

# Tweaks and Keys for Block Ciphers: The TWEAKEY Framework

Jérémy Jean, Ivica Nikolić, and Thomas Peyrin

Division of Mathematical Sciences, School of Physical and Mathematical Science,  
Nanyang Technological University, Singapore  
{JJean, INikolic, Thomas.Peyrin}@ntu.edu.sg

**Abstract.** We propose the TWEAKEY framework with goal to unify the design of tweakable block ciphers and of block ciphers resistant to related-key attacks. Our framework is simple, extends the key-alternating construction, and allows to build a primitive with arbitrary tweak and key sizes, given the public round permutation (for instance, the AES round). Increasing the sizes renders the security analysis very difficult and thus we identify a subclass of TWEAKEY, that we name STK, which solves the size issue by the use of finite field multiplications on low hamming weight constants. Overall, this construction allows a significant increase of security of well-known authenticated encryptions mode like  $\Theta$ CB3 from birthday-bound security to full security, where a regular block cipher was used as a black box to build a tweakable block cipher. Our work can also be seen as advances on the topic of secure key schedule design.

**Keywords:** tweak, block cipher, key schedule, authenticated encryption.

## 1 Introduction

Block ciphers are among the most scrutinized cryptographic primitives, used in many constructions as basic secure bricks that ensure data encryption and/or authenticity. In the last few decades, a lot of research has been conducted on this topic, and it is believed that building a secure and efficient block cipher is now a well-understood problem. In particular, designs that allowed to prove their security against classical differential or linear attacks have been a very important step forward, and have been incorporated in the current main worldwide standard AES-128 [34]. This topic is mature and the community has recently been focusing on other directions, such as the possibility to build ciphers dedicated to very constrained environments [9, 12, 23].

The security of the block ciphers, both Feistel and Substitution-Permutation networks, has been well studied when the key is fixed and secret, however, when the attacker is allowed to ask for encryption or decryption with different (and related) keys the situation becomes more complicated. In the past, many published ciphers have been broken in this so-called *related-key model* [4, 5] and it has even been demonstrated that the Advanced Encryption Standard (AES) has

flaws in this model [6, 7]. It is known how to design a cryptographically good permutation composed of several iterated rounds, but when it comes to keying this permutation with subkeys generated by the key schedule, it is hard to ensure that the overall construction remains secure. Most key schedule constructions are ad-hoc, in the sense that the designers came up with a key schedule that is quite different from the internal permutation of the cipher, in a hope that no meaningful structure is created by the interaction of the two components. This is the case of PRESENT [9] or AES [34] ciphers, where the key schedule is purposely made different from the round function. Some key schedules can be very weak but fast and lightweight (like in LED [23], where many rounds are required to ensure security against related-key attacks), while some can be very strong but slow (like in the internal cipher of the WHIRLPOOL hash function [3]). In order to partially ease this task of deriving a good schedule, some automatic tools analyzing the resistance of the ciphers against simple related-key differential attacks have been developed [8, 21, 32].

The Hasty Pudding cipher [37], proposed to the AES competition organized by the NIST, permitted the user to insert an additional input to the classical key and plaintext pair, called *spice* by the designers of this cipher. This extra input  $T$ , later renamed as *tweak*, was supposed to be completely public and to randomize the instance of the block cipher: to different values of  $T$  correspond different and independent families of permutations  $E_K$ . This feature was formalized in 2002 by Liskov et al. [30, 31], who showed that tweakable block ciphers are valuable building blocks if retweaking (changing the tweak value) is less costly than changing its secret key. Tweakable block ciphers (see MERCY [13], for example) found many different utilizations in cryptography, such as disk encryption where each block is ciphered with the same key, but the block index is used as tweak value.

Simple constructions of a tweakable block cipher  $E_K(T, P)$  based on a block cipher  $E_K(P)$ , like XORing the tweak into the key input and/or message input, are not satisfactory. For example, only XORing the tweak into the key input would result in an undesirable property that  $E_K(T, P) = E_{K \oplus X}(T \oplus X, P)$ . Liskov et al. propose instead to use universal hash families for that purpose. The XE and XEX constructions [36] (and the follow-up standard XTS [19]) are based on finite field multiplications in  $GF(2^n)$ , and present the particularity of being efficient if sequential tweaks are used. Nonetheless, even with such feature, these scheme might not be really efficient as the cipher execution is not negligible compared to a finite field multiplication in  $GF(2^n)$  (for example when AES is the internal block cipher and the scheme implementation uses AES-NI instructions). More importantly, *these methods ensure only security up to the birthday-bound (relative to the block cipher size)*. This can be a problem as the main block cipher standards only have 64- or 128-bit block size. Minematsu [33] partially overcomes this limitation by proving beyond birthday-bound security for his design, but at the expense of a very reduced efficiency. The same observation applies to more recent beyond-birthday constructions such as [29, 38]. Overall,

*none of the state-of-the-art block-cipher-based schemes provide both efficiency and beyond birthday-bound security.*

Ad-hoc constructions would be a solution, with the obvious drawback that security proofs regarding the construction would be very hard to obtain. So far, this direction has seen a surprisingly low number of proposals. The NIST SHA-3 competition for hash functions triggered a few, like SKEIN [20] (with its ad-hoc internal tweakable block cipher *Threefish*) and BLAKE2 [2]. It is interesting to note that both are Addition-Rotation-XOR (ARX) functions and thus offer less possibility of proofs with regard to classical differential-linear attacks. As of today, it remains an open problem to design an ad-hoc AES-like tweakable block cipher, which in fact would be very valuable for authenticated encryption as AES-NI instruction sets guarantee extremely fast software implementations. Such a primitive would enable very efficient authenticated encryption with beyond birthday-bound security and proof regarding the mode of operation.

Liskov et al. proposed to separate the roles of the secret key (which provides uncertainty to the adversary) from that of the tweak (which provides independent variability) – interestingly, almost all tweakable block cipher proposals (except *Threefish*) follow this rule. This might be seen as counter intuitive as it is required the tweak input to be somehow more efficient than the key input, but at the same time the security requirement on the tweak seem somehow stronger than on the key, since the attacker can fully control the former (even though tweak-recovery attacks are irrelevant). We argue in this article that, in practice, when one designs a block cipher these two inputs should be considered almost the same, as incorporating a tweak and a secret key shares in fact a lot of common ground, especially for the large family of key-alternating ciphers.

**Our Contributions.** In this article, we bring together key schedule design and tweak input handling for block ciphers in a common framework that we call *TWEAKEY* (Section 2). The idea is to provide a simple framework to design a tweakable block cipher with any key and any tweak sizes. Our construction is very simple and can be seen as a natural extension of key-alternating ciphers: a subtweakey (i.e. a value obtained from the key and the tweak inputs) is incorporated into the internal state at every round of the iterative cipher. One advantage of such a framework is that one can obtain a tweakable single-key block cipher or a double-key length block cipher with the very same primitive.

Not all instances of *TWEAKEY* are secure and, in particular, the case where the key and tweak material is treated exactly the same way does not lead to a secure cipher. However, handling the key and the tweak material the same way would be attractive in terms of performance, implementation, but more importantly it would greatly simplify the security analysis, which is currently the main difficulty designers have to face when constructing an ad-hoc tweakable block cipher. Indeed, the main challenge is to evaluate the appropriate number of rounds required to make the cipher secure – when the tweak size  $t$  and key size  $k$  are too large this problem becomes infeasible. We propose a solution in Section 3 and we give a subclass of *TWEAKEY* for AES-like ciphers, named *STK* (for Superposition *TWEAKEY*), where the key and the tweak materials are treated almost the same way – the small

difference between the linear key and the tweak schedules is sufficient to remove the aforementioned weakness. Due to the structure of STK, the security analysis is rendered much easier, and the number of rounds can be kept small. The STK construction leads to promising performances: in [24], a complete 128-bit tweak 128-bit key 128-bit block cipher proposal `Deoxys-BC` based on the AES round function is proposed as an instance of the STK construction. It is faster and more lightweight than other tentatives to build a tweakable block cipher from AES-128. When used in `OCB3` [28] authenticated encryption, `Deoxys-BC` runs at about 1.3 c/B on the latest Intel processors. This has to be compared to `OCB3`, which runs at 0.7-0.88 c/B when instantiated with AES-128, but only ensures birthday-bound security. Alternatively, `Deoxys-BC` could be a replacement for AES-256, which has related-key issues as shown in [7]. The STK construction offers a very lightweight tweakkey schedule (only composed of a substitution of bits), that even allows the key to be hardwired in hardware implementations. Similarly, one can mention `Joltik-BC`: a lightweight instance of the STK construction as a 64-bit tweak 64/128-bit key 64-bit tweakable block cipher.

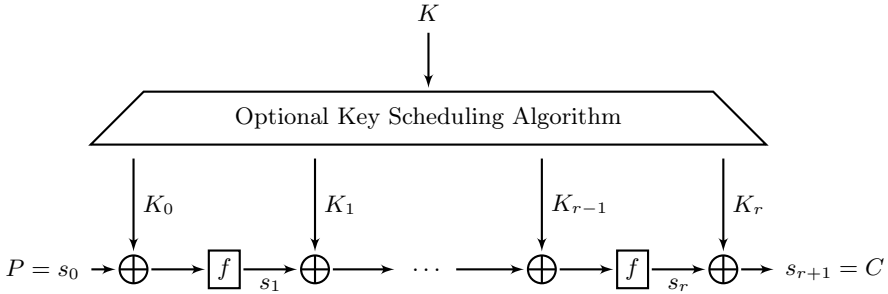
In [26], the problem of tweaking AES-128 without altering the key schedule is handled. The authors introduce `Kiasu-BC` as part of the TWEAKEY framework as a way to securely introduce a 64-bit tweak in the 10-round AES-128 block cipher.

## 2 The TWEAKEY Framework

In this section, we introduce the TWEAKEY construction framework that allows to add a tweak of (almost) any length to a key-alternating block cipher and/or to extend the key space of the block cipher to (almost) any size. In some sense, one can view the TWEAKEY framework as a simple generalization of key-alternating ciphers, offering more flexibility with regards to tweak and/or key sizes. Similarly to key-alternating ciphers, we emphasize that not all TWEAKEY instances are secure. We give in later sections natural instances of TWEAKEY that lead to secure ciphers.

### 2.1 Key-Alternating Ciphers

A symmetric primitive like a block cipher  $E$  is usually built upon a smaller building block  $f$  that is iterated a certain number of times – we refer to such a function  $f$  as a *round function*. Usually  $f$  is cryptographically weak, but its iterations bring security to  $E$ . The number  $r$  of iterations heavily depends on the targeted security of  $E$ , the structure of  $f$ , its differential properties, its algebraic degree, etc. In general, the function  $f$  takes two inputs: the first is the state, while the second is a round-dependent parameter called *round key* or *subkey*. The round keys are obtained by the expansion of a master secret  $K$  with an expansion (key schedule) algorithm:  $K \rightarrow (K_0, \dots, K_r)$ . Formally, for a non-negative  $i < r$ , we write  $f(s_i, K_i) = s_{i+1}$  the function that transforms the state  $s_i$  in one round into the state  $s_{i+1}$ , with the use of the round key  $K_i$ . Initially,



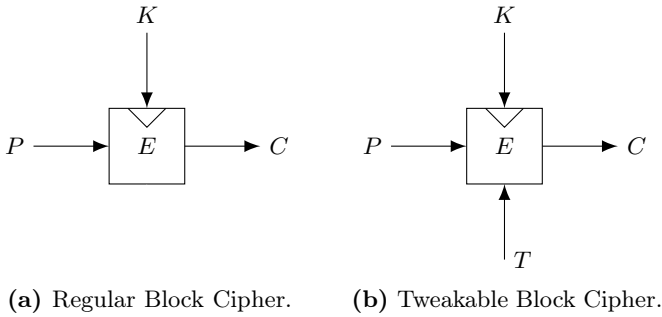
**Fig. 1.** Key-alternating cipher: the function  $f$  is applied  $r$  times, surrounded by subkey mixing operations

the state  $s_0$  is set to the plaintext value  $P$ , and state  $s_r$  at the output of the  $r$ -th round is the ciphertext  $C$ .

As a subclass of iterated block ciphers, we consider further the particular case of *key-alternating block ciphers*, which specify how the round keys are used (see Figure 1). The concept has been initially introduced by Daemen in [14, 16] and has later been reused in many block cipher designs, e.g. [9, 15, 23]. Specifically, we say that  $E$  is a key-alternating cipher when the general form  $f(s_i, K_i) = s_{i+1}$  for  $i < r$  becomes  $f(s_i \oplus K_i) = s_{i+1}$ , where the current state  $s_i$  and the incoming round key  $K_i$  are XORed prior to the application of the round function  $f$ . Moreover, a final round key  $K_r$  is added after the  $r$  applications of  $f$  to produce the ciphertext. The soundness of such a construction has been theoretically studied recently in [1, 10].

## 2.2 Tweakable Block Ciphers

The concept of tweakable block ciphers goes back to the Hasty Pudding cipher [37], and has later been formalized by Liskov, Rivest and Wagner in [30, 31], where they suggest to move the randomization of symmetric primitives brought by the high-level operations of the modes directly at the block-cipher level. The signature of standard block ciphers can be described as  $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  where an  $n$ -bit plaintext  $P$  is transformed into an  $n$ -bit ciphertext  $C = E(K, M)$  using a  $k$ -bit key  $K$ . On top on these inputs, tweakable block ciphers introduce an additional  $t$ -bit parameter  $T$  called tweak (see Figure 2). The signature for a tweakable block cipher therefore becomes  $E : \{0, 1\}^k \times \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ , the ciphertext  $C = E(K, T, P)$  where the tweak  $T$  does not need to be secret and thus can be placed in the public domain. Similarly to a regular block cipher where  $E(K, \cdot)$  is a permutation for all  $K \in \{0, 1\}^k$ , a tweakable block cipher preserves this behavior as  $E(K, T, \cdot)$  is a permutation for all  $(K, T) \in \{0, 1\}^k \times \{0, 1\}^t$ .



**Fig. 2.** Types of ciphers

Usually, the security notion expected from a tweakable block cipher is to be indistinguishable from a tweakable random permutation (a family of independent random permutations parameterized by  $T$ ). It is important to note that the security model considers that the attacker has full control over both the message and the tweak inputs.

**Adversarial Model.** Besides the classical single-key attack model, a typical model for block ciphers is the related-key model, where the adversary can ask for encryption/decryption of plaintext/ciphertext with a key related to the original one. In this article, we only consider the relation between the keys and tweaks to be the classical XOR difference, and refer to [17] for more details on this so-called *key access scheme*. Similarly to the related-key model, the related-tweak model denotes a situation where the adversary can ask for encryption/decryption of plaintext/ciphertext with a tweak related to the original one, while the key remains the original one. Continuing further, we can also combine these two models and consider the related-key related-tweak adversarial model. Moreover, instead of related-key or related-tweak model, one can consider open-key and/or open-tweak models, where the adversary has full control over the key/tweak. This model is reasonable to consider as in practice an active adversary might have a full control over the tweak. For the key, this model might be interesting when the block cipher is used in a hash function setting, where message blocks are usually inserted in the key input of the inner block cipher of the compression function. Since in this article we do not always separate key and tweak input, we sometimes denote *related-tweakey* or *open-tweakey* to refer to related-key related-tweak or open-key open-tweak model, respectively.

### 2.3 The TWEAKEY Construction

In theory, for a tweakable block cipher the distinction between the tweak input and the key input is clear: the former is public and can be fully controlled by the attacker, while the later is secret. This might indicate that in practice the tweak input must be handled more carefully than the key input, since the attacker is

given more power<sup>1</sup>. However, from the point of view of applications, what is intrinsically required for a tweakable block cipher is that computing consecutive cipher calls with different random tweak values should be very efficient, while not necessarily required for the key input. This tends to indicate that, in the contrary, the tweak input should not use more computations than the key input.

This contradiction regarding the proportion of computations between the tweak and key inputs should make tweakable block cipher designers handle both inputs almost equivalently (we note that this is the case for example in Threefish [20]). Moving in this direction, we introduce the TWEAKEY framework, that tries to bridge the gap between key and tweak inputs by providing a unified vision. This framework can be seen as a direct extension of the key-alternating cipher construction. As of today, building a tweakable block cipher with a key-alternating approach has never been considered, but we note that Goldenberg et al. [22] studied how to insert a tweak input inside a Luby-Rackoff cipher from a theoretical point of view.

The term *tweakey* refers to an input that can be both tweak or key material, *without distinction*. Using our framework, the obvious advantage is that one can leverage the work already done on key schedule design in order to build proper tweak schedule, or tweakey schedule more generally.

**The TWEAKEY construction** is a framework to build a  $n$ -bit tweakable block cipher with  $t$ -bit tweak and  $k$ -bit key. It consists of two states: the  $n$ -bit internal state  $s$  and the  $(t + k)$ -bit tweakey state  $tk$ , and we denote respectively as  $s_i$  and  $tk_i$  their values throughout the rounds. The state  $s_0$  is initialized with the plaintext  $P$  (or ciphertext  $C$  for decryption), and  $tk_0$  is initialized with the tweak and key material. Then, the cipher is composed of  $r$  successive rounds each composed of three steps:

- a subtweakey extraction function  $g$  from the tweakey state, and incorporation of this subtweakey to the internal state (for ease of description, we consider that the subtweakey incorporation is done with a simple XOR, but this can be trivially extended to other operations),
- an internal state update permutation  $f$ ,
- a tweakey state update function  $h$ .

This can be summarized as:  $s_{i+1} = f(s_i \oplus g(tk_i))$  followed by  $tk_{i+1} = h(tk_i)$ . At the end, the last subtweakey is incorporated to the last internal state and  $s_r \oplus g(tk_r)$  represents the ciphertext  $C$  (or plaintext  $P$  for decryption). The subtweakeys are usually of size  $n$  bits, but they might be smaller. The framework is depicted in Figure 3.

---

<sup>1</sup> One may argue that key recovery attacks are not to be considered for the tweak input, which makes the tweak and the key inputs fundamentally different. However, from a designer perspective, it seems easier to protect against key-recovery attacks, than against a known-key distinguisher. For example, for most ciphers, more rounds can be attacked in the open-key model than in the related-key model.

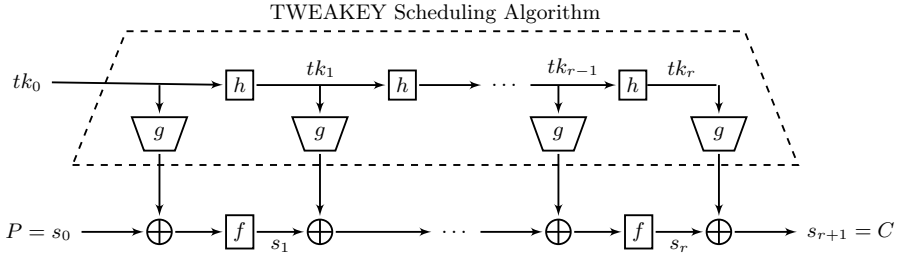


Fig. 3. The TWEAKEY framework

Increasing the amount of tweak or key material obviously renders the task of the designer much more complex in terms of security analysis. To separate these situations, we denote TK- $p$  the class of tweakable block ciphers when one handles  $p \times n$  of tweakkey material. For example, a simple single-key cipher would fit in TK-1, while an  $n$ -bit key,  $n$ -bit tweak block cipher (or for a double-key cipher with no tweak input) would fit in TK-2. By extension, a public permutation would fit in TK-0. The tweakkey material can be any amount of key and/or tweak. A tweak-only cipher can be an interesting primitive as well, for example when building a compression function (the members of the MD-SHA hash function family would actually fit in our framework, the subtweakey having smaller size than  $n$ ).

We emphasize that TWEAKEY is only a framework and, as such, will not guarantee a secure cipher. It is up to the designer to ensure picking a proper TWEAKEY instance. The functions  $f$ ,  $g$  and  $h$  must be chosen along with the number of rounds  $r$  such that no known attack can apply on the resulting primitives. More precisely, this must be true for any choice of the tweak/key size tradeoff inside the tweakkey input. A natural way to achieve this while keeping the same  $f$ ,  $g$  and  $h$  would be to set the number of rounds as the maximal number of required rounds over all the possible tweak/key size tradeoffs. By known attacks, we refer in particular to classical differential/linear attacks, even in related-tweakey or open-tweakey model, or meet-in-the-middle techniques. Moreover, the key schedule is often used to break inherent symmetries from the internal state update function and to break round similarities (for example in the case of AES), hence this has to be taken in account as well.

**Cipher Instances Separation.** Since the tweak and key material are not made distinct in our framework, one might argue that since the tweakable block cipher is always the same whatever is the amount of key or tweak inputs, there are some obvious relations between these different versions. If the designer would prefer to avoid these properties, this can be easily and securely done for example by encoding the various cipher versions on a few bits of the tweakkey state (with two distinct key/tweak sizes versions, one tweak bit would then have to be booked for that matter). Nevertheless, in the rest of this article, we do not consider related-cipher attacks [39].



### 3 The STK Construction

#### 3.1 Motivation

The TWEAKEY framework unifies the tweak and key input for a tweakable block cipher, but does not provide real instantiation of this construction, i.e. which functions  $f$ ,  $g$  and  $h$  (and number of rounds  $r$ ) one should choose. For instance, a trivial example resulting in a non-secure primitive consists in choosing the identity function for the update function  $h$  (i.e. the key schedule of LED), and defining  $g$  as the XOR of all  $n$ -bit tweakey words. In such a case, regardless of the choice of the function  $f$ , the construction would not be secure as cancellations of the tweakey words would lead to outputs of  $g$  consisting of zero bits.

One of the main causes for the low number of ad-hoc tweakable block ciphers is the fact that adding a tweak input makes the security analysis much harder. Building a block cipher secure in the related-key model is already not an easy task, and by incorporating an additional tweak or a double key, the task becomes even more difficult. In the case of AES, there exists tools [8, 21] to analyze the best differential characteristics in the related-key model, but they mainly work for TK-1. As soon as we switch to bigger keys or add tweak inputs, like TK-2 or TK-3, the searches might become infeasible, unless very good characteristics exist to speed up the search with branching cuts, which would mean that the cipher is insecure.

One research direction that we follow in this article consists in finding a construction within the TWEAKEY framework that simplifies this analysis. A potential and natural solution would be that all  $p = (t + k)/n$   $n$ -bit words of tweakey are handled the same way (i.e. the function  $h$  is symmetric with regards to the  $p$   $n$ -bit words of the tweakey state of TK- $p$ ), and that  $g$  simply XORs all these  $n$ -bit tweakey state words to the internal state. The security analysis is simplified as any analysis independently performed on one of the  $n$ -bit words of tweakey will hold for the other words as well (and thus the tools working for TK-1 could now do the analysis even for TK- $p$  with  $p > 1$ ). The problem is to understand what happens when all words are considered together as their interaction might cause potential weaknesses (e.g. if we insert differences in all the tweakey words). For example, assume we would like to build an AES-like cipher with double key: this would fit in TK-2 as  $k = 2n$  and  $t = 0$ . If the two  $n$ -bit tweakey words were treated equivalently, we could use the differential characteristic search tools to assess the security of the primitive with regards to classical differential attacks, and then use this information to pick an appropriate number of rounds. However, there is an obvious weakness if we strictly follow this strategy: starting with the two tweakey words equal would lead to zero being XORed to the internal state every round, since their value would always cancel each other in the XOR. Using constants to separate the two words would work, but only if the  $h$  function is strongly non-linear, which is something we would like to avoid for efficiency reasons. In fact, we would like to push even further the efficiency incentive and only consider nibble-wise substitutions for the  $h$  function.

In the remaining of the section, we propose a simple solution to overcome this issue for AES-like ciphers. The basic idea is to minimize possible differences cancellations between tweakey words by using small field multiplications. Following this mechanism still allows to apply the existing differential characteristic search tools, while avoiding the trivial characteristic in the tweakey scheduling algorithm.

### 3.2 The STK (Superposition TWEAKEY) Construction

The STK construction is a subclass of the TWEAKEY framework for AES-like ciphers defined over a finite field  $GF(2^c)$ . Recall that  $p = (t + k)/n$  denotes the number of  $n$ -bit words in the tweakey state composed of  $t$ -bit tweak and  $k$ -bit key. Assuming that the AES-like S-Box operates on  $c$  bits (thus we have  $n/c$  nibbles in a  $n$ -bit word), the STK construction further specifies the  $f$ ,  $g$  and  $h$  functions as follows (also see Figure 4):

- the function  $g$  simply XORs all the  $p$   $n$ -bit words of the tweakey state to the internal state (AddRoundTweakey, denoted ART), and then XORs a round-dependent constant  $C_i$ ,
- the function  $h$  first applies the same nibble position substitution function  $h'$  to each of the  $p$   $n$ -bit words of the tweakey state, and then multiply each  $c$ -bit cell of the  $j$ -th  $n$ -bit word by a nonzero coefficient  $\alpha_j$  in the finite field  $GF(2^c)$  (with  $\alpha_i \neq \alpha_j$  for all  $1 \leq i \neq j \leq p$ )
- the function  $f$  is an AES-like round.

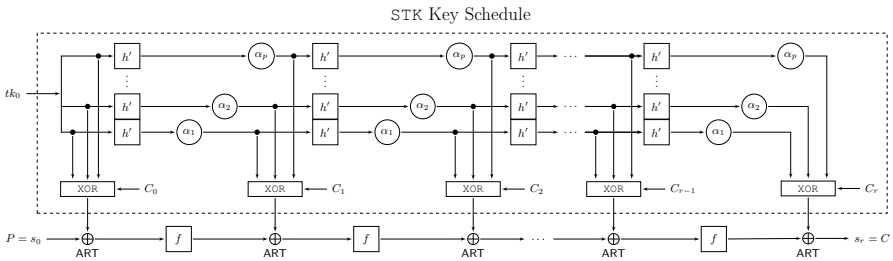


Fig. 4. The STK construction: example with TK- $p$

### 3.3 Rationale behind the STK Construction

Most automated differential analysis tools for AES-like ciphers (e.g., [8, 18, 21]) use truncated differential representation to make feasible the search for differential characteristics. In the truncated difference representation [27], the exact value of a difference in a nibble is not specified. Rather, only the presence (active nibble) or absence (inactive nibble) of a difference is kept track of. In the STK construction, the different subtweakey words will have precisely the same truncated representation of the difference if the input tweakey words have the same

difference. The reason behind this is that they all apply the same functions  $g$  and  $h$ , which are completely independent of the tweakable word considered. This feature already significantly simplifies the analysis for the designer, since a simple TK-1 differential analysis (already known to be possible with the current tools) will ensure the security for all situations in which only a single tweakable word contains a difference. Having all the tweakable words treated almost equivalently is therefore very helpful for the designer.

The issue, however, is to understand what happens when differences are placed in several tweakable words at the same place (in the same nibbles). In particular, the difficulty lies in the cancellations that might happen in the nibbles at the output of  $g$  (recall that  $g$  will XOR all the sub-tweakable word to the state). These cancellations are the reason why having exactly the same update function for all tweakable words leads to a design that is not secure. The trick we use is to apply a nibble-wise multiplication with a distinct coefficient  $\alpha_j$  for all tweakable words. This prevents the large number of cancellation of differences in a particular nibble position at the output of  $g$ . To explain this, first observe that as we apply the very same nibble position substitution function  $h'$  to each of the  $p$  tweakable words, the relative position of the nibble between the tweakable words is always the same (i.e. two nibbles at the same position inside their tweakable word will always keep that property). Thus, we can divide the tweakable nibbles into  $n/c$  fully independent subgroups (according to the nibble position in the  $n$ -bit tweakable words), and to each of these subgroups will correspond one and only one nibble at the output of  $g$  at every round. More precisely, in each subgroup, we have  $p$  input nibbles  $\mathbf{x} = [x_1, \dots, x_p]$  (one in each tweakable word) and  $r + 1$  output nibbles  $\mathbf{y} = [y_0, \dots, y_r]$  (since we have to generate  $r + 1$  sub-tweakeys). Our STK construction ensures that whenever a non-null difference is inserted in the input nibbles of the subgroup, there will always be at least  $r + 1 - p$  active output nibbles. These output nibbles  $\mathbf{y}$  can be expressed in terms of  $\mathbf{x}$  by using a right-matrix multiplication  $\mathbf{y} = \mathbf{x} \times \mathbf{V}$  with the following  $p \times (r + 1)$  Vandermonde matrix:

$$\mathbf{V} = \left( \alpha_i^j \right)_{i,j} = \begin{pmatrix} \alpha_1^0 & \alpha_1^1 & \dots & \alpha_1^r \\ \alpha_2^0 & \alpha_2^1 & \dots & \alpha_2^r \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_p^0 & \alpha_p^1 & \dots & \alpha_p^r \end{pmatrix},$$

In order to minimize the number of nonzero elements in  $\mathbf{y}$  for  $\mathbf{x} \neq 0$ , we need to ensure that all the columns in  $\mathbf{V}$  are linearly independent. This is true as long as the  $\alpha_i$  coefficients,  $1 \leq i \leq p$  are pairwise distinct. Using for example the specific distinct coefficients  $\alpha_j = j \in GF(2^c)$ ,  $1 \leq j < 2^c$ , in TK- $p$ ,  $1 \leq p \leq c - 1$ , then at most  $p$  elements of  $\mathbf{y}$  can be zero for  $\mathbf{x} \neq 0$ , which is the property that we targeted.

To summarize, when we deal with differences in several tweakable words (which is supposedly very hard to analyze due to the important number of nibbles), the study of the STK construction is again the same as for a classical TK-1 analysis,

except that at most  $p$  active output nibbles can be erased in each subgroup. This extra constraint in the search is rather easy to include in the existing analysis tools [8, 21] and this is precisely why we believe the STK construction to be interesting. It has been created with this criteria in mind, so as to ease a systematic cryptographic analysis by existing tools, rather than only relying on ad-hoc constructions, which are *de facto* more difficult to evaluate.

As a side note, the constants  $C_i$  in the STK construction prevent obvious issues regarding symmetries in the internal state for an AES-like cipher, as the RCON constants do for the original AES key scheduling algorithm. The choice of these constants are left at the discretion of the designers, but one could recommend for instance to use the AES RCON constants, based of the exponentiation of 2 in  $GF(2^c)$ , or the exponentiation of any other primitive element in that field.

The nibble positions permutation  $h'$  is also left at the discretion of the designers, but it must be carefully chosen so as to provide the best resistance against classical differential/linear attacks. This will permit the designers to safely choose an appropriate number of rounds  $r$ . This number will of course strongly depend on the amount of tweakable material, since more tweakable material makes it harder for the designer to create a secure tweakable block cipher. Our analysis tools indicate that using identity function instead of  $h'$  would lead to designs that require a great number of rounds. Therefore, we recommend  $h'$  to be a nibble positions permutation so as to prevent the existence of very good differential characteristics, but yet remaining a very efficient function to compute.

### 3.4 Performances

The performance of the STK construction is very high due to the simple transformations used in the schedules – all of them are linear and lightweight. The cost of the nibble position permutation  $h'$  is very low, however, the choice of the coefficients  $\alpha_j$  might have a significant impact on the performances. For optimal efficiency, one should typically use  $\alpha_1 = 1$  and  $\alpha_2 = 2$  in the case of TK-2. For larger instances, TK- $p$  with  $p > 2$ , one could use powers of 2 as coefficients  $\alpha_j$  in order to maintain high efficiency in the computations of the coefficients multiplications. In most of the applications, the tweak is changed more frequently than the key. For instance, in a number of authenticated encryption schemes, the key is the same across different calls to the tweakable cipher, while the tweak is different in each call. Thus, it is reasonable to make the tweak schedule more efficient than the key schedule. Therefore, the tweak schedule should use the most efficiently implementable coefficients  $\alpha_j$  ( $\alpha_1 = 1$  would be the first choice). However, for some particular use-cases, it can be better to assign coefficient  $\alpha_1 = 1$  to a key input. Indeed, for hardware implementations, it might be very valuable in certain scenarios to hard-wire the key in order to greatly reduce the area required (this is a feature of several lightweight ciphers). Yet, this would be possible only if the key input is not modified during the execution of the entire cipher and this is ensured only if  $\alpha_1 = 1$  is assigned to this key input. The efficiency of the STK constructions can best be measured in term of key/tweak agility, i.e. how well

the construction behaves when the key and/or the tweak are frequently changed. Due to the very low number of transformations, and all being completely linear, this construction has obviously one of the simplest possible schedules.

## 4 Conclusion

We have introduced the TWEAKEY framework, which helps designers to build a secure tweakable block cipher by bringing together key schedule design and tweak input. Inside this framework, we have identified a new type of construction, named STK, that is simple and generic and which provides efficient schemes, as shown by the two STK instances Deoxys-BC and Joltik-BC. We have also shown how to directly tweak the AES-128 block cipher, with the very simple and extremely efficient Kiasu-BC tweakable block cipher. The three candidates Kiasu [26], Joltik [25] and Deoxys [24] to the CAESAR authenticated encryption competition by the same authors are based on three instances of either the TWEAKEY or the STK constructions and are claimed secured against classical class of cryptanalytic attacks, as differential and meet-in-the-middle attacks.

We believe this work opens many questions and future works. First, it would be interesting to prove the soundness of our framework and the STK construction. Namely, can we generalize the recent proofs done on key-alternating ciphers? Secondly, we believe that several nibble positions permutation  $h'$  might be of particular interest for the STK construction. The search space is quite large, thus a smart method in order to prune bad candidates is necessary, as well as very optimized search tools. This problem is actually even more complex, since the best permutation for TK- $i$  might not necessarily be the best for TK- $j$  with  $i \neq j$ . Then, a very valuable advance would be to find a way to tweak directly the AES-128 (keeping the original key schedule) with a 128-bit tweak, since the best achievable option to date only handles up to 64-bit tweak (Kiasu-BC). Our searches led us to the conclusion that this seems quite hard to achieve. Finally, the problem of designing a simple, secure and efficient key schedule for AES-like ciphers remains an open problem. Is it possible to find an efficient key schedule that could lead to simple human-readable proofs on the minimal number of active S-Boxes in a differential characteristic in the related-tweakey model?

## References

1. Shamir, A.: How to share a secret. *Communications of the ACM* 22(11), 612–613 (1979)
2. Aumasson, J.-P., Neves, S., Wilcox-O’Hearn, Z., Winnerlein, C.: BLAKE2: Simpler, smaller, fast as MD5. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) *ACNS 2013*. LNCS, vol. 7954, pp. 119–135. Springer, Heidelberg (2013)
3. Barreto, P.S.L.M., Rijmen, V.: The WHIRLPOOL Hashing Function. Submitted to NESSIE (September 2000)
4. Biham, E.: New Types of Cryptanalytic Attacks Using Related Keys. In: Helleseht, T. (ed.) *EUROCRYPT 1993*. LNCS, vol. 765, pp. 398–409. Springer, Heidelberg (1994)

5. Biham, E., Dunkelman, O., Keller, N.: A Unified Approach to Related-Key Attacks. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 73–96. Springer, Heidelberg (2008)
6. Biryukov, A., Khovratovich, D.: Related-Key Cryptanalysis of the Full AES-192 and AES-256. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 1–18. Springer, Heidelberg (2009)
7. Biryukov, A., Khovratovich, D., Nikolić, I.: Distinguisher and Related-Key Attack on the Full AES-256. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 231–249. Springer, Heidelberg (2009)
8. Biryukov, A., Nikolić, I.: Automatic Search for Related-Key Differential Characteristics in Byte-Oriented Block Ciphers: Application to AES, Camellia, Khazad and Others. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 322–344. Springer, Heidelberg (2010)
9. Bogdanov, A.A., Knudsen, L.R., Leander, G., Paar, C., Poschmann, A., Robshaw, M., 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)
10. Bogdanov, A., Knudsen, L.R., Leander, G., Standaert, F.-X., Steinberger, J., Tischhauser, E.: Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 45–62. Springer, Heidelberg (2012)
11. Canetti, R., Garay, J.A. (eds.): CRYPTO 2013, Part I. LNCS, vol. 8042. Springer, Heidelberg (2013)
12. 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)
13. Crowley, P.: Mercy: A Fast Large Block Cipher for Disk Sector Encryption. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 49–63. Springer, Heidelberg (2001)
14. Daemen, J., Govaerts, R., Vandewalle, J.: Correlation Matrices, vol. 35, pp. 275–285
15. Daemen, J., Knudsen, L.R., Rijmen, V.: The Block Cipher SQUARE. In: Biham, E. (ed.) FSE 1997. LNCS, vol. 1267, pp. 149–165. Springer, Heidelberg (1997)
16. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard. Springer (2002)
17. Daemen, J., Rijmen, V.: On the related-key attacks against AES. Proceedings of the Romanian Academy, Series A 13(4), 395–400 (2012)
18. Derbez, P., Fouque, P.-A., Jean, J.: Improved Key Recovery Attacks on Reduced-Round AES in the Single-Key Setting. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 371–387. Springer, Heidelberg (2013)
19. Dworkin, M.J.: SP 800-38E. Recommendation for Block Cipher Modes of Operation: The XTS-AES Mode for Confidentiality on Storage Devices (2010)
20. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: The SKEIN Hash Function Family (2009)
21. Fouque, P.-A., Jean, J., Peyrin, T.: Structural evaluation of AES and chosen-key distinguisher of 9-round AES-128. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 183–203. Springer, Heidelberg (2013)
22. Goldenberg, D., Hohenberger, S., Liskov, M., Schwartz, E.C., Seyalioglu, H.: On Tweaking Luby-Rackoff Blockciphers. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 342–356. Springer, Heidelberg (2007)

23. 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)
24. Jean, J., Nikolić, I., Peyrin, T.: Deoxys v1.1, Submission to the CAESAR competition (2014), <http://www1.spms.ntu.edu.sg/~syllab/Deoxys>
25. Jean, J., Nikolić, I., Peyrin, T.: Joltik v1.1, Submission to the CAESAR competition (2014), <http://www1.spms.ntu.edu.sg/~syllab/Joltik>
26. Jean, J., Nikolić, I., Peyrin, T.: Kiasu v1.1, Submission to the CAESAR competition (2014), <http://www1.spms.ntu.edu.sg/~syllab/Kiasu>
27. Knudsen, L.R.: Truncated and Higher Order Differentials, vol. 35, pp. 196–211
28. Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 306–327. Springer, Heidelberg (2011)
29. Landecker, W., Shrimpton, T., Terashima, R.S.: Tweakable Blockciphers with Beyond Birthday-Bound Security. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 14–30. Springer, Heidelberg (2012)
30. 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)
31. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. *Journal of Cryptology* 24(3), 588–613 (2011)
32. Matsui, M.: On Correlation between the Order of S-Boxes and the Strength of DES. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 366–375. Springer, Heidelberg (1995)
33. Minematsu, K.: Beyond-Birthday-Bound Security Based on Tweakable Block Cipher. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 308–326. Springer, Heidelberg (2009)
34. National Institute of Standards and Technology (NIST): Advanced Encryption Standard (AES). FIPS PUB 197, U.S. Department of Commerce (November 2001)
35. Preneel, B. (ed.): FSE 1994. LNCS, vol. 1008. Springer, Heidelberg (1995)
36. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 16–31. Springer, Heidelberg (2004)
37. Schroepel, R.: The Hasty Pudding Cipher (1998)
38. Shrimpton, T., Terashima, R.S.: A Modular Framework for Building Variable-Input-Length Tweakable Ciphers. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 405–423. Springer, Heidelberg (2013)
39. Wu, H.: Related-Cipher Attacks. In: Deng, R.H., Qing, S., Bao, F., Zhou, J. (eds.) ICICS 2002. LNCS, vol. 2513, pp. 447–455. Springer, Heidelberg (2002)