

A Generic Approach to Invariant Subspace Attacks: Cryptanalysis of Robin, iSCREAM and Zorro

Gregor Leander¹(✉), Brice Minaud², and Sondre Rønjom³

¹ Horst Görtz University for IT Security, Ruhr-Universität Bochum, Bochum, Germany

gregor.leander@rub.de

² Agence Nationale de la Sécurité des Systèmes d'Information, Paris, France

brice.minaud@gmail.com

³ Nasjonal sikkerhetsmyndighet, Oslo, Norway

sondrer@gmail.com

Abstract. Invariant subspace attacks were introduced at CRYPTO 2011 to cryptanalyze PRINTCIPHER. The invariant subspaces for PRINTCIPHER were discovered in an ad hoc fashion, leaving a generic technique to discover invariant subspaces in other ciphers as an open problem. Here, based on a rather simple observation, we introduce a generic algorithm to detect invariant subspaces. We apply this algorithm to the CAESAR candidate iSCREAM, the closely related LS-design Robin, as well as the lightweight cipher Zorro. For all three candidates invariant subspaces were detected, and result in practical breaks of the ciphers. A closer analysis of independent interest reveals that these invariant subspaces are underpinned by a new type of self-similarity property. For all ciphers, our strongest attack shows the existence of a weak key set of density 2^{-32} . These weak keys lead to a simple property on the plaintexts going through the whole encryption process with probability one. All our attacks have been practically verified on reference implementations of the ciphers.

Keywords: Cryptanalysis · Lightweight cryptography · Invariant subspace · Self-similarity · iSCREAM · LS-designs · Zorro · CAESAR

Introduction

Block ciphers are one of the most essential cryptographic primitives. Our understanding of how to build secure block ciphers has greatly advanced in the last 20 years. Nowadays, analyzing a given block cipher with respect to a large class of non-trivial attacks, including linear and differential attacks and their variants, is a well-understood process for a large class of block ciphers. However,

G. Leander—The work of Gregor Leander was funded by the BMBF UNIKOPS project.

when it comes to designing block ciphers with strong performance requirements, often less conservative approaches are chosen. Examples of such performance requirements that have recently been studied extensively include low hardware footprint (e.g. PRESENT [6], LED [20], KATAN [10]), low memory consumption on small embedded processors (e.g. ITUBee [21], SPECK [5], PRIDE [2]), low latency (e.g. PRINCE [7]) and ease of side-channel protection (e.g. Zorro [14], LS-Designs [15]).

In order to fit within constrained settings, many of these ciphers feature innovative designs: they may rely on simpler round functions, or minimal key schedules. While in most cases, guarantees against traditional linear or differential attacks are still offered, the simpler structure of many of these ciphers may lend itself to new attacks. Careful cryptanalysis is required in order to assess the security of these new designs; in this process, new techniques have emerged.

One such technique is the invariant subspace attack, introduced in [22] for the cryptanalysis of PRINTCIPHER. The general idea behind this attack is the following: assume that the round function of a cipher maps a coset A of some vector subspace of the inner state to a coset B of the same space, and a fixed key belonging to $A - B$ is added in every round. Then the set A is preserved by the round function, and hence remains stable through the whole encryption process. This property holds for a large set of keys in PRINTCIPHER, breaking the cipher in a practical setting. This type of attack seems particularly well-suited to substitution-permutation networks (SPN) with a minimal key schedule or cryptographic permutations with highly structured round constants.

Invariant subspace attacks are unusual in that they rely on an unexpected form of symmetry in the round function, and yield attacks that are independent of the number of rounds. On the other hand, the same attributes are shared by attacks based on self-similarity properties. These properties were first formally defined in [4] to study alternative descriptions of AES, and later used in [8] to cryptanalyze the SHA-3 candidate *Lesamnta* and the lightweight cipher XTEA.

Interestingly, despite its fundamental nature, the understanding of symmetries and invariant subspaces, in block ciphers or cryptographic permutations, is rather limited. The invariant subspace attack on PRINTCIPHER was found in an ad hoc fashion and no general approach to detect or avoid such invariant spaces is known. This is even more surprising as for more involved attacks like differential and linear attacks and their variations our general understanding of detection and avoiding those attacks by design is much more evolved.

Our Contribution. In this paper we aim at increasing the general understanding of invariant subspaces. For this purpose, and as our first main result, we present a generic algorithm that is able to detect invariant subspaces. The running time of this algorithm depends on the block size of the primitive and the density of the existing invariant subspaces. In particular, it is especially efficient if relatively large invariant subspaces exist. As the impact of an invariant subspace increases with its dimension, this can be seen as detecting stronger attacks significantly faster than minor attacks.

We apply this generic algorithm to the lightweight cipher Robin introduced at FSE 2014 as a concrete instance of the LS-design framework [15], the closely related CAESAR [1] candidate iSCREAM [18], as well the lightweight cipher Zorro presented at CHES 2013 [14]. In all cases the algorithm is able to detect invariant subspaces.¹ Attacks resulting from these invariant subspaces break all three ciphers in a practical setting. All attacks have been verified on reference implementations of the ciphers.

As our second main contribution, we show that the invariant subspaces we have discovered are underpinned by a type of self-similarity property of independent interest, stemming from a linear map commuting with the round function. Surprisingly, such a map exists for all three ciphers, despite Robin and Zorro having quite different structures. As a result, we obtain stronger attacks on our target ciphers. We also hope to provide useful insight for the design and analysis of ciphers with minimal key schedules as well as for the choice of round constants in cryptographic permutations.

More specifically, our attacks show the existence of weak keys in Robin, Zorro, as well as iSCREAM in the chosen-tweak scenario. In all cases, the proportion of weak keys is 2^{-32} within the set of all keys. Encryption of a single chosen plaintext is enough to determine whether a key is weak, making our weak key setting very practical. Once a key is recognized as weak, a simple property on plaintexts goes through the whole encryption process with probability one, breaking plaintext confidentiality. In addition, the full 128-bit encryption key can be recovered using one chosen plaintext in time complexity 2^{64} . We also show that all three ciphers are instantly broken in the related-key setting, without any weak key requirement; although only iSCREAM claims security in this model.

In the case of LS-designs, we furthermore present a second attack, based on S-box-dependent invariant subspaces, without an underlying self-similarity. We obtain a new set of weak keys with the same properties as above, including the fact that they can be detected using a single chosen plaintext. However when applying this second attack to Robin and iSCREAM, we obtain a much rarer set of weak keys, with only one key in 2^{80} being weak. We then fine-tune this attack against iSCREAM, and obtain a ratio of weak keys of 2^{-48} , at the cost of requiring 2^{32} chosen-tweak chosen plaintexts in order to detect whether a key is weak; once a weak key is detected, the full 128-bit key can be recovered in time complexity 2^{48} with no additional data.

Regarding LS-designs, it should be pointed out that while our first attack breaks Robin and iSCREAM in a practical setting, Fantomas and SCREAM appear to be safe. Moreover, in the case of Robin and iSCREAM, a careful tweak of the ciphers should be able to prevent our attacks. Thus, the security of the LS-design framework in general is not called into question. On the contrary, in the case of Zorro, our attack adds to other attacks suggesting that partial nonlinear layers should be approached with caution.

¹ The source code of our tool is available at invariant-space.forge.inria.fr.

Related Work. Invariant subspace attacks were introduced in [22]. Their application to PRINTCIPHER relies on undesirable properties induced by its 3-bit S-boxes. By contrast, most of our attacks (except the second attack on Robin and iSCREAM) are actually independent of a particular choice of S-boxes.

A thorough analysis of invariant subspaces in PRINTCIPHER was subsequently carried out in [9]. Using a dedicated tool, the authors were able to enumerate all invariant subspaces of PRINTCIPHER, of the type uncovered in the previous article. However their approach is tailored to PRINTCIPHER, and does not extend to other ciphers.

An older line of work has studied “linear factors” of DES [11, 13, 28], which bear some resemblance to invariant subspaces. The existence of a linear factor is an even stronger property than that of an invariant subspace: essentially, it asks that a (linearly defined) portion of the ciphertext only depend on (linearly defined) portions of the plaintext and key. Nonetheless, it is interesting to note that our attacks do uncover a linear factor in Robin and Zorro (the subcipher in Section 3.3), although only in a weak key setting.

Along a similar line, the attack on SAFER in [25] should be mentioned. It exploits the action of the cipher on cosets of a vector space as a whole, rather than isolating a specific trail or characteristic.

Self-similarity properties were used to attack hash functions and block ciphers in [8]. It should be noted that self-similarity is a very wide framework, encompassing attacks ranging from probability one related-key differentials to slide attacks. To the best of our knowledge, the commutation property we consider here is very different from any previous work.

There is no prior cryptanalysis of LS-designs. As for iSCREAM, an issue with the padding in the original CAESAR submission of SCREAM and iSCREAM was pointed out in [29] and subsequently corrected. Our attacks have caused iSCREAM to be temporarily withdrawn from the CAESAR competition for a redesign [17].

By contrast, many attacks have been carried out against Zorro, mostly differential or linear in nature [3, 19, 27, 30, 31]. The best attack in [3] is a differential attack requiring 2^{41} data and time complexity 2^{45} to break the full cipher. Our attack is of a different nature: it holds in the weak key setting (with 2^{96} weak keys out of 2^{128}), requires minimal data and time, and is independent of the number of rounds. Similar to [3], our attack can be readily extended to Zorro-like ciphers, as shown in the ePrint version of this work [23].

Structure of the Paper

In Section 1, we recall the definitions of invariant subspace and present our generic algorithm for detecting such invariant subspaces. In Section 2, we provide a description of LS-designs, including our targets Robin and iSCREAM. In Sections 3, 4 and 5, we develop our attacks against LS-designs, introduce a particular self-similarity property, the resulting invariant subspaces, and finally describe a different invariant subspace attack not underpinned by self-similarity. In Section 6, we apply our self-similarity and invariant subspace attacks to Zorro. Finally, in Section 7, we conclude with a discussion of our results and outline interesting open problems.

1 A Generic Algorithm to Detect Invariant Subspaces

In this section we first recall the invariant subspace attack and later present our algorithmic approach to detect invariant subspaces in a generic manner.

1.1 Invariant Subspace Attacks

Invariant subspace attacks were introduced and applied to PRINTCIPHER in [22]. We briefly recall the basic principle here.

Consider a n -bit block cipher with round function F_K consisting of a key addition and a SP layer $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$. That is, F_K is defined by $F_K(x) = F(x + K)$. Assume the SP-layer F is such there exists a subspace $A \subseteq \mathbb{F}_2^n$ and two constants $u, v \in \mathbb{F}_2^n$ with the property:

$$F(u + A) = v + A$$

Then, given a (round) key $K \in u - v + A$, i.e. $K = K_A + u - v$ with $K_A \in A$, the following holds:

$$F_K(v + A) = F(v + A + u - v) = F(u + A) = v + A$$

i.e. the round function maps the affine subspace $v + A$ onto itself. If all round keys are in $u - v + A$, in particular if identical round keys are used as in LS-designs and Zorro, then this property is iterative over an arbitrary number of rounds.

In the case where an identical key is added in every round (there is no key schedule), a key is said to be weak iff it belongs to $u - v + A$. Whenever a key is weak, plaintexts in $v + A$ are mapped to ciphertexts in $v + A$, breaking plaintext confidentiality. The number of weak keys is the cardinality of A .

In order to detect whether an unknown key is weak, it is enough to encrypt one plaintext in $v + A$, and test whether the resulting ciphertext is in the same space. Indeed, over the set of all keys, false positives will occur with the same frequency as true positives, and can be discarded with a second chosen plaintext.

1.2 A Generic Algorithm

In this section we present a simple and entirely generic probabilistic algorithm able to discover invariant subspaces for a given round function. The algorithm gives instant results for vector subspaces, and is able to discover affine subspaces in time proportional to their density. Despite its simplicity, this algorithm is enough to automatically discover all invariant subspace attacks to be elaborated upon in the following sections.

The algorithm will identify minimal invariant subspaces and thereby identify invariant subspace attacks automatically. However, further analysis usually allows to significantly improve upon the attacks recovered automatically by the algorithm and gain further insights in the structure of the detected weakness. Furthermore, as the expected running time is determined by the density of invariant subspaces, it might well be that not all possible attacks are detected. Thus, for the moment, this generic algorithm cannot be used to fully exclude the existence of invariant subspaces.

Identifying Minimal Subspaces. Assume we are given a permutation $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$. Here F could be a (keyless) round of a block cipher or a cryptographic permutation (like Keccak- f). Our goal is to find affine subspaces $u + A \subset \mathbb{F}_2^n$ such that:

$$F(u + A) = v + A$$

for some $v \in \mathbb{F}_2^n$.

Our algorithm is based on the following trivial observation.

Lemma 1. *Assume $u + A$ is an affine subspace such that $F(u + A)$ is also an affine subspace $v + A$. Then for any subset $X \subseteq A$, the linear span of $(F(u + X) - v) \cup X$ is contained in A .*

The idea is to first guess one possible offset u' of the affine space to be found and use $v' = F(u')$. Next, we guess a one-dimensional subspace of A , denote this by A_0 . The algorithm will succeed if and only if $u' + A_0$ is contained in $u + A$.

1. We compute A_{i+1} from A_i as:

$$A_{i+1} = \text{span}\{(F(u' + A_i) - v') \cup A_i\}$$

2. If the dimension of A_{i+1} equals the dimension of A_i , we found an invariant subspace and exit.
3. If not, we continue with step 1.

Thus, the idea is to start with what we denote *nucleon* of A and map it using F until it stabilizes. In the case that our initial guess was wrong and $u' + A_0$ is not contained in some non-trivial invariant subspace we will end up with the full space after at most n iterations of the above.

Note that it is not necessary to really map the complete spaces A_i using F but a randomly chosen subset of relatively small size is enough for our purpose and significantly speeds up the process.

If the largest invariant subspace of F has dimension d , the algorithm will detect this space (or any invariant subspaces of this space) after an expected number of $2^{2(n-d)}$ guesses for A_0 and u' . Thus, in this basic form, the algorithm becomes quickly impractical. However, in the case of round functions of a cipher (or a cryptographic permutation) that differ by round constants only, its running time can be greatly improved as described next.

Knowing the Nucleon. For block ciphers with identical round keys or cryptographic permutations, we actually have a very good idea about the nucleon we want to be included in the space A , namely the round constants. More precisely, we consider round functions $F_i : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ that differ only by the addition of constants, i.e.

$$F_i(x) = F(x) + c_i$$

for $c_i \in \mathbb{F}_2^n$, where for simplicity we assume $c_0 = 0$. We are looking for affine subspaces $u + A$ that are mapped to $v + A$ by all round functions. In particular

$$F_0(u + A) = F(u + A) = v + A$$

and

$$F_i(u + A) = F(u + A) + c_i = v + A$$

which implies

$$v + A = c_i + v + A$$

and thus $c_i \in A$. Thus, given the situation as above, any subspace that is invariant under all round functions must necessarily contain the linear span of all round constants c_i .

For the algorithm outlined above this has significant consequences. Here, the only thing we have to guess is the offset. Therefore, the expected number of iterations of the algorithm is reduced from $2^{2(n-d)}$ to 2^{n-d} .

Moreover, after running the algorithm for m iterations with randomly chosen guesses for the offset, the probability that an invariant subspace of dimension d is not detected by the approach is given by

$$p_{m,n,d} := (1 - 2^{n-d})^m$$

which can be approximated by

$$\log p_{m,n,d} \approx -m2^{d-n}.$$

The Algorithm.

```

1: procedure CLOSURE(function  $F$ , nucleon  $A$ , offset  $u$ )
2:    $v \leftarrow F(u)$ 
3:   StableCount  $\leftarrow 0$ 
4:   while StableCount  $< N$  do
5:     Pick a random  $x \xrightarrow{\$} u + \text{span}\{A\}$ 
6:     if  $F(x) - v \in \text{span}\{A\}$  then
7:       StableCount = StableCount + 1
8:     else
9:       Add  $F(x) - v$  to  $A$ 
10:      StableCount  $\leftarrow 0$ 
11:    end if
12:  end while
13:  return  $u + \text{span}\{A\}$ 
14: end procedure

```

For offset u and nucleon A , the above procedure outputs the smallest affine subspace containing $u + \text{span}\{A\}$, that is mapped to a coset of the same space by F (with high probability). The algorithm depends on a global parameter N that controls the risk of error. Namely, when the algorithm exits, elements of $u + \text{span}\{A\}$ are mapped to $v + \text{span}\{A\}$ with probability greater than $1 - 2^{-N}$. This probabilistic result is enough for an invariant subspace attack to go through even for moderate choices of N .

Guessing the Offset. If we are actually looking for stable *vector* spaces rather than affine spaces, as will be the case in the S-box independent setting described in Section 3.2, guessing the offset is not needed: we can choose zero as the offset. Then the algorithm above finds the smallest invariant subspace instantly.

In the general case where we are looking for any (affine) invariant subspace, we need to guess one offset u belonging to the affine space we are searching for. Then we can run the procedure above to find the generated invariant subspace, if it exists (otherwise, the algorithm will simply output the full space). If the space we are looking for has dimension d , guessing such an offset u by brute force will require 2^{n-d} tries on average. Of course we just require *one* invariant subspace; so in general 2^{n-d} can be replaced by the density of vectors belonging to (non-trivial) invariant subspaces.

Each iteration of the algorithm requires Gaussian reduction to determine whether a certain n -bit vector belongs to some subspace, amounting to n^2 operations. Hence the overall running time to find an invariant subspace of dimension d is roughly $n^2 \cdot 2^{n-d}$. Thus if n is large, the above approach will only work if $n - d$ is relatively small, or more generally the density of invariant subspaces is large. The case where n is small is also useful in order to find invariant subspaces through a single S-box: this is how we found spaces in Appendix B (after making the algorithm deterministic and exhaustive, which is affordable for small n).

1.3 Applications

We applied the algorithm to the block ciphers Zorro, Robin, Fantomas, LED and NOEKEON, as well as to the CAESAR candidate iSCREAM. We chose $N = 50$ to be very conservative. We ran the algorithm with approximately 2^{34} iterations for each primitive, stopping earlier in the case where an invariant subspace was detected. The results are summarized in the table below.

Table 1. Experimental Results: Here n is the block size and $d_{0.001}$ is the smallest dimension of an invariant subspace that has a probability to exist upper bounded by 0.001

Primitive	n	Dimension found	$d_{0.001}$	Running Time (h)
LED	64	-	34	24
NOEKEON [12]	128	-	98	40
Fantomas	128	-	98	40
Robin	128	96	-	22
iSCREAM	128	96	-	22
Zorro	128	96	-	<1

For LED, NOEKEON and Fantomas, no invariant subspaces were detected given our limited iterations. In that case, Table 1 indicates the dimension $d_{0.001}$ of the largest invariant subspace that has a probability to exist upper bounded by 0.001. More precisely, if x denotes the codimension of the largest invariant

subspace, each random guess of an offset has probability 2^{-x} of falling into this subspace. After T tries, the probability of not having found the subspace is thus $(1-2^{-x})^T \approx e^{-T2^{-x}}$. We want this probability to be $1/1000$ within $T = 2^{32}$ tries, which yields $x = 32 - \log(\ln(1000)) \approx 30$, so $d_{0.001} \approx n - 30$. Thus it is unlikely that invariant subspaces of dimension above 98 exist for NOEKEON. However, the existence of smaller subspaces cannot be excluded with high probability by our results.

As we will show below, for Zorro, Robin and the CAESAR candidate iSCREAM the largest invariant subspace has dimension 96 out of 128, i.e. density 2^{-32} . Thus the time complexity is expected to be 2^{32} Gaussian eliminations on 128×128 binary matrices. Our experiments confirm this estimation. Discovering the invariant subspace took 22 hours on a single desktop PC equipped with an Intel Xeon Core i7 with 12 virtual cores used in parallel.

In the case of Zorro, we chose to use a single round as target function, rather than the four rounds separating key addition. It turns out many cosets of the invariant subspace in Appendix A are sent to another coset by a single round (namely, all cosets stemming with offsets where cells 0 and 3 are equal). Our generic approach discovers this fact and the associated subspace instantly, hence the “< 1” time in the previous table.

As mentioned in the introduction, a detailed analysis of the findings of the generic algorithm allows to understand the underlying structure of the invariant subspaces we have found, and improve the attacks. We present those findings in the following sections.

2 Description of LS-Designs, Robin, and iSCREAM

2.1 LS-Designs

LS-designs were introduced by Grosso, Leurent, Standaert and Varici at FSE 2014 [15]. We refer the interested reader to their article for a detailed presentation of LS-designs and their design rationale. For our purpose, a brief technical description suffices.

An LS-design is a block cipher encrypting n -bit plaintext blocks using a n -bit key. The inner state of the cipher, as well as the plaintext, ciphertext, and key, are all represented as an $r \times c$ bit array, with r the number of rows and c the number of columns. A concrete LS-design is parametrized by the following components:

- A choice of r and c . The size of the key and message blocks is $n = r \cdot c$.
- An r -bit S-box s .
- A bijective linear map ℓ on c -bit vectors, called the L-box.
- A number of rounds t .
- A choice of k -bit round constants $C(i)$ for $1 \leq i \leq t$.

In order to encrypt a given n -bit plaintext block, the plaintext is first loaded into the inner state of the cipher, and the master key is added in (all additions

are bitwise XORs). Then a round function is applied successively for rounds 1 to t . At that point the ciphertext is equal to the inner state. The round function at round i proceeds as follows:

1. The round constant $C(i)$ is added to the inner state.
2. The S-box s is applied to each column of the state.
3. The L-box ℓ is applied to each row of the state.
4. The n -bit master key K is added to the state.

2.2 Notation

When dealing with LS-designs, we will always use the previous notation; that is:

- r the number of rows of the state.
- c the number of columns.
- n the size of the state; that is, $n = r \cdot c$.
- s the r -bit S-box.
- S the S-box step; that is, the application of s on each column of the state.
- ℓ the $c \times c$ binary matrix representing the linear layer, identified with the corresponding linear map on \mathbb{F}_2^c .
- L the L-box step; that is, application of ℓ on each row of the state.

2.3 Robin

In [15], two concrete LS-designs are proposed, Robin and Fantomas. The idea behind Robin is that both the S-box and L-box are involutive. This allows the same circuitry to be reused when computing these components and their inverse operation, i.e. when encrypting and decrypting. This saves valuable space on embedded devices when both encryption and decryption capabilities are required. The trade-off is that involutive components have more structure, resulting in a slightly higher number of rounds to reach the same security level as an LS-design based on non-involutive components.

Robin strictly fits within the LS-design framework recalled in the previous section. As such it can be fully described by the following parameters:

- The inner state of Robin has 8 rows and 16 columns, resulting in 128-bit blocks and a 128-bit key.
- The 8-bit involutive S-box is given in [15].
- The 16-bit involutive L-box is depicted as a 16×16 binary matrix on Fig. 1.
- The number of rounds is 16.
- At round i (starting from 1), the round constant $C(i)$ is zero outside of the first row, where it is equal to $\ell(i)$, with ℓ the L-box matrix.

2.4 iSCREAM

SCREAM and iSCREAM [18] are two authenticated ciphers closely related to LS-designs. In fact iSCREAM is essentially a tweaked version of Robin, together with a Tweakable Authenticated Encryption (TAE) mode of operation [24].




Fig. 1. Matrix representing the L-box of Robin and iSCREAM. Dark cells stand for 1's and white cells for 0's.

Meanwhile SCREAM is similar to Fantomas, with a different linear layer. The TAE mode of operation requires a tweakable block cipher [24]. Accordingly, the difference between the block cipher underlying iSCREAM and Robin stems from the introduction of a 128-bit tweak T into the (previously non-existent) key schedule.

In the remainder of this article we focus on weaknesses of the block cipher on which iSCREAM is built, independently of the mode of operation. We may abuse notations and write iSCREAM to mean its underlying block cipher.

This block cipher can be described as an LS-design, except for the fact that during the key addition phase, instead of adding in K every round: at odd rounds, $K + T$ is added; while at even rounds, $T \lll_c 1$ is added, where $T \lll_c 1$ denotes a circular shift of the columns of T by one column towards the left. The combination of two rounds is called a step. Beside that, iSCREAM can be described by the following parameters:

- The inner state of iSCREAM has 8 rows and 16 columns, resulting in 128-bit blocks and a 128-bit key.
- The S-box and L-box are those of Robin.
- The number of rounds depends on the required security level. The original article lists six variants. However the primary recommendation for iSCREAM as per CAESAR requirements is 12 steps (24 rounds) [16]. A secondary recommendation claiming related-key security has 14 steps (28 rounds). Since our attacks are essentially independent of the number of rounds, we omit other variants.
- At round i (starting from 1), the round constant $C(i)$ is zero outside of the first row, where it is equal to $27 \cdot i$ modulo 256 (affecting only the first 8 bits of the row).

3 Invariant Permutation Attack

In the next two sections, we analyze the invariant subspace discovered by the generic algorithm on Robin and iSCREAM. This subspace is actually induced by a particular type of self-similarity of independent interest, as is the invariant subspace of Zorro. Using this self-similarity directly results in an even stronger attack, as will be discussed below.

We start by recalling the concept of self-similarity and explaining a link to commutating linear maps. These concepts will afterwards be applied to LS-designs in general and to Robin and iSCREAM in particular (as well as to Zorro in Section 6).

3.1 Self-Similarity Properties and Linear Commutant

In [4] and [8], self-similarity in general is defined as:

Definition 1 (Self-similarity in a block cipher). *For a fixed block cipher E , let $E_K(x)$ denote the ciphertext block resulting from the encryption of plaintext block x under key K . A self-similarity relation is given by invertible and efficiently computable mappings ϕ, ψ, θ such that:*

$$\forall K, x : \theta(E_K(x)) = E_{\psi(K)}(\phi(x))$$

What we are interested in is the case where $M = \phi = \psi = \theta$ is a linear map. This situation will arise if the cipher follows a generalized Even-Mansour structure where key-independent round functions F_i alternate with the addition of a fixed key K (i.e. no key schedule); and M commutes with the round functions F_i . This last condition is very demanding; but this is precisely what happens in both Robin and Zorro, despite their different structure. We expand on why this might be the case in the discussion (Section 7). The following lemma sums up the attack.

Lemma 2. *Consider a block cipher composed of round functions F_i separated by addition of a fixed key K . Suppose there exists a linear map M such that M commutes with the F_i 's. Then:*

$$\forall x : M(E_K(x)) = E_{M(K)}(M(x))$$

In particular, if $K = M(K)$:

$$\forall x : M(E_K(x)) = E_K(M(x))$$

The commutativity of M and the round functions can be interpreted from the invariant subspace perspective. Indeed, if we let $A = \ker(M^i + Id)$ for any i , A is an invariant subspace². Of course self-similarity is a stronger property stemming from a stronger requirement on the cipher.

² It may be that a non-trivial commuting matrix leads only to trivial invariant subspaces, as evidenced by the 2×2 binary matrix with rows [01] and [11]. However if M is involutive, $\ker(M + Id)$ is at least half of the space.

In our applications, M will be involutive, so we focus on the case $i = 1$. In the remainder, whenever two plaintext blocks (or ciphertext blocks, or inner states, or keys) satisfy $x_2 = M(x_1)$, we say that they are *related*. If a plaintext block (or ciphertext block, or inner state, or key) is related to itself, we say that it is *self-related*. A *weak key* is a self-related key. In short, our attack states that weak keys map self-related plaintexts to self-related ciphertexts; while related keys map pairs of related plaintexts to pairs of related ciphertexts.

3.2 S-box-Independent Setting

We now focus on the case where the cipher is a substitution-permutation network (SPN), whose round function F_i consists of an S-box layer with identical S-boxes, a linear map L , addition of a round constant $C(i)$, and addition of a fixed key K . From the invariant subspace (resp. self-similarity) perspective, we are interested in subspaces (resp. linear maps) that traverse (resp. commute with) each of these components.

It is quite apparent that the main roadblock is the non-linear S-box layer. However even in a generic setting where we do not take into account a particular choice of S-box, any permutation of the S-box inputs will commute with the S-box layer (due to S-boxes being identical). Thus we restrict our attention to permutations of S-box inputs rather than general linear maps.

In terms of invariant subspaces, this corresponds to subspaces containing those vectors whose coordinates belonging to the same cycle in the permutation are equal; that is, subspaces that only require S-box inputs to be equal to some other input, or independent. We call such spaces *equality spaces*. Note that these are vector subspaces and no longer affine subspaces. Our strongest attacks actually occur in this setting.

As for constant and key addition, asking that their addition commutes with M amounts to asking that they belong to $\ker(M + Id)$. Now it remains to find permutations that commute with the linear layer. An efficient algorithm to do so is provided in the ePrint version of this work [23]. The invariant subspace variant seems more difficult, as we do not know an algorithm able to efficiently enumerate equality spaces that traverse a linear map.

3.3 Key Recovery

The self-similarity attack above breaks plaintext confidentiality. In addition, if the commuting permutation P is involutive (as will be the case in our applications), efficient key recovery may be possible. In short, the part of the key corresponding to fixed points of the permutation can be guessed independently of the rest.

Intuitively, this is because if two self-related inner states differ only outside the fixed points of the permutation P , this difference will never be propagated to the fixed points of P . This is clear for the S-box layer (because the permutation operates on entire S-box inputs), but also holds for the linear layer. A general

statement and proof are provided in the ePrint version of this work [23]. In fact the proof also encompasses the case where the S-box layer is partial.

What we show is that the cipher contains an embedded subcipher operating on the fixed points F of the permutation: we can project self-related plaintexts and ciphertexts on F and obtain a well-defined map. Note that this embedded subcipher may lend itself to further attacks; this is a direction we have not investigated, as we believe ciphers are sufficiently broken at that point.

3.4 Invariant Permutation Attack on LS-Designs

Notation. In the S-box-independent setting of Section 3.2, for an LS-design, a permutation of S-box inputs is simply a permutation of the columns of the state. Let us write P for such a permutation. We always denote by the lowercase p its effect on a single row. Thus, P is the application of p on each row of the state. We identify p with the corresponding $c \times c$ permutation matrix. We adopt notations from Section 2.2.

The particular structure of LS-designs means that P commutes with L iff p commutes with ℓ . This is still a strong requirement, but we expect the L-box of an LS-design to have some structure in order to provide a good branch number, especially if it is involutive. In the case of Robin for instance, the linear layer is built from a Reed-Muller code and provides plenty of structure. Applied to LS-designs, Lemma 2 becomes:

Lemma 3. *For an LS-design, assume there exists a permutation P with the following properties:*

- P commutes with L .
- $P(C(i)) = C(i)$ for all round indices i .

Then for any plaintext message m :

$$\text{Enc}_{P(K)}(P(m)) = P(\text{Enc}_K(m))$$

In particular, if $K = P(K)$:

$$\text{Enc}_K(P(m)) = P(\text{Enc}_K(m))$$

Note that the identity permutation trivially satisfies the above requirements. Hereafter we always assume P is non-trivial. If $\text{ncycles}(p)$ is the number of cycles of p , weak keys form a proportion $2^{-r \cdot (c - \text{ncycles}(p))}$ of all keys (namely, those keys whose columns are equal on each cycle of p).

Key Recovery. The previous attack breaks plaintext confidentiality. In addition, when P is involutive, efficient key recovery is possible, as announced in Section 3.3. A general statement and proof are provided in the ePrint version of this work [23].

It may still be worthwhile to provide a simpler statement dedicated to LS-designs. This is what we propose below.

Lemma 4. *Consider an LS-design, and assume there exists a permutation P with the same requirements as in Lemma 3. Also assume that P is an involution. Consider a weak key $K = P(K)$. Denote by F the set of fixed points of P .*

Take any self-related plaintext $m = P(m)$. Then the value of the ciphertext $\text{Enc}_K(m)$ on the columns in F only depends on the value of m and K on the same columns.

Proof. Since P is an involution, all of its cycles have length 1 or 2. Hence we can partition the columns of the state into three subsets F, A, B , such that P is the identity on F , and maps A and B into each other. Take any self-related message m that is zero on F . Then the linear layer maps m to a self-related state $L(m)$ that is also zero on F . To see this, write $m = m_A + m_B$, where m_A is equal to m on A , and zero elsewhere, and likewise m_B is equal to m on B and zero elsewhere. Then $P(m_A) = m_B$, hence $P(L(m_A)) = L(m_B)$ by commutativity of P and L . Since P is the identity on F , this implies that $L(m_A) + L(m_B)$ is zero on F , so $L(m)$ is zero on F .

Thus, if $m = P(m)$ is zero on F , so is $L(m)$. By linearity, this implies that if m_1 and m_2 are self-related and equal on F , then so are $L(m_1)$ and $L(m_2)$. Thus, the property that two self-related states are equal on F goes through the linear layer. This property automatically goes through the S-box layer since it is column-wise. Since the same key and round constants are added to both sides, they have no impact. Hence this property goes through the whole cipher. \square

As a direct consequence, the value of the key on the columns corresponding to fixed points of P can be guessed independently of the rest of the key by using any self-related plaintext. In addition, the embedded subcipher is a smaller LS-design, and may lend itself to further attacks. As a side note, both this lemma and the previous one also show that the cipher is malleable in a strong sense.

Permutation Characteristic. Instead of considering only permutations P commuting with L , we can naturally look for pairs of permutations (P, Q) such that $L \cdot P = Q \cdot L$. We denote this by $P \rightarrow Q$, representing the fact that if two inner states are related by P before the linear layer, then after the linear layer they are related by Q .

From there we can hope to build a form of characteristic $P_0 \rightarrow P_1 \rightarrow P_2 \rightarrow P_3 \rightarrow \dots$. The commutative case in the previous section corresponds to $P \rightarrow P$. Note that the set of permutations P such that $Q = L \cdot P \cdot L^{-1}$ is a permutation forms a group. Also note that if L is involutive, $P \rightarrow Q$ is equivalent to $Q \rightarrow P$: indeed $L \cdot P = Q \cdot L$ implies $P^{-1} \cdot L = L \cdot Q^{-1}$, implies $L \cdot Q = P \cdot L$: hence any transition $P \rightarrow Q$ yields an iterative characteristic of length at most 2.

A particularly interesting case occurs whenever $P \rightarrow P^\alpha$ for some $\alpha \neq 0$. Indeed, in that case we automatically have a cyclic characteristic $P \rightarrow P^\alpha \rightarrow P^{\alpha^2} \rightarrow \dots \rightarrow P^{\alpha^i} = P$. Moreover the attack from Lemma 3 goes through with exactly the same requirements on the key and round constants (namely they are self-related by P).

Application to Robin. Applying our attack to Robin amounts to finding a permutation p commuting with the matrix ℓ in Fig. 1, such that P leaves all round constants $C(i)$ invariant. More generally, as pointed out just above, we can actually look at transitions $P \rightarrow Q$, i.e. permutations p, q such that $\ell \cdot p = q \cdot \ell$. It turns out there are 720 such transitions, and all of them are of the form $P \rightarrow P^{-1}$. Moreover 76 of these permutations are involutive, and hence commute with L .

Recall that the round constants of Robin are defined as $C(i) = \ell(i)$ on the first row, and zero on the others, for $1 \leq i \leq t$. Hence we want $p(\ell(i)) = \ell(i)$, which amounts to $p(i) = i$ by commutativity. Since i ranges from 1 to 16, what we are looking for is simply permutations leaving the first 5 columns fixed. It turns out there exists exactly one such permutation, namely the involutive permutation P_0 switching columns 8, 9, 10, 11 respectively with columns 12, 13, 14, 15. Looking at Fig. 1, one can indeed see that permuting the rows and columns of the matrix of ℓ by p_0 leaves the matrix invariant, which is the same as saying p_0 commutes with ℓ .

With P_0 , weak keys are simply keys whose last four columns are equal to the previous four. In particular the proportion of weak keys is 2^{-32} . Furthermore P_0 leaves the first 8 columns fixed, so Lemma 4 shows that for self-related plaintexts, the first 8 columns of ciphertexts only depend on the first 8 columns of plaintext and key. This makes it possible to guess the value of the master key on the first 8 columns independently of the rest of the key. This means 64 bits of the key can be guessed separately; then the remaining 64 bits are symmetric through P_0 , so only 32 bits remain to be guessed. Thus the full key can be recovered in time complexity 2^{64} by encrypting any self-related message. This may yield a few solutions, which can be checked against any other plaintext/ciphertext pair.

Table 2. Permutations p_0, p_1 and p_2 . Fixed points are omitted.

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
p_0									12	13	14	15	8	9	10	11
p_1					8	9	10	15	4	5	6					7
p_2					12	13	14	11				7	4	5	6	

Beside P_0 , two other permutations P_1 and P_2 commuting with L leave the round constants invariant up to the very last round (cf. Table 2). This means related plaintexts are mapped to related inner states after 15 encryption rounds; followed by the final constant addition, S-box layer, L-box layer, and key addition. The final linear layer can be reversed, and the resulting states will agree on pairs of columns transposed by P on which $C(16)$ is equal. In both cases, there is one such pair, so self-related keys with respect to P_1 and P_2 can still be detected easily by encrypting a few self-related plaintexts, reversing the last linear layer, and checking that these two columns agree.

Permutations P_1 and P_2 both leave 8 columns fixed and hence yield an attack with essentially the same properties as P_0 . Actually some key bits can be recovered faster than with P_0 thanks to the one-round differential at the end, but this

involves the symmetric part of the key (that is, outside the fixed points of the permutation) and thus the overall key recovery time is still 2^{64} .

Application to iSCREAM. Recall that iSCREAM and Robin share the same linear layer. Round constants only affect the first eight columns of the state, and so we are looking for permutations commuting with L and leaving the first eight columns unchanged. As a matter of fact, there exists exactly one such permutation, namely the same permutation P_0 as above, which switches the last four columns of the state with the previous four.

Another difference between Robin and iSCREAM is the number of rounds, but that is actually irrelevant for our attack. The last difference is the presence of a tweak in the key schedule. Recall that at odd rounds, $T + K$ is added, while at even rounds, $T \lll_c 1$ is added, where T is a 128-bit tweak. In a chosen-tweak scenario, we can simply set T to zero, or any other value such that T and $T \lll_c 1$ are invariant by P . Then the attack against Robin from the previous section applies to iSCREAM essentially unchanged, with the same consequences.

A small variant of our attack is also possible when using P_0 as the commuting permutation. What we truly want is that $K + T$ and $T \lll_c 1$ should be self-related. This amounts to asking that columns 8, 9, 10, 11 should be equal to columns 12, 13, 14, 15. Since $T \lll_c 1$ is a column-wise shift of T by one column towards the left, this means that columns 9, 10, 11, 12 of T should be equal to columns 13, 14, 15, 0. Note that there is no condition on column 8 of T . As a consequence, for $K + T$ to be self-related for some choice of T , it is enough to ask that columns 9, 10, 11 of K should be equal to columns 13, 14, 15. Indeed in that case, we can fix T to be all-zero, except for column 8 which can take any value: exactly one such choice of T will satisfy that $K + T$ is self-related. Thus we obtain a larger set of weak keys (with ratio 2^{-24}), at the cost of requiring 2^8 chosen-tweak messages in order to detect whether a fixed unknown key is weak.

In addition, some variants of iSCREAM claim related-key security. If two keys are related by P_0 , then our attack applies immediately without any weak key requirement, following the first consequence of Lemma 3. That is, related plaintexts are mapped by the related keys to related ciphertexts. Thus it is easy to check whether a pair of keys is related, and the cipher is broken in a strong sense.

Generalizations of the Permutation Attack. There appears to be a few simple ways in which our attack could be generalized. We discuss them briefly here.

We could consider a probabilistic version of the attack. Instead of requiring $L \cdot P = P \cdot L$, we could consider P 's such that the kernel of $L \cdot P - P \cdot L$ is almost the full space. In the case of Robin or iSCREAM, this would incur a cost at least 2^{-8} per round.

Another natural extension is to consider cases where all round constants are P -invariant except for the last few rounds (or first few rounds). Then our attack goes through most of the encryption process, and eventually yields a differential

attack on the remaining rounds. When encrypting self-related plaintexts, this differential attack turns into an inner differential.

4 Invariant Equality Space Attack

In this section we study invariant subspaces for LS-designs following the S-box-independent setting of Section 3.2. We begin by defining equality spaces, and then present our results on Robin and iSCREAM. On the way, we will recover the invariant subspace detected by our generic algorithm (Section 1.3), and link it to the commuting permutations from the previous section.

4.1 Equality Spaces

As always, we use notations from Section 2.2. We always view n -bit vectors as an $r \times c$ matrix. In Section 3.2, we defined equality spaces in general terms for an SPN; we now provide a more specific definition suited to LS-designs.

Definition 2. *A subspace E of $\{0,1\}^c$ is an equality space iff there exists a partition of $\{0, \dots, c-1\}$ such that E is the set of vectors whose values on coordinates belonging to the same class in the partition are equal.*

The dimension of E is the number of classes of the partition. By E^r we denote the set of n -bit states whose columns belonging to the same class in the partition underlying E are equal. Equivalently, this means that every row of the state belongs to E , hence the notation E^r . By extension we also call E^r an equality space. The point of this definition is that equality spaces are preserved by the S-box layer. The question is to determine which equality spaces are also preserved by the linear layer. That is, we are looking for equality spaces $E \subset \{0,1\}^c$ such that $\ell(E) = E$.

As pointed out in Section 3.1, when a permutation P commutes with L , the equality space defined by the cycles of P is preserved by the linear layer. The idea is that equality spaces preserved by the linear layer do not necessarily stem from a commuting permutation. Conversely, commuting permutations are an interesting special case, since they lead to a stronger property: indeed, when considering equality spaces rather than permutations, we are looking at a property of a single state, and there is no equivalent to the property that distinct related plaintexts are mapped to related ciphertexts; there is also no equivalent to Lemma 4. Meanwhile, Lemma 3 becomes:

Lemma 5. *For an LS-design, assume there exists an equality space E such that:*

- $\ell(E) = E$.
- $C(i) \in E^r$ for all round indexes i .

Then for any key K and plaintext message m :

$$\text{If } K \in E^r \text{ and } m \in E^r \text{ then } \text{Enc}_K(m) \in E^r$$

The lemma trivially holds if E is the full space $\{0,1\}^c$; hereafter we assume this is not the case. Then we have an attack in the weak key setting, where weak keys are keys in E^r . Hence the proportion of weak keys is $2^{-r \cdot (c - \dim(E))}$.

4.2 Variants of the Attack

Essentially the same extensions as in Section 3.4 apply to equality spaces.

Characteristics: if the image $F = L(E)$ of an equality space E is also an equality space, we write $E \rightarrow F$. As with permutations, we can aim to build a characteristic $E_0 \rightarrow E_1 \rightarrow \dots$ over several rounds. Note that the set of equality spaces is closed under intersection, and as a direct consequence, the set of equality spaces E such that $L(E)$ is an equality space is also closed under intersection. If L is involutive, $E \rightarrow F$ is equivalent to $F \rightarrow E$, so characteristics are automatically cyclic.

Probabilistic attack: Instead of asking $F = L(E)$, we can require the dimension of the quotient space $F/L(E)$ to be small.

Differential ending: If all round constants are in the required equality spaces except for the last few (or first few) rounds, it may be possible to cover the remaining rounds with an inner differential characteristic. Indeed in the case $E \rightarrow E$, the equality space attack may be seen as an all-zero inner differential attack.

Differential attack: the entire attack itself may be transposed into the differential world, at the expense of becoming probabilistic. Consider a state difference living in E^r with $L(E) = E$. Then at each round, require that the S-box layer preserves this equality; that is, the output of some S-boxes which receive equal input, should remain equal. Note that if E stems from an involutive permutation commuting with L , the columns corresponding to fixed points of the permutation can be set to a zero difference: this will be preserved by the linear layer (cf. the proof of Lemma 4). This attack avoids key and round constant requirements, at the cost of much lower probability, and hence high data requirements. In practice this would lead to a weaker attack against Robin than truncated differential product trails in the original article [15], because the branch number is 8 and the non-fixed points of P_0 involve 8 S-boxes.

4.3 Application to Robin and iSCREAM

Since Robin and iSCREAM share the same linear layer L , we consider them together. We enumerated all equality spaces E such that $L(E)$ is an equality space (there are around 2^{33} partitions of 16 elements, so this is feasible), and analyzed the results.

Our first observation is that there are many more *well-behaved* equality spaces E (in the sense that $L(E)$ is also an equality space), than *well-behaved* permutations P (in the sense that $Q = L \cdot P \cdot L^{-1}$ is also a permutation). Namely, there are 720 well-behaved permutations for L , while there are 30162 well-behaved spaces of dimension 8 or more. Even if we remove from this list spaces that are an intersection of larger well-behaved spaces (and thus could have a chance of indirectly resulting from well-behaved permutations), 7746 well-behaved spaces remain.

Recall that L is involutive, so any transition $E \rightarrow F$ (i.e. $L(E) = F$ with E and F two equality spaces) yields a cyclic characteristic $E \rightarrow F \rightarrow E$. Hence all

well-behaved spaces belong to cycles of length 1 or 2. The aforementioned 7746 intersection-reduced well-behaved spaces of dimension at least 8 form 2506 cycles of length 1 (that is, $E \rightarrow E$) and 2620 cycles of length 2 (that is $E \rightarrow F \rightarrow E$). Thus equality spaces offer considerably more potential attacks, depending on round constants.

However, all equality spaces compatible with actual round constants for Robin minus the last round, and hence directly usable in an attack, stem from commuting permutations. There exist four such spaces: three of them correspond to permutations P_0 , P_1 and P_2 from Table 2, and the last one is a space of dimension 8 resulting from the composition of any two of the previous permutations (any combination yields the same permutation or its inverse). As for iSCREAM, the only well-behaved space compatible with round constants is the one resulting from P_0 . Thus, our previous attack is not improved. Moreover, the largest well-behaved spaces have dimension 12 and all stem from involutive permutations (there are 15 of them). The largest well-behaved equality spaces not stemming from a well-behaved permutation have dimension 10. This may be interpreted to mean that the strongest phenomenon is due to commuting permutations.

Thus for both Robin and iSCREAM, the equality space induced by P_0 is the only equality space that goes through the whole cipher, including the last round. This space has dimension 96 over \mathbb{F}_2 , and it is the invariant subspace automatically discovered by the generic algorithm from Section 1.

4.4 A note on Fantomas and SCREAM

The matrix L of Fantomas is a permutation of the lines and columns of the matrix of Robin. As a consequence, they have the same number of well-behaved permutations and spaces. However we found no cycle among well-behaved spaces of Fantomas of dimension 6 or more (lower dimensions would yield very weak attacks); and no characteristic of length more than 2. Hence Fantomas seems safe from this attack.

The same is true for SCREAM. However, it is worth noting that there exists no well-behaved permutation for the matrix of SCREAM, while we found 5404 well-behaved spaces of dimension 8 or more.

5 A Second Invariant Subspace Attack on LS-Designs

In this section we present a different invariant subspace attack on LS-designs, which may be regarded as a form of dual of the previous attack. This attack does not stem from an underlying permutation; nor does it have an equivalent for Zorro. Thus, this section is specific to LS-designs, and takes advantage of their particular structure: namely, the fact that LS-designs not only rely on a layer of identical S-boxes, but also on a layer of identical L-boxes.

Now that we have understood the invariant subspace discovered by our generic algorithm as being an equality space, i.e. a space that is automatically preserved by the S-box layer, it is natural to ask if something similar can be

done with the L-box layer. That is, we are now going to look for a property that is automatically preserved by the L-box layer.

This gives us more freedom, since we can leverage linearity. Essentially, if all columns of the state live in the same linear subspace, this will remain true after the linear layer (in the ePrint version of this work [23], we prove that this is in fact the most general property generically preserved by the linear layer); whereas in the previous case, we were limited to equality spaces. Beside this difference, the attack is essentially a dual version of the previous one, reversing the roles of the L-box and S-box layers.

5.1 Description of the Attack

In the previous attack, we searched for equality spaces $E \subset \{0, 1\}^c$ on the rows of the state such that $\ell(E) = E$. Instead, we are now interested in general linear subspaces $A \subset \{0, 1\}^r$ on the columns of the state such that $s(A) = A$. Once again, if A is a linear space on the columns (or one of its cosets), we denote by A^c the set of states whose columns all belong to A .

The core of the attack is the following: assume $s(A) = A$ for some linear space A . If the inner state lies in A^c , this will remain true after the S-box layer. Moreover, this property is automatically preserved by the linear layer. Indeed, the linear layer of an LS-design is not truly “line-wise”: precisely because the same linear map is applied to each row, the linear layer may be seen as directly adding together column vectors. From this point of view, it becomes clear that if all columns lie in the same linear space A , this remains true after the linear layer.

Thus we are still within the invariant subspace framework, and follow the corresponding strategy: we choose A such that all round constants belong to A^c , and we consider a weak key scenario by requiring that the key also lie in A^c . If these requirements are fulfilled, plaintexts in A^c are mapped to ciphertexts in A^c .

More generally, we can consider cosets of linear spaces (i.e. affine spaces) rather than just linear spaces: indeed, as long as each coordinate at the output of ℓ is the sum of an odd number of coordinates at the input, the linear layer still preserves the property that all columns belong to a fixed coset. The following lemma sums up the attack.

Lemma 6. *Let u, v, w be r -bit vectors, and A be a linear subspace of r -bit vectors. Assume the following conditions hold:*

- *The S-box s maps all vectors in $u + A$ to vectors in $v + A$.*
- *Either $v = 0$ or all rows of the matrix of ℓ have an odd number of 1's.*
- *The columns of all round constants are in $w + A$.*
- *The columns of the key are in $(u + v + w) + A$.*

Then any plaintext in $(u + w) + A$ is encrypted into a ciphertext in $(u + w) + A$ (and conversely).

Weak keys are keys in $(u + v + w) + E$. This means a proportion $2^{-c \cdot (r - \dim(A))}$ of keys is weak.

5.2 Application to Robin and iSCREAM

In the case of Robin, the second condition in Lemma 6 is automatically true. In order to satisfy the third condition (round constants), since round constants only affect the first row of the state, we require that the r -bit vector denoted by 1 , with 1 on the first row and 0 elsewhere, belongs to E . To instantiate the attack, it remains to look for affine spaces whose direction contains the vector 1 , that are mapped by the S-box to affine spaces with the same direction.

It turns out the largest such spaces have dimension 3, and are mapped into themselves. We list all six choices in Table 3. Since these spaces have dimension 3, and the state has 8 rows and 16 columns, a proportion $2^{-16 \cdot 5} = 2^{-80}$ of keys are weak. This means our attack is considerably weaker than the first one against Robin. By comparison, a generic multi-target time-memory trade-off with 2^{48} memory would lead to key recovery for the same proportion of keys. Of course our attack requires no memory or table lookup.

Table 3. Six affine spaces of dimension 3 invariant through s

Values in A								Dir(A)		
00	01	26	27	84	85	a2	a3	01	26	84
18	19	7c	7d	9e	9f	fa	fb	01	64	86
28	29	32	33	8a	8b	90	91	01	1a	a2
3c	3d	5e	5f	b2	b3	d0	d1	01	62	8e
44	45	66	67	c8	c9	ea	eb	01	22	8c
4e	4f	54	55	6c	6d	76	77	01	1a	22

We now turn to iSCREAM. Recall that its S-box is the same as that of Robin, and round constants still only affect the first row of the state. We want both $K + T$ and T to live in the same coset, so we require T to lie in $(u + v + w + A)^c$, and K to lie in A^c . In our actual attack we have $u = v$ and $w = 0$ so in the end, we can set the tweak to zero (or any value in A^c), and the attack goes through with the same parameters as before.

5.3 Taking Advantage of the iSCREAM Tweak Schedule

In the case of iSCREAM, it is possible to leverage the tweak schedule to create a trade-off between the ratio of weak keys and the number of chosen-tweak messages required to detect a weak key. To simplify notations, we explain this technique using vector spaces; it extends to their cosets in a straightforward manner. Assume we have two vector spaces A and B with $S(A) = B$. As before, we assume $1 \in A$ and $1 \in B$ so that round constants belong to A^c and B^c . Since S is involutive, we have $S(B) = A$, so $A \rightarrow B \rightarrow A$ is a characteristic for the the S-box.

In order for this characteristic to traverse encryption, we need $K + T \in A^c$, and $T \lll_c 1 \in B^c$, which is equivalent to $T \in B^c$. For this it is enough to ask $K \in A^c + B^c = (A + B)^c$. Indeed in that case, write $K = K_A + K_B$ with

$K_A \in A^c$ and $K_B \in B^c$. Then for $T = K_B$, we have $K + T \in A^c$ and $T \in B^c$, which is precisely what we want. Of course the key is unknown to the attacker, so she cannot compute T in this way. Instead, she can try every value in the supplementary space of A^c in $(A+B)^c$ (which is smaller than B^c , if only because $1 \in A \cap B$). For exactly one such value of the tweak, every plaintext in A^c will be encrypted to a ciphertext in B^c .

Now the question is to find two spaces A and B as above. Actually we look for cosets of linear spaces with the same properties, since the linear layer of iSCREAM also preserves these cosets. In summary, we look for affine spaces $u + A \neq v + B$ such that $S(u + A) = v + B$, and 1 belongs to $A \cap B$.

It turns out the largest such spaces have dimension 3. There are 11 such spaces (counting only 1 for $u + A \rightarrow v + B$ and $v + B \rightarrow u + A$), listed in Appendix B. Furthermore, 8 of these spaces satisfy $\dim(A + B) = 5$, which is the maximal possible value since 1 belongs to $A \cap B$. Thus $K \in (A + B)^c$ yields a ratio of weak keys of $2^{-c \cdot (r - \dim(A+B))} = 2^{-48}$.

In order to detect whether a key is weak, one needs to encrypt a message for each tweak in the supplementary of A^c in $(A + B)^c$, which is of dimension $2 \cdot c$, hence 2^{32} chosen-tweak messages are required (for a random key and a given choice of the tweak, a false positive has probability only 2^{-80} , and can be discarded by one additional chosen-tweak message). Finally, once a weak key is detected in this way, we know $K + T \in A^c$ for one specific T , hence $K = T + A^c$, so only $2^{c \cdot \dim(A)} = 2^{48}$ possibilities remain for the value of the key.

5.4 Variants of the Attack

It seems natural to consider a probabilistic version of the attack, where instead of requiring that every vector in $u + A$ be mapped by the S-box to a vector in $v + A$, we only require *most* of them to comply. If only x elements in $u + A$ are not mapped to $v + A$, the probability to pass an S-box is $1 - x/2^r$. The cost for each round is then $(1 - x/2^r)^c$. In the case of Robin, there is no A of dimension 4 with $x < 3$, so there does not appear to be an obvious interesting probabilistic version of the attack.

6 Commuting Permutation and Invariant Subspace for Zorro

6.1 Description of Zorro

The block cipher Zorro was introduced at CHES 2013 [14]. Like LS-designs, the design goal is to offer a cipher that can efficiently be made resistant to side-channel attacks through masking [26]. This is achieved by two main techniques: first, a carefully constructed 8-bit S-box; and second, an AES-like structure where S-boxes are only applied on the first row of the state.

The 128-bit state is represented as a 4×4 array of 8-bit cells. The round function applies the following transformations:

- **SubBytes**: A fixed 8-bit S-box is applied to the first row of the state.
- **AddConstant**: At round i , the constants i , i , i and $i \ll 3$ are added to the four cells of the first row (from left to right).
- **ShiftRow**: This step is identical to AES. Row i , counting from zero, is shifted by i cells to the left.
- **MixColumns**: This step is again identical to AES. A fixed 4×4 circulant matrix on \mathbb{F}_{2^8} is applied to each column of the state. The matrix is the same as that of AES.

Four consecutive rounds are called a step. After each step, the 128-bit master key is simply added to the inner state: there is no key schedule. Encryption consists in key addition, followed by 6 steps (24 rounds), each followed by key addition.

6.2 Self-Similarity and Invariant Subspace

We are interested in an S-box-independent commuting linear map, as in Section 3.1. To simplify, we focus on a single round: commuting with every round is a sufficient condition to commute with every step. Thus we are looking for a linear map M acting as a permutation on the S-boxes, and commuting with the linear layer.

Since there are only four S-boxes, there are only 24 choices for the permutation. In fact, because the constant added to the fourth S-box is different from the others, we impose that this S-box should remain fixed by the permutation, leaving only 6 possibilities. In this way, our linear map will automatically commute with both the S-box and constant addition layers.

For each of the 6 permutation choices on the first 3 S-boxes, the set of linear maps behaving as this particular permutation on the first 4 cells, and independently on the other cells, is itself a vector space. Furthermore the commutant of the linear layer is naturally a vector space. Thus, it suffices to intersect these two spaces to find a solution, if it exists.

It turns out there exists exactly one solution, for the permutation swapping the first and third S-boxes, and leaving the other two fixed. This solution is given in Appendix A, together with the resulting invariant subspace. This subspace has dimension 12 over \mathbb{F}_{2^8} , that is, 96 over \mathbb{F}_2 . Hence the proportion of weak keys is 2^{-32} .

In the ePrint version of this work [23], we show how to enumerate all invariant subspaces for Zorro, and deduce that the previous space is in fact the only invariant subspace (in the S-box-independent setting). The strategy used to enumerate spaces extends naturally to any SPN with a partial S-box layer of only a few S-boxes per round.

6.3 Key Recovery

The key recovery strategy from Section 3.4 extends to partial S-box layers such as Zorro. In brief, if an involutive linear map commutes with the components of an SPN, and acts as a permutation on the S-box inputs, part of the key may

be recovered independently of the rest. When the S-box layer is full, i.e. the commuting map is simply a permutation, this part of the key corresponds to the fixed points of the permutation. When the S-box layer is partial, and hence the commuting map M is not fully a permutation, the role of the non-fixed points is essentially played by $I = \text{Im}(M + Id)$. A formal statement and proof are provided in the ePrint version of this work [23].

The consequence for Zorro is that once a key is recognized as weak, 64 bits of the key can be guessed independently of the rest using one chosen plaintext (any self-related plaintext). Indeed, the part of the key in I only influences the part of the ciphertext in I . After these 64 bits have been recovered by brute force, only 32 bits remain to be guessed, due to the key being weak. Thus key recovery requires only one chosen plaintext and a time complexity of 2^{64} offline encryptions.

7 Conclusion

In this article, we present a unified cryptanalysis of several ciphers based on invariant properties traversing the cipher under certain conditions, while providing generic tools for this type of attack. Our attacks are able to break lightweight ciphers Robin, iSCREAM and Zorro in a practical setting.

Our attacks from sections 4 and 5 are quite similar in principle. The state of an LS-design is a rectangular array. A fixed line-wise operation is performed in each direction. Each attack looks for properties of the inner state that would be *structurally* preserved in one direction (in the sense that this does not depend on the specificities of the S-box or linear layer), that would happen to also be preserved in the other (this time due to the particular choices of S or L).

In the case where the generic direction is linear, any linear space is preserved, and under some conditions any coset; if it is nonlinear, only equality spaces are preserved. In the ePrint version of this work[23], we prove that these are in fact the most general properties structurally preserved in each direction, so our attacks fully realize the program outlined in the previous paragraph. It remains an open question whether a similar attack could in some way combine information from both directions; that is, neither direction would preserve the invariant property in a fully generic way.

Concerning our first attack on LS-designs from sections 3 and 4 (encompassing both invariant permutations and invariant equality spaces), the structure of the linear map is a key component. It seems unlikely that the attack could succeed in cases where the linear layer is not involutive. Indeed, as shown by the matrices of SCREAM and Fantomas, even in the presence of a large number of well-behaved equality spaces, it appears that iterative characteristics do not occur by accident. By contrast, if the linear layer is involutive, any well-behaved equality space (or permutation) yields a cyclic characteristic of length at most 2;

and indeed, in the case of Robin and iSCREAM, thousands of iterative characteristics exist. Of course, the matrix of Robin and iSCREAM has much more structure than a generic involutive matrix.

It is quite striking that exactly the same attacks exist on Zorro, despite its quite different structure (byte-oriented vs. bit-oriented, partial S-box layer vs. full, AES-like vs. somewhat SERPENT-like). It is worth noting however that both ciphers attempt precisely the same goal, namely to offer efficient masked implementations. As a result both reduce non-linear operations to a minimum per round, while giving more weight to the linear layer; LS-designs achieve this by parallelizing the S-box through bit slicing; Zorro by resorting to a partial S-box layer. In both cases the contribution of the non-linear layer is very structured with respect to the linear layer; this, together with the minimal key schedule and simple round constants leads to our attacks.

We note that all our attacks can be prevented by a careful choice of round constants. One needs only ensure that no weaker (such as probabilistic or differential) version of the attack is left behind. This is particularly true when claiming related-key security (as in iSCREAM), since in this setting our attacks do not require weak keys, and hence weaker probabilistic versions are quite relevant.

Going back to the generic algorithm used to find the attacks, an interesting open problem is to specialize it to SPN structures, hoping to achieve better time complexity. In particular, it may be worthwhile to find an algorithm that is able to enumerate all invariant subspaces through a layer of n S-boxes, given n and the S-box. With improvements in time complexity, it may become possible to entirely disprove the existence of invariant subspaces for some SPNs.

Finally, we hope our analysis contributes some insight for the design of future ciphers with minimal key schedules and the choice of round constants in cryptographic permutations.

Acknowledgments. The authors would like to thank Henri Gilbert for many fruitful discussions related to the attacks presented in this article.

A Commuting Linear Map and Invariant Subspace for Zorro

The commuting linear map M is represented as a 16×16 matrix over \mathbb{F}_{2^8} , using the AES representation of \mathbb{F}_{2^8} as $\mathbb{F}_2[x]/(x^8 + x^4 + x^3 + x + 1)$.

$$\begin{bmatrix}
 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
 0 & 0 & 0 & 0 & 34 & 101 & 35 & 101 & 50 & 249 & 50 & 249 & 249 & 116 & 249 & 116 \\
 0 & 0 & 0 & 0 & 101 & 35 & 101 & 34 & 249 & 50 & 249 & 50 & 116 & 249 & 116 & 249 \\
 0 & 0 & 0 & 0 & 35 & 101 & 34 & 101 & 50 & 249 & 50 & 249 & 249 & 116 & 249 & 116 \\
 0 & 0 & 0 & 0 & 101 & 34 & 101 & 35 & 249 & 50 & 249 & 50 & 116 & 249 & 116 & 249 \\
 0 & 0 & 0 & 0 & 17 & 86 & 17 & 86 & 1 & 0 & 0 & 0 & 249 & 50 & 249 & 50 \\
 0 & 0 & 0 & 0 & 86 & 17 & 86 & 17 & 0 & 0 & 0 & 1 & 50 & 249 & 50 & 249 \\
 0 & 0 & 0 & 0 & 17 & 86 & 17 & 86 & 0 & 0 & 1 & 0 & 249 & 50 & 249 & 50 \\
 0 & 0 & 0 & 0 & 86 & 17 & 86 & 17 & 0 & 1 & 0 & 0 & 50 & 249 & 50 & 249 \\
 0 & 0 & 0 & 0 & 51 & 190 & 51 & 190 & 86 & 17 & 86 & 17 & 35 & 101 & 34 & 101 \\
 0 & 0 & 0 & 0 & 190 & 51 & 190 & 51 & 17 & 86 & 17 & 86 & 101 & 34 & 101 & 35 \\
 0 & 0 & 0 & 0 & 51 & 190 & 51 & 190 & 86 & 17 & 86 & 17 & 34 & 101 & 35 & 101 \\
 0 & 0 & 0 & 0 & 190 & 51 & 190 & 51 & 17 & 86 & 17 & 86 & 101 & 35 & 101 & 34
 \end{bmatrix}$$

The invariant subspace $\ker(M + Id)$ is generated by the following 12 row vectors, in the same representation.

$$\begin{aligned}
 &(1 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 &(0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 &(0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 &(0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 38 \ 0 \ 0 \ 159 \ 0) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 3) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 38 \ 0 \ 0 \ 159 \ 0) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 3) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 79 \ 0 \ 0 \ 38 \ 1) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 79 \ 0 \ 0 \ 38 \ 1) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1 \ 0) \\
 &(0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 1)
 \end{aligned}$$

B Well-Behaved Affine Spaces for the Robin and iSCREAM S-Box

Only spaces whose direction contains 1 are listed.

Values in $u + A$			Basis of A			Values in $v + B = S(u + A)$			Basis of B			$\dim(A + B)$										
00	01	26	27	84	85	a2	a3	01	26	84	a3	01	26	84	3							
18	19	7c	7d	9e	9f	fa	fb	01	64	86	fa	fb	01	64	86	3						
28	29	32	33	8a	8b	90	91	01	1a	a2	32	33	28	29	01	1a	a2	3				
3c	3d	5e	5f	b2	b3	d0	d1	01	62	8e	d0	d1	3c	3d	5e	5f	01	62	8e	3		
44	45	66	67	c8	c9	ea	eb	01	22	8c	ea	eb	44	45	66	67	01	22	8c	3		
4e	4f	54	55	6c	6d	76	77	01	1a	22	76	77	6c	6d	54	55	4f	4e	01	1a	22	3
28	29	32	33	6c	6d	76	77	01	1a	44	8a	8b	54	55	4e	4f	01	1a	c4	4		
28	29	32	33	4e	4f	54	55	01	1a	66	8a	8b	76	77	6c	6d	01	1a	e6	4		
2e	2f	38	39	8c	8d	9a	9b	01	16	a2	6e	6f	63	62	cd	cc	c1	c0	01	0c	a2	4
08	09	2e	2f	8c	8d	aa	ab	01	26	84	4d	4c	6f	6e	c1	c0	e3	e2	01	22	8c	5
08	09	38	39	9a	9b	aa	ab	01	30	92	4d	4c	63	62	cd	cc	e3	e2	01	2e	80	5
0a	0b	12	13	c6	c7	de	df	01	18	cc	2c	2d	36	37	68	69	72	73	01	1a	44	5
0e	0f	16	17	c2	c3	da	db	01	18	cc	bd	bc	ad	ac	f1	f0	e1	e0	01	10	4c	5
20	21	3e	3f	86	87	98	99	01	1e	a6	59	58	5d	5c	e9	e8	ed	ec	01	04	b0	5
22	23	34	35	80	81	96	97	01	16	a2	78	79	74	75	d8	d9	d4	d5	01	0c	a0	5
24	25	3a	3b	82	83	9c	9d	01	1e	a6	47	46	43	42	fd	fc	f9	f8	01	04	ba	5
4a	4b	50	51	8e	8f	94	95	01	1a	c4	e4	e5	f4	f5	a8	a9	b8	b9	01	10	4c	5

References

1. CAESAR- Competition for Authenticated Encryption: Security, Applicability, and Robustness. General secretary Daniel J. Bernstein (2013). <http://competitions.cr.yp.to/caesar.html>
2. Albrecht, M.R., Driessen, B., Kavun, E.B., Leander, G., Paar, C., Yalçın, T.: Block ciphers – focus on the linear layer (feat. PRIDE). In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 57–76. Springer, Heidelberg (2014)
3. Bar-On, A., Dinur, I., Dunkelman, O., Lallemand, V., Tsaban, B.: Improved analysis of Zorro-like ciphers. Cryptology ePrint Archive, Report 2014/228 (2014). <http://eprint.iacr.org/>
4. Barkan, E., Biham, E.: In how many ways can you write rijndael? In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 160–175. Springer, Heidelberg (2002)
5. 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)
6. Bogdanov, A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M.J.B., Seurin, Y., Vikkelsoe, C.: PRESENT: an ultra-lightweight block cipher. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 450–466. Springer, Heidelberg (2007)
7. Borghoff, J., et al.: PRINCE – a low-latency block cipher for pervasive computing applications. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 208–225. Springer, Heidelberg (2012)
8. Bouillaguet, C., Dunkelman, O., Leurent, G., Fouque, P.-A.: Another look at complementation properties. In: Hong, S., Iwata, T. (eds.) FSE 2010. LNCS, vol. 6147, pp. 347–364. Springer, Heidelberg (2010)
9. Bulygin, S., Walter, M., Buchmann, J.: Many weak keys for PRINTCIPHER: fast key recovery and countermeasures. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 189–206. Springer, Heidelberg (2013)
10. 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)
11. Chaum, D., Evertse, J.-H.: Cryptanalysis of des with a reduced number of rounds. In: Williams, H.C. (ed.) CRYPTO 1985. LNCS, vol. 218, pp. 192–211. Springer, Heidelberg (1986)
12. Daemen, J., Peeters, M., Van Assche, G., Rijmen, V.: NESSIE proposal: NOEKEON (2000). <http://gro.noekeon.org/>
13. Evertse, J.-H.: Linear structures in block ciphers. In: Price, W.L., Chaum, D. (eds.) EUROCRYPT 1987. LNCS, vol. 304, pp. 249–266. Springer, Heidelberg (1988)
14. Gérard, B., Grosso, V., Naya-Plasencia, M., Standaert, F.-X.: Block ciphers that are easier to mask: how far can we go? In: Bertoni, G., Coron, J.-S. (eds.) CHES 2013. LNCS, vol. 8086, pp. 383–399. Springer, Heidelberg (2013)
15. Grosso, V., Leurent, G., Standaert, F.-X., Varici, K.: LS-designs: bitslice encryption for efficient masked software implementations. To appear in the Proceedings of FSE 2014 (2014). <http://www.uclouvain.be/crypto/people/show/382>
16. Grosso, V., Leurent, G., Standaert, F.-X., Varici, K., Durvaux, F., Gaspar, L., Kerckhof, S.: Addendum to the CAESAR submission for SCREAM and iSCREAM. Posted on the official CAESAR submission list (2014). <http://competitions.cr.yp.to/round1/scream-ordering.txt>

17. Grosso, V., Leurent, G., Standaert, F.-X., Varici, K., Durvaux, F., Gaspar, L., Kerckhof, S.: CAESAR candidate SCREAM. Presentation by Gaëtan Leurent at DIAC 2014 (2014). <http://2014.diac.cr.yp.to/slides/leurent-scream.pdf>
18. Grosso, V., Leurent, G., Standaert, F.-X., Varici, K., Durvaux, F., Gaspar, L., Kerckhof, S.: SCREAM & iSCREAM. Entry in the CAESAR competition [1] (2014). <http://competitions.cr.yp.to/round1/screamv1.pdf>
19. Guo, J., Nikolić, I., Peyrin, T., Wang, L.: Cryptanalysis of Zorro. Cryptology ePrint Archive, Report 2013/713 (2013). <http://eprint.iacr.org/>
20. 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)
21. 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)
22. Leander, G., Abdelraheem, M.A., AlKhzaimi, H., Zenner, E.: A cryptanalysis of PRINTCIPHER: the invariant subspace attack. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 206–221. Springer, Heidelberg (2011)
23. Leander, G., Minaud, B., Rønjom, S.: A generic approach to invariant subspace attacks: Cryptanalysis of Robin, iSCREAM and Zorro. Cryptology ePrint Archive, Report 2015/068 (2015). <http://eprint.iacr.org/2015/068>
24. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable block ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer, Heidelberg (2002)
25. Murphy, S.: An analysis of SAFER. Journal of Cryptology 11(4), 235–251 (1998)
26. Prouff, E., Rivain, M.: Masking against side-channel attacks: a formal security proof. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 142–159. Springer, Heidelberg (2013)
27. Rasoolzadeh, S., Ahmadian, Z., Salmasizadeh, M., Aref, M.R.: Total break of Zorro using linear and differential attacks. Cryptology ePrint Archive, Report 2014/220 (2014). <http://eprint.iacr.org/>
28. Reeds, J.A., Manferdelli, J.L.: Des has no per round linear factors. In: Blakely, G.R., Chaum, D. (eds.) CRYPTO 1984. LNCS, vol. 196, pp. 377–389. Springer, Heidelberg (1985)
29. Sim, S.M., Wang, L.: Practical forgery attacks on SCREAM and iSCREAM (2014). Posted on the crypto competitions mailing list at <https://groups.google.com/d/forum/crypto-competitions>, report available at <https://www1.spms.ntu.edu.sg/~syllab/m/images/b/b3/ForgeryAttackOnSCREAM.pdf>
30. Soleimany, H.: Probabilistic slide cryptanalysis and its applications to LED-64 and Zorro. To appear in the Proceedings of FSE 2014 (2014). <http://research.ics.aalto.fi/publications/bibdb2014/pdf/fse2014.pdf>
31. Wang, Y., Wu, W., Guo, Z., Yu, X.: Differential cryptanalysis and linear distinguisher of full-round zorro. In: Boureau, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 308–323. Springer, Heidelberg (2014)