keys.

where I is the $n \times n$ identity matrix. The length of the code is $2n$ and its dimension is n (here dimension corresponds to $\log_{2^b}(|C|)$ as it is not necessarily a linear code). Thus, C is a $(2n, 2^n)$ additive code over \mathbb{F}_2^b with minimal distance $d = \mathcal{B}_d(L)$.

The linear branch number corresponds in the same way to the minimal distance of the \mathbb{F}_2 -linear code C^\perp with generator matrix

$$G^* = [L \mid I].$$

Note that C^\perp is the dual code of C and in general the minimal distances of C^\perp and C do not need to be identical.

Finally, given linear maps L_1 and L_2 , we denote by $L_1 \times L_2$ the direct sum of the mappings, i.e.

$$(L_1 \times L_2)(x, y) := (L_1(x), L_2(y)).$$

3 The Interleaving Construction

Following the wide-trail strategy, we construct linear layers by constructing a $(2n, 2^n)$ additive codes with minimal distance d over \mathbb{F}_2^b . The code needs to have a generator matrix G in standard form, i.e.

$$G = [I \mid L^T]$$

where the submatrix L is invertible, and corresponds to the linear layer we are using.

Hence, the main question is how to construct “efficient” matrices L with a given branch number. Our construction allows to combine small matrices into bigger ones. We hereby drastically reduce the search-space of possible linear layers. This in turn makes it possible to construct efficient linear layers for various trade-offs, as demonstrated in the following sections.

As mentioned above, the construction described in [21] can be seen as a special case of our construction. The main difference (except the generalization) is that we shift the focus of the construction in [21] from the 4 round super-box view to a 2 round-view. While Daemen and Rijmen focused on the bounds for 4 rounds, we make use of their ideas to actually construct linear layers. Moreover, a particular instance of the general construction we elaborate on here, was already used in the linear layer of the hash function Whirlwind [7]. There, several small MDS matrices are used to construct a larger one.

We give a simple illustrative example of our approach in [2], Appendix A.

3.1 The General Construction

We are now ready to give a formal description of our approach. First define the following isomorphism

$$P_{b_1, \dots, b_k}^n : \left(\mathbb{F}_2^{b_1} \times \mathbb{F}_2^{b_2} \times \dots \times \mathbb{F}_2^{b_k} \right)^n \rightarrow \left(\mathbb{F}_2^{b_1} \right)^n \times \left(\mathbb{F}_2^{b_2} \right)^n \times \dots \times \left(\mathbb{F}_2^{b_k} \right)^n$$

$$(x_1, \dots, x_n) \mapsto \left((x_1^{(1)}, \dots, x_n^{(1)}), \dots, (x_1^{(k)}, \dots, x_n^{(k)}) \right)$$

$$\text{where } x_i = \left(x_i^{(1)}, \dots, x_i^{(k)} \right) \text{ with } x_i^{(j)} \in \mathbb{F}_2^{b_j}.$$

This isomorphism performs the transformation of mapping Sbox outputs to our small linear layers L_i . For example, in Appendix A of [2], we considered individual bits (i.e. $b_1, \dots, b_k = 1$) from 4 (i.e., $k = 4$) 4-bit Sboxes (i.e $n = 4$).

Note that, for our purpose, there are in fact many possible choices for P . In particular, we may permute the entries within $(\mathbb{F}_2^{b_i})^n$. Given this isomorphism we can now state our main theorem. The construction of P follows the idea of a diffusion-optimal mapping as defined in [21, Definition 5].

Theorem 1. *Let $G_i = [I \mid L_i^T]$ be the generator matrix for an \mathbb{F}_2 -linear $(2n, 2^n)$ code with minimal distance d_i over $\mathbb{F}_2^{b_i}$ for $0 \leq i < k$. Then the matrix $G = [I \mid L^T]$ with*

$$L = (P_{b_1, \dots, b_k}^n)^{-1} \circ (L_0 \times L_1 \times \dots \times L_{k-1}) \circ P_{b_1, \dots, b_k}^n$$

is the generator matrix of an \mathbb{F}_2 -linear $(2n, 2^n)$ code with minimal distance d over \mathbb{F}_2^b where

$$d = \min_i d_i \quad \text{and} \quad b = \sum_i b_i.$$

Proof. Since P_{b_1, \dots, b_k}^n and $(P_{b_1, \dots, b_k}^n)^{-1}$ are permutation matrices, by construction L has full rank. To see that $\text{wt}_b(w) + \text{wt}_b(v) \geq \min_i d_i$ for any $v \in \mathbb{F}_2^b \setminus \{0\}$ and $w = L \cdot v$, observe that $\text{wt}_b(w) + \text{wt}_b(v)$ is minimal when all entries in v are zero except those mapped to the positions acted on by L_j where L_j is the matrix with the minimal branch number. \square

Remark 1. The interleaving construction allows to construct non-equivalent codes even when starting with equivalent L_i 's. This is shown in a particular case in Appendix C of [2], where different choices of (equivalent) L_i 's lead to different numbers of minimum-weight codewords.

A special case of the construction above is implicitly already used in AES. In the case of AES, it is used to construct a $[8, 4, 5]$ code over \mathbb{F}_2^{32} from 4 copies of the $[8, 4, 5]$ code over \mathbb{F}_2^8 given by the *MixColumn* operation. In the Superbox view on AES, the *ShiftRows* operation plays the role of the mapping P (and its inverse) and *MixColumns* corresponds to the mappings L_i .⁴

In the following, we use this construction to design efficient linear layers. Besides the differential and linear branch number, we hereby focus mainly on three criteria:

- Maximize the diffusion (cf. Section 3.3)
- Minimize the density of the matrix (cf. Section 4)
- Software-efficiency (cf. Section 5)

⁴ Note that the cipher PRINCE implicitly uses the construction twice. Once for generating the matrix M as in Appendix A of [2] and second for the improved bound on 4 rounds, just like in AES.

The strategy we employ is as follows. We first find candidates for L_0 , i.e., $(2n, 2^n)$ additive codes with minimal distance d_0 over $\mathbb{F}_{2^{b_0}}$. In this stage, we ensure that the branch number is d_0 and our efficiency constraints are satisfied. We then apply permutations to L_0 to produce L_i for $i > 0$. This stage maximizes diffusion.

3.2 Searching for L_0

The following lemma (which is a rather straightforward generalization of Theorem 4 in [53]) gives a necessary and sufficient condition that a given matrix L has branch number d over \mathbb{F}_2^b .

Lemma 1. *Let L be a $bn \times bn$ binary matrix, decomposed into $b \times b$ submatrices $L_{i,j}$.*

$$L = \begin{pmatrix} L_{0,0} & L_{0,1} & \cdots & L_{0,n-1} \\ L_{1,0} & L_{1,1} & \cdots & L_{1,n-1} \\ \vdots & \vdots & \ddots & \vdots \\ L_{n-1,0} & L_{n-1,1} & \cdots & L_{n-1,n-1} \end{pmatrix} \quad (1)$$

Then, L has differential branch number d over \mathbb{F}_2^b if and only if all $i \times (n-d+i+1)$ block submatrices of L have full rank for $1 \leq i < d-1$. Moreover, L has linear branch number d if and only if all $(n-d+i+1) \times i$ block submatrices of L have full rank for $1 \leq i < d-1$.

Based on Lemma 1 we may instantiate various search algorithms which we will describe in Section 4 and Section 5. In our search we focus on cyclic matrices, i.e. matrices where row $i > 0$ is constructed by cyclic shifting row 0 by i indices. These matrices have the advantage of being efficient both in software and hardware. Furthermore, since these matrices are symmetric, considering the dual code C^\perp to $C = [I \mid L^T]$ is straightforward.

3.3 Ensuring High Dependency

In this section, we assume we are given a matrix L_0 and wish to construct L_1, \dots, L_{k-1} that maximize the diffusion of the map $L = \left(P_{b_1, \dots, b_k}^n\right)^{-1} \circ (L_0 \times L_1 \times \cdots \times L_{k-1}) \circ P_{b_1, \dots, b_k}^n$.

Given an $bn \times bn$ binary matrix L decomposed as in Eq. (1), we define its support as the $n \times n$ binary matrix $\text{Supp}(L)$ where

$$\text{Supp}(L)_{i,j} = \begin{cases} 1 & \text{if } L_{i,j} \neq 0 \\ 0 & \text{else} \end{cases}$$

Now assume that $\text{Supp}(L_0)$ has a zero entry at index i', j' . If we apply the same L_i in all k positions this means that the outputs from the i' 'th Sbox have no impact on the inputs of the j' 'th Sbox after the linear layer. In other words, a

linear-layer following the construction of Theorem 1 ensure full dependency if and only if

$$\left(\bigvee_{0 \leq i < k} \text{Supp}(L_i) \right)_{i', j'} = 1 \quad \forall 0 \leq i', j' < n.$$

Hence, we want to apply different matrices L_i in each of the k positions, such that in at least one $\text{Supp}(L_i)$ has a non-zero entry at index i', j' for all $0 \leq i', j' < n$. In order to construct matrices L_i for $i > 0$ from a matrix L_0 we may apply block-permutation matrices from the left and right to L_0 as these clearly neither impact the density nor the branch number. Hence, we focus on finding permutation matrices P_i, Q_i such that the density of $\bigvee_{0 \leq i < b} \text{Supp}(P_i \cdot L_0 \cdot Q_i)$ is maximized. In Appendix F of [2], we give two strategies for finding such P_i, Q_i , one is heuristic but computationally cheap, the other is guaranteed to return an optimal solution – based on Constraint Integer Programming – but can be computationally intensive.

We note that the difficulty of the problem depends on the size of the Sbox and the density of L_i . As MDS matrices always have density 1, the problem of full dependency does not occur when combining such matrices. Finally, if the construction ensures full dependency for a given k , it is always possible to achieve full dependency for any $k' \geq k$.

In contrast with the branch number, if a linear layer ensures high dependency, its inverse does not necessarily achieve the same dependency. Thus, it is in general necessary to check the dependency of the inverse separately.

4 Optimizing for Hardware

In this section, we give examples of $[2n, n, d]$ codes over \mathbb{F}_2^b and give algorithms for finding such instances. First, the following lemma gives a lower bound on the density of a matrix with branch number d . Our aim here is to find linear layers that are efficiently implementable in hardware. More precisely, we aim for an implementation in one cycle. PHOTON and LED demonstrated that there is a trade-off between clock cycles and number of gate equivalence for the linear layer. The trade-off we consider here is, complementary to PHOTON and LED, between efficient implementation in one clock cycle and the (linear and differential) branch number. Note that in our setting, the cost of implementation is directly connected to the number of ones in the matrix representation of the linear layer.

Lemma 2. *Let matrix $G = [I \mid L^T]$ be the generator matrix for an \mathbb{F}_2 -linear $(2n, 2^n)$ code with minimal distance d such that the dual code has minimum distance d as well. Then L has at least $d - 1$ ones per row and per column.*

Proof. Computing $w = L \cdot v$ where v is a vector with one non-zero entry 1, we have that w must be a vector with $d-1$ non-zero entries if the minimum distance of $[I \mid L^T]$ is d . Hence, there must be at least $d-1$ ones per row. Applying the same argument to $w = L^T \cdot v = v \cdot L$ shows that at least $d-1$ entries per column must be non-zero. \square

The main merit of the above lemma is that it allows to determine the optimal solutions in terms of efficiency. This is in contrast to the case for software implementation, where the optimal solution is unknown.

Lemmas 1 and 2 give rise to various search strategies for finding $(2n, 2^n)$ additive codes with minimal distance d over \mathbb{F}_2^b . We discuss those strategies in Appendix B of [2] and present results of those strategies next.

4.1 Hardware-Optimal Examples

Below we give some examples for our construction. We hereby focus on $[2n, n, d]$ codes over \mathbb{F}_2 , i.e. we use $b_i = 1$.⁵ Note that this naturally limits the achievable branch number. For binary linear codes the optimal minimal distance is known for small length (cf. [27] for more information). We give a small abridgement of the known bounds on the minimal distance for linear $[2n, n]$ codes over $\mathbb{F}_2, \mathbb{F}_4$, and \mathbb{F}_8 in Appendix E of [2]. As can be seen in this table, in order to achieve a high branch number, it might be necessary to consider linear codes over \mathbb{F}_{2^m} , or (more general) additive codes over \mathbb{F}_2^m for some small $m > 1$.

The examples in Figure 1 are optimal in the sense that they achieve the best possible branch number (both linear and differential) for the given length (with the exception of $n = 11, 13$, and 14) with the least possible number of ones in the matrix (cf. Lemma 2). The number D corresponds to the average number of ones per row/column and D_{inv} to the average number of ones per row/column of the inverse matrix. The only candidate which does not satisfy $D = d - 1$ is $n = 8$. This candidate was found using the approach from Appendix B.3 of [2], which guarantees to return the optimal solution. Hence, we conclude that $4\frac{1}{8}$ is indeed the lowest density possible. That is, there is no 8×8 binary matrix with branch number 5 with only 32 ones, but the best we can do is 33 ones.

For each example we list the dimension (i.e. the number of Sboxes), the achieved branch number and the minimal k such that it is possible to achieve full dependency with two Sbox layers interleaved with one linear layer. These values were found using the CIP approach in Section 3.3. Note that in this case (i.e. $b_i = 1$) the value k actually corresponds to the minimal Sbox size that allows full dependency. Finally, k_{inv} is the minimum Sbox size to achieve full diffusion for the inverse matrix. Note that for all these examples, the corresponding code is actually equivalent to its dual. In particular this implies that the linear and differential branch number are equal.

⁵ We refer to Appendix C of [2] for an exemplary comparison of the set of linear layers constructed by Theorem 1 and the entire space with the same criteria for $[8, 4, 4]$ codes over \mathbb{F}_2^4 .

n	$\max(d)$	d	D	D_{inv}	k	k_{inv}	Technique	Matrix
2	2	2	1	1	2	2	[2], App. B.1	cyclic shift (10) to the left.
3	2	2	1	1	3	3	[2], App. B.1	cyclic shift (100) to the left.
4	4	4	3	3	2	2	[2], App. B.1	cyclic shift (1110) to the left.
5	4	4	3	3	2	2	[2], App. B.1	cyclic shift (11100) to the left.
6	4	4	3	3	2	2	[2], App. B.1	cyclic shift (110100) to the left.
7	4	4	3	$4\frac{3}{7}$	3	2	[2], App. B.3	in Figure 2
8	5	5	$4\frac{1}{8}$	$4\frac{7}{8}$	3	2	[2], App. B.3	in [2], Appendix F.3
9	6	6	5	$5\frac{6}{9}$	2	2	[2], App. B.3	in Figure 2
10	6	6	5	5	2	2	[2], App. B.1	cyclic shift (1111010000) to the left.
11	7	6	5	5	3	3	[2], App. B.1	cyclic shift (11110100000) to the left.
12	8	8	7	7	2	2	[2], App. B.1	cyclic shift (110111101000) to the left.
13	7	6	5	5	≤ 4	≤ 4	[2], App. B.1	cyclic shift (1110110000000) to the left.
14	8	6	5	5	≤ 4	≤ 4	[2], App. B.1	cyclic shift (11101010000000) to the left.
15	8	8	7	7	3	3	[2], App. B.1	cyclic shift (101101110010000) to the left.
16	8	8	7	7	3	3	[2], App. B.1	cyclic shift (1111011010000000) to the left.

Fig. 1. Examples of hardware efficient linear layers over \mathbb{F}_2

$$\begin{pmatrix} 0110001 \\ 1000011 \\ 0100011 \\ 0001110 \\ 1001100 \\ 0110100 \\ 1011000 \end{pmatrix} \begin{pmatrix} 001110110 \\ 100101110 \\ 010010111 \\ 111101000 \\ 100110101 \\ 001111001 \\ 111010010 \\ 011001011 \\ 110001101 \end{pmatrix}$$

Fig. 2. Examples of $[14, 7, 4]$ and $[18, 9, 6]$ codes over \mathbb{F}_2

5 Software-Friendly Examples and the Cipher PRIDE

In this section, we describe our new lightweight software-cipher PRIDE, a 64-bit block cipher that uses a 128-bit key. We refer to Appendix D of [2] for a sketch of the security analysis and to the full version for more details.

We chose to design an SPN block cipher because it seems that this structure is better understood than e.g. ARX designs. We are, unsurprisingly, making use of the construction given in Theorem 1. We here decided on a linear layer with high dependency and a linear & differential branch number of 4. One key-observation is that the construction of Theorem 1 fits naturally with a bit-sliced implementation of the cipher, in particular with the Sbox layer. As a bit-sliced implementation of the Sbox layer is advantageous on 8-bit micro-controllers, in any case this is a nice match.

The target platform of PRIDE is Atmel’s AVR micro-controller [4], as it is dominating the market along with PIC [44] (see [45]). Furthermore, many implementations in literature are also implemented in AVR, we therefore opt for this platform to provide a better comparison to other ciphers (including SIMON and SPECK [10]). However, the reconfigurable nature of our search architecture (cf. Section 5.1) to find the basic layers of the cipher allows us to extend the search to various platforms in the future.

5.1 The Search for the Linear Layer

A natural choice in terms of Theorem 1 is to choose $k = 4$ and $b_1 = b_2 = b_3 = b_4 = 1$. Thus, the task reduces to find four 16×16 matrices forming one 64×64 matrix (to permute the whole state) of the following form:

$$\begin{pmatrix} L_0 & 0 & 0 & 0 \\ 0 & L_1 & 0 & 0 \\ 0 & 0 & L_2 & 0 \\ 0 & 0 & 0 & L_3 \end{pmatrix}$$

Each of these four 16×16 matrices should provide branch number 4 and together achieve high dependency with the least possible number of instructions. Instead of searching for an efficient implementation for a given matrix, we decided to search for the most efficient solution fulfilling our criteria.

To find such matrices (L_i) that could be implemented very efficiently given the AVR instruction set, we performed an extensive and hardware-aided tree search. Our search engine was optimized to look for AVR assembly code segments utilizing a limited set of instructions that would result in linear behaviour at matrix level. These are namely CLC, EOR, MOV, MOVW, CLR, SWAP, ASR, ROR, ROL, LSR, and LSL instructions. As we are looking for 16×16 matrices, the state to be multiplied with each L_i is stored in two 8-bit registers, which we call X and Y . We also allowed utilization of four temporary registers, namely $T0$, $T1$, $T2$, and $T3$. We designed and optimized our search engine according to these registers. Our search engine checks the resulting matrix L_i after N instructions to see if it provides the desired characteristics. While trying to reach instruction N , we try all possible instruction-register combinations in each step. This of course comes with an impractical time complexity, especially when N is increased further. To deal with this time complexity, we came up with several optimizations. As a first step, we limited the utilization of certain instruction-register combinations. For example, we excluded CLC and CLR instructions from the combinations for the first and last instructions. Also, EOR is not considered in the first instruction. Again, for the first and last instructions, SWAP, ASR, ROR, ROL, LSR, and LSL instructions are only used with X and Y . Furthermore, we did not allow temporary registers as the destination while trying MOV and MOVW instructions in the last instruction and $X - Y$ registers as the destination while trying MOV and MOVW instructions in the first instruction.

However, such optimizations were not enough to reduce the time complexity. We therefore applied further optimizations, i.e., when the matrices of all registers do not give full rank, we stop the search as we know that we cannot find an invertible linear layer any more.

In the end, we found matrices that fulfil all of our criteria starting from 7 instructions.

We implemented our search architecture on a Xilinx ML605 (Virtex-6 FPGA) evaluation board. The reconfigurable nature of the FPGA allowed us to change easily between different parameters, i.e. the number of instructions. The details of this search engine can be found in [33].

5.2 An Extremely Efficient Linear Layer

As a result of the search explained in Section 5.1, we achieved an extremely efficient linear layer. The cheapest solution provided by our search needed 36 cycles for the complete linear layer, which is what we opted for. The optimal matrices forming the linear layer are given in the Appendix G of [2]. Of these four matrices, L_0 and L_3 are involutions with the cost of 7 instructions (in turn, clock cycles), while L_1 and L_2 require 11 and 13 instructions for true and inverse matrices, respectively. The assembly codes are given in Appendix H of [2] to show the claimed number of instructions.

Comparing to linear layers of other SPN-based ciphers clearly demonstrated the benefit of our approach. Note however, that these comparisons have to be taken with care as not all linear layers operate on the same state size and do not offer the same security level. The linear layer of the ISO-standard lightweight cipher PRESENT [15] costs 144 cycles (derived from the total cycle count given in [24]). MixColumns operation of NIST-standard AES⁶ costs 117 instructions (but 149 cycles because of 3-cycle data load instruction utilizations, as MixColumns constants are implemented as look-up table – which means additional 256 bytes of memory, too) [6]. Note that ShiftRows operation was merged with the look-up table of Sbox in this implementation, so we take only MixColumns cost as the linear layer cost. The linear layer of another ISO-standard lightweight cipher CLEFIA [49] (again 128-bit cipher) costs 146 instructions and 668 cycles. Bit-sliced oriented design Serpent (AES finalist, 128-bit cipher) linear layer costs 155 instructions and 158 cycles. Other lightweight proposals, KLEIN [26] and mCrypton linear layers cost 104 instructions (100 cycles) and 116 instructions (342 cycles), respectively [23]. Finally, the linear layer cost of PRINCE is 357 instructions and 524 cycles⁷, which is even worse than AES. One of the reasons for this high cost is the non-cyclic 4×4 matrices forming the linear layer. The other reason is the ShiftRows operation applied on 4-bit state words, which makes coding much more complex than that of AES on an 8-bit micro-controller.

⁶ It is of course not fair to compare a 128-bit cipher with a 64-bit cipher. However, we provide AES numbers as a reference due to the fact that it is a widely-used standard cipher and its cost is much better compared to many lightweight ciphers.

⁷ We implemented this cipher on AVR, as we could not find any AVR implementations in the literature.

5.3 Sbox Selection

For our bit-sliced design, we decided to use a very simple (in terms of software-efficiency – the formulation is given in Appendix I of [2]) 10-instruction Sbox (which makes $10 \times 2 = 20$ clock cycles in total for the whole state). It is at the same time an involution Sbox, which prevents the encryption/decryption overhead. Besides being very efficient in terms of cycle count, this Sbox is also optimal with respect to linear and differential attacks. The maximal probability of a differential is $1/4$ and the best correlation of any linear approximation is $1/2$. The PRIDE Sbox is given below.

x	0x0	0x1	0x2	0x3	0x4	0x5	0x6	0x7	0x8	0x9	0xa	0xb	0xc	0xd	0xe	0xf
$S(x)$	0x0	0x4	0x8	0xf	0x1	0x5	0xe	0x9	0x2	0x7	0xa	0xc	0xb	0xd	0x6	0x3

The assembly codes are given in Appendix H of [2] to show the claimed number of instructions.

5.4 Description of PRIDE

Similar to PRINCE, the cipher makes use of the FX construction [34,13]. A pre-whitening key k_0 and post-whitening key k_2 are derived from one half of k , while the second half serves as basis k_1 for the round keys, i.e.,

$$k = k_0 || k_1 \quad \text{with} \quad k_2 = k_0.$$

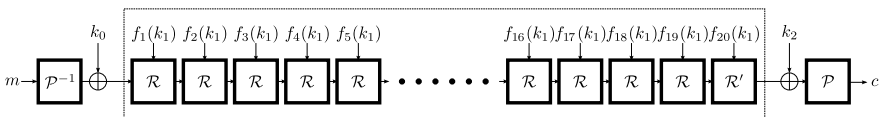
Moreover, in order to allow an efficient bit-sliced implementation, the cipher starts and ends with a bit-permutation. This clearly does not influence the security of PRIDE in any way. Note that in a bit-sliced implementation, none of the permutations P nor P^{-1} used in PRIDE has to be actually implemented explicitly. The cipher has 20 rounds, of which the first 19 are identical. Subkeys are different for each round, i.e., the subkey for round i is given by $f_i(k_1)$. We define

$$f_i(k_1) = k_{1_0} || g_i^{(0)}(k_{1_1}) || k_{1_2} || g_i^{(1)}(k_{1_3}) || k_{1_4} || g_i^{(2)}(k_{1_5}) || k_{1_6} || g_i^{(3)}(k_{1_7})$$

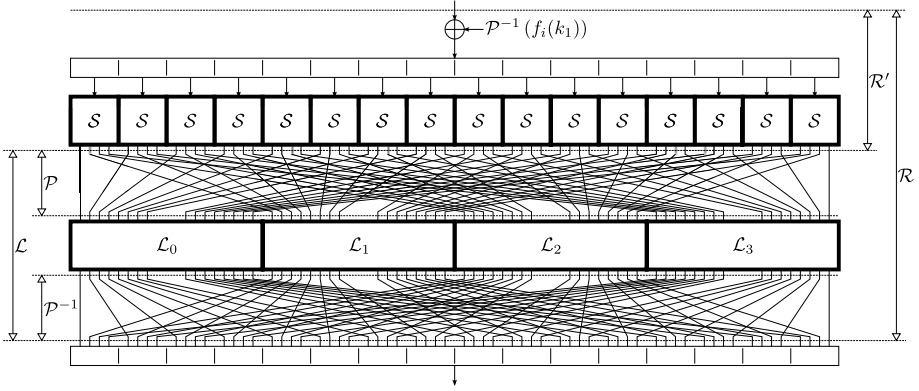
as the subkey derivation function with four byte-local modifiers of the key as

$$\begin{aligned} g_i^{(0)}(x) &= (x + 193i) \bmod 256, & g_i^{(1)}(x) &= (x + 165i) \bmod 256, \\ g_i^{(2)}(x) &= (x + 81i) \bmod 256, & g_i^{(3)}(x) &= (x + 197i) \bmod 256, \end{aligned}$$

which simply add one of four constants to every other byte of k_1 . The overall structure of the cipher is depicted here:



The round function \mathcal{R} of the cipher shows a classical substitution-permutation network: The state is XORed with the round key, fed into 16 parallel 4-bit Sboxes and then permuted and processed by the linear layer.



The difference between \mathcal{R} and \mathcal{R}' is that in the latter no more diffusion is necessary, therefore the last round ends after the substitution layer. With the software-friendly matrices we have found as described above, the linear layer is defined as follows (cf. Theorem 1 and Appendix G of [2]):

$$L := P^{-1} \circ (L_0 \times L_1 \times L_2 \times L_3) \circ P \quad \text{where} \quad P := P_{1,1,1,1}^{16}.$$

The test vectors for the cipher are provided in the Appendix J of [2].

5.5 Performance Analysis

As depicted above, one round of our proposed cipher PRIDE consists of a linear layer, a substitution layer, a key addition, and a round constant addition (key update). In a software implementation of PRIDE on a micro-controller, we also perform branching in each round of the cipher in addition to the previously listed layers. Adding up all these costs gives us the total implementation cost for one round of the cipher. The total cost can roughly be calculated by multiplying the number of rounds with the cost of each round. Note that we should subtract the cost of one linear layer from the overall cost, as PRIDE has no linear layer in the last round. The software implementation cost of the round function of PRIDE on Atmel AVR ATmega8 8-bit micro-controller [4] is presented in the following:

	Key update	Key addition	Sbox Layer	Linear Layer	Total
Time (cycles)	4	8	20	36	68
Size (bytes)	8	16	40	72	136

Comparing PRIDE to existing ciphers in literature, we can see that it outperforms many of them significantly both in terms of cycle count and code size. Note that we are not using any look-up tables in our implementation, in turn no RAMs⁸. The comparison with existing implementations is given below:

	AES-128 [24]	SERPENT-128 [24]	PRESENT-128 [24]	CLEFIA-128 [24]	SEA-96 [50]	NOEKEON-128 [23]	PRINCE-128	ITUbee-80 [32]	SIMON-64/128 [10]	SPECK-64/96 [10]	SPECK-64/128 [10]	PRIDE
$t(\text{cyc})$	3159	49314	10792	28648	17745	23517	3614	2607	2000	1152	1200	1514
bytes	1570	7220	660	3046	386	364	1108	716	282	182	186	266
eq.r.	5/10	1/32	4/31	1/18	8/92	1/16	5/12	12/20	33/44	34/26	34/27	

In the table, the first row is the time (performance) in clock cycles, the second row is the code size in bytes, and the third row is the *equivalent rounds*. The third row expresses the number of rounds for the given ciphers that would result in a total running time similar to PRIDE.

Note that, as we did not come across to any reference implementations in the literature, we implemented PRINCE in AVR for comparison. We also do not list the RAM utilization for the ciphers under comparison in the table.

In the implementation of PRIDE, our target was to be fast and at the same time compact. Note that we do not exclude data & key read and data write back as well as the whitening steps in our results (these are omitted in SIMON and SPECK numbers). Although the given numbers are just for encryption, decryption overhead is also acceptable: It costs 1570 clock cycles and 282 bytes. A cautionary note is indicated for the above comparison for several reasons. AES, SERPENT, CLEFIA, and NOEKOEN are working on 128-bit blocks; so, for a cycle per byte comparison, their cycle count has to be divided by a factor of two. Moreover, the ciphers differ in the claimed security level and key-size. PRIDE does not claim any resistance against related-key attacks (and actually can be distinguished trivially in this setting) and also generic time-memory trade-offs are possible against PRIDE in contrast to most other ciphers. Besides those restrictions, the security margin in PRIDE in terms of the number of rounds is (in our belief) sufficient.

One can see that PRIDE is comparable to SPECK-64/96 and SPECK-64/128 (members of NSA's *software*-cipher family), which are based on a Feistel structure and use modular additions as the main source of non-linearity.

In addition to the above table, the recent work of Grosso et al. [28] presents LS-Designs. This is a family of block ciphers that can systematically take advantage

⁸ Which has the additional advantage of increased resistance against cache-timing attacks.

of bit-slicing in a principled manner. In this paper, the authors make use of look-up tables. Therefore, a direct comparison with PRIDE is not fair as the use of look-up tables does not minimize the linear layer cost. However, to have an idea, we can try to estimate the cost of the 64-bit case of this family. They suggest two options: The first uses 4-bit Sbox with 16-bit Lbox, and the second uses 8-bit Sbox with 8-bit Lbox. The first option has 8 rounds, which results in 64 non-linear operations, 128 XORs, and 128 table look-ups in total. The second one has 6 rounds, which takes 72 non-linear operations, 144 XORs, and 48 table look-ups. For linear layer cost, we consider the XOR cost together with table look-ups. Unfortunately, it is not easy to estimate the overall cost of the given two options on AVR platform as the table look-ups take more than one cycle compared to the non-linear and linear operations. Another important point here to mention is that the use of look-up tables result in a huge memory utilization.

Finally, we note that, despite its target being software implementations, PRIDE is also efficient in hardware. It can be considered a hardware-friendly design, due to its cheap linear and Sbox layers.

6 Conclusion

In this work, we have presented a framework for constructing linear layers for block ciphers which allows to trade security against efficiency. For a given security level, in our case we focused on the branch number, we demonstrated techniques to find very efficient linear layers satisfying this security level. Using this framework, we presented a family of linear layers that are efficient in hardware. Furthermore, we presented a new cipher PRIDE dedicated for 8-bit micro-controllers that offers competitive performance due to our new techniques for finding linear layers.

One important question is on the optimality of a given construction for a linear layer. In particular, in the case of our construction, the natural question is if the reduction of the search space excludes optimal solutions and only sub-optimal solutions remain. For the hardware-friendly examples presented in Section 4 and Appendix C of [2], it is easy to argue that those constructions are optimal. Thus, in this case the reduction of the search space clearly did not have a negative influence on the results. In general, and for the linear layer constructed in Section 5 in particular, the situation is less clear. The main reason is that, again, the construction of linear layers is understudied and hence we do not have enough prior work to answer this question satisfactorily at the moment. Instead we view the PRIDE linear layer as a strong benchmark for efficient linear layers with the given parameters and encourage researchers to try to beat its performance.

Along these lines, we see this work as a step towards a more rigorous design process for linear layers. Our hope is that this framework will be extended in future. In particular, we would like to mention the following topic for further investigations. It seems that using an Sbox with a non-trivial branch number has the potential to significantly increase the number of active Sboxes when combined with a linear layer based on Theorem 1. Finding ways to easily prove such a result is worth investigating.

Finally, regarding PRIDE, we obviously encourage further cryptanalysis.

References

1. AES. Advanced Encryption Standard. FIPS PUB 197, Federal Information Processing Standards Publication (2001)
2. Albrecht, M.R., Driessen, B., Kavun, E.B., Leander, G., Paar, C., Yalçın, T.: Block Ciphers – Focus On The Linear Layer (feat. PRIDE): Full Version. IACR Cryptology ePrint Archive, 2014:453 (2014)
3. Anderson, R., Biham, E., Knudsen, L.: Serpent: A Proposal for the Advanced Encryption Standard (1998)
4. Atmel AVR. ATmega8 Datasheet, <http://www.atmel.com/images/doc8159.pdf>
5. Augot, D., Finiasz, M.: Direct Construction of Recursive MDS Diffusion Layers using Shortened BCH Codes. In: Fast Software Encryption (FSE). LNCS. Springer (to appear, 2014)
6. AVRAES: The AES block cipher on AVR controllers, <http://point-at-infinity.org/avraes/>
7. Barreto, P.S.L.M., Nikov, V., Nikova, S., Rijmen, V., Tischhauser, E.: Whirlwind: A New Cryptographic Hash Function. Des. Codes Cryptography 56(2-3), 141–162 (2010)
8. Barreto, P.S.L.M., Rijmen, V.: The Anubis Block Cipher. Submission to the NESSIE project (2001)
9. Barreto, P.S.L.M., Rijmen, V.: The Khazad Legacy-level Block Cipher. Submission to the NESSIE project (2001)
10. Beaulieu, R., Shors, D., Smith, J., Treatman-Clark, S., Weeks, B., Wingers, L.: The SIMON and SPECK Families of Lightweight Block Ciphers. IACR Cryptology ePrint Archive, 2013:414 (2013)
11. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Keccak Specifications (2009)
12. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. In: Menezes, A., Vanstone, S.A. (eds.) CRYPTO 1990. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (1991)
13. Biryukov, A.: DES-X (or DESX). In: Encyclopedia of Cryptography and Security, 2nd edn., p. 331. Springer (2011)
14. Biryukov, A., De Cannière, C., Braeken, A., Preneel, B.: A Toolbox for Cryptanalysis: Linear and Affine Equivalence Algorithms. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 33–50. Springer, Heidelberg (2003)
15. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., Seurin, Y., Vikkelsø, C.: PRESENT: An Ultra-Lightweight Block Cipher. In: Paillier, P., Verbaauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
16. Borghoff, J., et al.: PRINCE – A Low-Latency Block Cipher for Pervasive Computing Applications - Extended Abstract. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 208–225. Springer, Heidelberg (2012)
17. Brinkmann, M., Leander, G.: On the Classification of APN Functions Up to Dimension Five. Des. Codes Cryptography 49(1-3), 273–288 (2008)
18. Carlet, C.: Vectorial Boolean Functions for Cryptography. In: Boolean Methods and Models. Cambridge University Press (2010)
19. Daemen, J.: Cipher and Hash Function Design, Strategies Based On Linear and Differential Cryptanalysis. PhD thesis, Katholieke Universiteit Leuven (1995)
20. Daemen, J., Knudsen, L., Rijmen, V.: The Block Cipher SQUARE. In: Biham, E. (ed.) FSE 1997. LNCS, vol. 1267, pp. 149–165. Springer, Heidelberg (1997)

21. Daemen, J., Rijmen, V.: The Wide Trail Design Strategy. In: Honary, B. (ed.) *Cryptography and Coding 2001*. LNCS, vol. 2260, pp. 222–238. Springer, Heidelberg (2001)
22. DES: Data Encryption Standard. FIPS PUB 46, Federal Information Processing Standards Publication (1977)
23. Eisenbarth, T., et al.: Compact Implementation and Performance Evaluation of Block Ciphers in ATtiny Devices. In: Mitrokotsa, A., Vaudenay, S. (eds.) *AFRICACRYPT 2012*. LNCS, vol. 7374, pp. 172–187. Springer, Heidelberg (2012)
24. Engels, S., Kavun, E.B., Mihajloska, H., Paar, C., Yalçın, T.: A Non-Linear/Linear Instruction Set Extension for Lightweight Block Ciphers. In: *ARITH'21: 21st IEEE Symposium on Computer Arithmetics*. IEEE Computer Society (2013)
25. Gauravaram, P., Knudsen, L., Matusiewicz, K., Mendel, F., Rechberger, C., Schläer, M., Thomsen, S.: Grøstl. SHA-3 Final-round Candidate (2009)
26. 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)
27. Grassl, M.: Bounds On the Minimum Distance of Linear Codes and Quantum Codes (2007), <http://www.codetables.de>
28. Grosso, V., Leurent, G., Standaert, F.-X., Varıcı, K.: LS-Designs: Bitslice Encryption for Efficient Masked Software Implementations. In: *Fast Software Encryption (FSE)*. LNCS. Springer (to appear, 2014)
29. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON Family of Lightweight Hash Functions. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 222–239. Springer, Heidelberg (2011)
30. Guo, J., Peyrin, T., Poschmann, A., Robshaw, M.J.B.: The LED Block Cipher. In: Preneel, B., Takagi, T. (eds.) *CHES 2011*. LNCS, vol. 6917, pp. 326–341. Springer, Heidelberg (2011)
31. Intel. Advanced Encryption Standard Instructions, Intel AES-NI (2008)
32. Karakoç, F., Demirci, H., Harmancı, A.E.: ITUbee: A Software Oriented Lightweight Block Cipher. In: Avoine, G., Kara, O. (eds.) *LightSec 2013*. LNCS, vol. 8162, pp. 16–27. Springer, Heidelberg (2013)
33. Kavun, E.B., Leander, G., Yalçın, T.: A Reconfigurable Architecture for Searching Optimal Software Code to Implement Block Cipher Permutation Matrices. In: *International Conference on ReConFigurable Computing and FPGAs (ReConFig)*. IEEE Computer Society (2013)
34. Kilian, J., Rogaway, P.: How to Protect DES Against Exhaustive Key Search (An Analysis of DESX). *J. Cryptology* 14(1), 17–35 (2001)
35. 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)
36. 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)
37. Lee, R.B., Fıŝkıran, M., Wang, M., Hilewitz, Y., Chen, Y.-Y.: PAX: A Cryptographic Processor with Parallel Table Lookup and Wordsize Scalability. Princeton University Department of Electrical Engineering Technical Report CE-L2007-010 (2007)
38. Lee, R.B., Shi, Z., Yang, X.: Efficient Permutation Instructions for Fast Software Cryptography. *IEEE Micro* 21(6), 56–69 (2001)

39. 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)
40. Lin, S., Costello, D.J. (eds.): Error Control Coding, 2nd edn. Prentice Hall (2004)
41. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: Helleseht, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994)
42. McGregor, J.P., Lee, R.B.: Architectural Enhancements for Fast Subword Permutations with Repetitions in Cryptographic Applications. In: 19th International Conference on Computer Design (ICCD 2001), pp. 453–461 (2001)
43. Nyberg, K.: Differentially Uniform Mappings for Cryptography. In: Helleseht, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 55–64. Springer, Heidelberg (1994)
44. PIC. 12-Bit Core Instruction Set
45. PIC vs. AVR, <http://www.ladyada.net/library/picvsavr.html>
46. Saarinen, M.-J.O.: Cryptographic Analysis of All 4×4 -Bit S-Boxes. In: Miri, A., Vaudenay, S. (eds.) SAC 2011. LNCS, vol. 7118, pp. 118–133. Springer, Heidelberg (2012)
47. Sajadieh, M., Dakhilalian, M., Mala, H., Sepehrdad, P.: Recursive Diffusion Layers for Block Ciphers and Hash Functions. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 385–401. Springer, Heidelberg (2012)
48. Shi, Z.J., Yang, X., Lee, R.B.: Alternative Application-Specific Processor Architectures for Fast Arbitrary Bit Permutations. IJES 3(4), 219–228 (2008)
49. Shirai, T., Shibutani, K., Akishita, T., Moriai, S., Iwata, T.: The 128-Bit Block-cipher CLEFIA (Extended Abstract). In: Biryukov, A. (ed.) FSE 2007. LNCS, vol. 4593, pp. 181–195. Springer, Heidelberg (2007)
50. 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)
51. Suzaki, T., Minematsu, K., Morioka, S., Kobayashi, E.: TWINE: A Lightweight Block Cipher for Multiple Platforms. In: Knudsen, L.R., Wu, H. (eds.) SAC 2012. LNCS, vol. 7707, pp. 339–354. Springer, Heidelberg (2013)
52. Ullrich, M., De Cannière, C., Indesteege, S., Küçük, Ö., Mouha, N., Preneel, B.: Finding Optimal Bitsliced Implementations of 4×4 -Bit S-boxes. In: Symmetric Key Encryption Workshop (2011)
53. Wu, S., Wang, M., Wu, W.: Recursive Diffusion Layers for (Lightweight) Block Ciphers and Hash Functions. In: Knudsen, L.R., Wu, H. (eds.) SAC 2012. LNCS, vol. 7707, pp. 355–371. Springer, Heidelberg (2013)
54. Wu, W., Zhang, L.: LBlock: A Lightweight Block Cipher. In: Lopez, J., Tsudik, G. (eds.) ACNS 2011. LNCS, vol. 6715, pp. 327–344. Springer, Heidelberg (2011)