Insuperability of the Standard Versus Ideal Model Gap for Tweakable Blockcipher Security

Bart Mennink

[1] Digital Security Group, Radboud University, Nijmegen, The Netherlands
`b.mennink@cs.ru.nl`
[2] CWI, Amsterdam, The Netherlands

Abstract. Two types of tweakable blockciphers based on classical blockciphers have been presented over the last years: non-tweak-rekeyable and tweak-rekeyable, depending on whether the tweak may influence the key input to the underlying blockcipher. In the former direction, the best possible security is conjectured to be $2^{\sigma n/(\sigma+1)}$, where n is the size of the blockcipher and σ is the number of blockcipher calls. In the latter direction, Mennink and Wang et al. presented optimally secure schemes, but only in the ideal cipher model. We investigate the possibility to construct a tweak-rekeyable cipher that achieves optimal security in the standard cipher model. As a first step, we note that all standard-model security results in literature implicitly rely on a generic standard-to-ideal transformation, that replaces all keyed blockcipher calls by random secret permutations, at the cost of the security of the blockcipher. Then, we prove that if this proof technique is adopted, tweak-rekeying will not help in achieving optimal security: if $2^{\sigma n/(\sigma+1)}$ is the best one can get *without* tweak-rekeying, optimal 2^n provable security *with* tweak-rekeying is impossible.

Keywords. Optimal security, standard model, ideal model, impossibility, tweakable blockciphers.

1 Introduction

A blockcipher $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ is a family of permutations on $\mathcal{M}$ indexed by a key $k \in \mathcal{K}$. Tweakable blockciphers generalize over the classical ones by the additional input of a *tweak*. More detailed, a tweakable blockcipher $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ satisfies the property that for every key $k \in \mathcal{K}$ and tweak $t \in \mathcal{T}$, $\widetilde{E}(k, t, \cdot)$ is a permutation on $\mathcal{M}$. The key is usually secret, but the tweak is a parameter that is known or even chosen by the user. In 2002, Liskov, Rivest, and Wagner [36] formalized the principle of tweakable blockciphers, and they have gained broad attention since then.

A well-established way of designing a tweakable blockcipher is by building it on top of a conventional blockcipher $E : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such as AES (other approaches will be discussed in Section 1.3). In their seminal work,

Liskov et al. proposed two such constructions:

$$\mathsf{LRW1}(k, t, m) = E(k, E(k, m) \oplus t), \tag{1}$$

$$\mathsf{LRW2}([k, h], t, m) = E(k, m \oplus h(t)) \oplus h(t), \tag{2}$$

where for the latter scheme, h is a universal hash function taken from a family of hash functions H. Related to $\mathsf{LRW2}$ is Rogaway's XEX [50] and its generalizations by Chakraborty and Sarkar [15] and Minematsu [42]: these constructions replace the masking $h(t)$ by a tweaking function based on $E(k, \cdot)$, and therewith eliminate the use of h. All of these constructions, however, only achieve birthday bound $2^{n/2}$ security.

1.1 Quest for Beyond Birthday Bound Security

Various attempts have been made to achieve security beyond the birthday bound, and we identify two approaches: *non-tweak-rekeyable* schemes and *tweak-rekeyable* schemes. In a non-tweak-rekeyable scheme, the key inputs to the underlying blockciphers are independent of the tweak, while in a tweak-rekeyable scheme, the tweak value may have an influence on the key input to the underlying blockcipher.

In the direction of non-tweak-rekeyable schemes, the state of the art centers around the security of $\sigma \geq 1$ round $\mathsf{LRW2}$:

$$\mathsf{LRW2}[\sigma]([\underline{k}, \underline{h}], t, m) = \mathsf{LRW2}([k_\sigma, h_\sigma], t, \cdots \mathsf{LRW2}([k_1, h_1], t, m) \cdots),$$

where $\underline{k} = (k_1, \ldots, k_\sigma)$ are blockcipher keys and $\underline{h} = (h_1, \ldots, h_\sigma)$ instantiations of a universal hash function family H. Landecker et al. [35] and Procter [48] showed that this construction achieves approximately $2^{2n/3}$ security for two rounds, and Lampe and Seurin [34] proved security up to about $2^{\sigma n/(\sigma+2)}$ for an arbitrary even number of rounds. It is conjectured that this scheme achieves $2^{\sigma n/(\sigma+1)}$ security for any $\sigma \geq 1$ [34].

Tweak-rekeyable schemes on the other hand tend to achieve higher levels of security easier, but require a different model. Minematsu [43] introduced the following scheme:

$$\mathsf{Min}(k, t, m) = E(E(k, t\|0^{n-\ell_t}), m), \tag{3}$$

where ℓ_t denotes the length of the tweak, and proved that it achieves security up to $\max\{2^{n/2}, 2^{n-\ell_t}\}$. It is straightforward to derive an attack on Min matching this bound. Note that the scheme only achieves beyond birthday bound security if $\ell_t < n/2$. The tweak size can be elegantly extended using the XTX construction of Minematsu and Iwata [45] at the cost of an extra universal hash function evaluation.

Mennink [38] introduced two constructions based on one, resp. two, blockcipher calls (for Men2 we use the adjusted function from the full version [39], see also Section 5.2):

$$\mathsf{Men1}(k,t,m) = E(k \oplus t, m \oplus z) \oplus z, \text{ where } z = k \otimes t \,, \qquad (4)$$
$$\mathsf{Men2}(k,t,m) = E(k \oplus t, m \oplus z) \oplus z, \text{ where } z = E(2k,t) \,. \qquad (5)$$

The former is proven secure up to about $2^{2n/3}$ queries, the latter approximately optimally 2^n secure. Wang et al. [56] generalized the approach of Mennink and derived a wide class of optimally secure schemes. However, on the downside, these constructions are all analyzed in the ideal cipher model, meaning that the underlying blockcipher is assumed to be perfectly random.

1.2   Optimal Security in Standard Model?

The usage of the ideal cipher model for tweakable blockciphers (and for symmetric-key schemes in general) can be considered controversial: the model is significantly stronger and allows for better security bounds, as evidenced by Mennink's and Wang et al.'s constructions. In this work, we investigate the distinction between the standard and ideal model for the case of tweakable blockciphers, and show the existence of an insuperable gap: whereas in the ideal model optimal security is possible fairly efficiently, we prove under reasonable assumptions that this cannot be achieved in the standard model.

Generic Standard-to-Ideal Reduction. All results on tweakable blockciphers in the standard cipher model [15,34–36,42,43,48,50], implicitly rely on a generic standard-to-ideal reduction, where the keyed blockcipher calls are replaced with secret ideal permutations. This step usually costs $\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D})$, where $\mathcal{D}$ is some strong related-key PRP distinguisher with a certain amount of resources, usually q queries to the keyed oracle $E_{\phi(k)}$ and τ time, and Φ is the set of related-key deriving functions ϕ that $\mathcal{D}$ is allowed to choose. This reduction is in fact also broadly used beyond the area of tweakable blockciphers, such as in authenticated encryption schemes [1,3,11,21,28,33,37,44,50,51] and message authentication codes [4,13,16,24,29,30,41,47,57–59], and in fact, we are not aware of any security result of a construction based on a standard-model blockcipher that uses a structurally different approach. Inspired by this, we investigate what level of tweakable blockcipher security can be achieved if this proof technique is employed.

Lower bound on $\mathbf{Adv}^{\mathbf{srkprp}}_{\Phi,E}(\mathcal{D})$. The generic reduction particularly means that $\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D})$ becomes a necessary term in the derivation, and we derive a lower bound on this advantage, i.e. to see how much the loss is.

Pivotal to the analysis is the set of related-key deriving functions Φ, which differs depending on the application. For instance, for LRW1 and LRW2 we would

have $\Phi_{\mathsf{LRW}} = \{k \mapsto k\}$ and the cost of the reduction is simply the strong PRP security of E. For the cascade $\mathsf{LRW2}[\sigma]$, we would have

$$\Phi_{\mathsf{LRW2}[\sigma]} = \{\underline{k} \mapsto k_i \mid i \in \{1, \dots, \sigma\}\}.$$

As the σ keys are independent this implies a reduction loss of σ times the strong PRP security of E (see also [34, 35]). In both cases, it is fair to assume that the strong PRP security of E is small. The situation gets more technical for tweak-rekeyable schemes. For Min and $\mathsf{Men2}$ we would have larger sets of key-deriving functions:[3]

$$\Phi_{\mathsf{Min}} = \{k \mapsto E(k, t\|0^{n-\ell_t}) \mid t \in \{0,1\}^{\ell_t}\},$$
$$\Phi_{\mathsf{Men2}} = \Phi_{\oplus} \cup \{k \mapsto 2k\} = \{k \mapsto k \oplus \delta \mid \delta \in \mathcal{K}\} \cup \{k \mapsto 2k\}.$$

If the size of Φ increases, the related-key insecurity increases. In more detail, we show that for *any* Φ and *any* E,

$$\mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}) \geq \Omega\left(\frac{\min\{q, |\Phi|\} \cdot r}{2^n}\right),$$

where $\mathcal{D}$ can make q related-key queries to $E_{\phi(k)}$ for random key k and has time to make r offline evaluations of E. (The bound is in fact a bit more fine-grained, cf. Proposition 1, but above simplification is adequate for a proper understanding of the result, and for $\Phi = \Phi_{\mathsf{Min}}$ and $\Phi = \Phi_{\mathsf{Men2}}$ above bound matches the one of Proposition 1.)

For Min, this bound entails a "minimal loss" of $\min\{q, 2^{\ell_t}\} \cdot r/2^n$, a term which in hindsight perfectly explains the security level of Min. For $\mathsf{Men2}$ the loss is even worse: $q \cdot r/2^n$. (Also if the "subkey" $2k$ in $\mathsf{Men2}$ is replaced by an independent key k', the same loss applies.) Concretely, this means that the usage of the generic standard-to-ideal reduction entails impossibility of beyond birthday bound security on $\mathsf{Men2}$. Clearly, this does not invalidate the security of $\mathsf{Men2}$: this negative result is purely due to the lossiness of the generic reduction.

This issue is in fact not new: already in 1998, Bellare et al. encountered it in their seminal paper on Luby-Rackoff backwards [8], and reverted to an analysis in the ideal cipher model. A formal treatment of the situation, however, has not been given. The issue also appeared for schemes based on primitives other than blockciphers. Most prominently, the security of the HMAC message authentication code is based on the PRF security of the underlying function [5,6]. As recently argued by Gaži et al. [23], this standard-model approach might be too pessimistic, and [25] approached the security of HMAC in the ideal compression function model.

Generalized Impossibility. We additionally demonstrate that the issue is not specific to $\mathsf{Men2}$, but applies to a broad spectrum of schemes. In more detail,

---

[3] The generic reduction does not directly apply to $\mathsf{Men1}$ as the same key is used for masking and encrypting, making the usage of the underlying cipher and the overlying mode mutually dependent. This is usually resolved by using two independent keys, such as in $\mathsf{LRW2}$. In this case, $\Phi_{\mathsf{Men1}} = \Phi_{\oplus}$.

we consider a generalized construction of a tweakable blockcipher based on a blockcipher, and show that, if the generic standard-to-ideal reduction is employed, achieving optimal standard-model security *with* tweak-rekeying is at least as hard as *without* tweak-rekeying. Given the state of the art on non-tweak-rekeyable schemes, and particularly the conjecture on LRW2[σ], this shines a negative light on the possibility to find a tweakable blockcipher that is secure in the standard cipher model. Note that the result does not imply that the generic standard-to-ideal reduction is unavoidable, nor that optimal security cannot be achieved, but *if* this reduction is employed and *if* the conjecture on LRW2[σ] is true, optimality seems impossible for this generalized class of functions. The approach followed for this impossibility result may be generalizable to different types of primitives.

Discussion. It is reasonable to question the relevance of any result in any of both models (other questions are discussed in detail in Section 8). It appears that, while the ideal-model results may sometimes be a bit too promising, standard-model results may be extremely loose. This is for instance the case for Men2, where the ideal-model results seem more representative than the standard-model ones. A similar observation was made by Shrimpton and Terashima [55], who introduce the ideal model under key-oblivious access as a weakened version of the ideal cipher model. As a general rule, it is always wise to interpret security results in any of the models with care.

1.3 Other Ways of Tweakable Blockcipher Design

We briefly elaborate on approaches to tweakable blockcipher design, other than constructing them from conventional blockciphers. One approach is to build them "from scratch," as is done for the Hasty Pudding Cipher [53], Mercy [20], Threefish [22], and TWEAKEY [31]. This approach, however, does not allow for any reductionist security argument. Goldenberg et al. [26] and Mitsuda and Iwata [46] transformed generalized Feistel schemes into tweakable generalized Feistel schemes. These constructions only achieve birthday bound security. A novel approach is to build tweakable blockciphers from public permutations, as is done by Sasaki et al. [52], Cogliati et al. [17, 18], Granger et al. [27], and Mennink [40]. This approach achieves comparable levels of security to the non-tweak-rekeyable schemes of above, but the security analysis is inherently done in the ideal permutation model.

1.4 Outline

Our model and the security of (tweakable) blockciphers are formalized in Section 2. In Section 3 we define what we consider a reduction and what we mean with optimal security. This section also includes a formalization of the generic standard-to-ideal reduction. We derive a lower bound on the strong related-key PRP security in Section 4. We revisit LRW2 and Men2 using these formalizations and results in Section 5. In Section 6 we present a generalized tweakable

blockcipher design, and in Section 7 we derive our impossibility result on the optimal security of a generalized tweakable blockcipher. We present an elaborate discussion of the results in Section 8.

2 Notation and Model

For a positive integer n, $\{0,1\}^n$ denotes the set of bit strings of length n. If $\mathcal{X}$ is some set, $x \xleftarrow{\$} \mathcal{X}$ denotes the uniformly random drawing of x from $\mathcal{X}$. The size of $\mathcal{X}$ is denoted by $|\mathcal{X}|$.

2.1 Blockciphers

A blockcipher $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ is a mapping such that for every key $k \in \mathcal{K}$, $E_k(\cdot) = E(k, \cdot)$ is a permutation on $\mathcal{M}$. For fixed k, its inverse is denoted by $E_k^{-1}(\cdot)$. We denote by $\mathsf{BC}(\mathcal{K}, \mathcal{M})$ the set of all such blockciphers. Letting $\mathsf{P}(\mathcal{M})$ be the set of all permutations on $\mathcal{M}$, the strong PRP security of E is defined as

$$\mathbf{Adv}_E^{\mathrm{sprp}}(\mathcal{D}) = \left| \mathbf{Pr}\left(\mathcal{D}^{E_k^{\pm}} = 1 \right) - \mathbf{Pr}\left(\mathcal{D}^{\pi^{\pm}} = 1 \right) \right|,$$

where the probabilities are over $k \xleftarrow{\$} \mathcal{K}$ and $\pi \xleftarrow{\$} \mathsf{P}(\mathcal{M})$, and the random coins of $\mathcal{D}$. Distinguisher $\mathcal{D}$ is typically bounded to have limited resources, such as τ time and q queries to its oracle.

We will consider a generalized security notion that captures the case where a distinguisher can perform related-key attacks. We follow the theoretical framework of Bellare and Kohno [7] and its generalization to tweakable blockciphers by Cogliati and Seurin [19]. Let Φ be a set of permitted related-key deriving functions that map $\mathcal{K}' \to \mathcal{K}$. Define the function $\mathrm{rk}[E] : \mathcal{K}' \times \Phi \times \mathcal{M} \to \mathcal{M}$ as

$$\mathrm{rk}[E](k, \phi, m) = E(\phi(k), m).$$

Note that $\mathrm{rk}[E]$ is invertible for fixed (k, ϕ), and the inverse is defined the straightforward way. The strong related-key PRP security of E is defined as

$$\mathbf{Adv}_{\Phi, E}^{\mathrm{srkprp}}(\mathcal{D}) = \left| \mathbf{Pr}\left(\mathcal{D}^{\mathrm{rk}[E]_k^{\pm}} = 1 \right) - \mathbf{Pr}\left(\mathcal{D}^{\mathrm{rk}[rE]_k^{\pm}} = 1 \right) \right|,$$

where the probabilities are over $k \xleftarrow{\$} \mathcal{K}'$ and $rE \xleftarrow{\$} \mathsf{BC}(\mathcal{K}, \mathcal{M})$, and the random coins of $\mathcal{D}$. Distinguisher $\mathcal{D}$ is typically bounded to have limited resources, such as τ time and q queries to its oracle.

Note that, for the sake of generality, the definition explicitly allows the domain $\mathcal{K}'$ and range $\mathcal{K}$ of the function ϕ to be distinct, although in many cases one simply has $\mathcal{K}' = \mathcal{K}$. If $\mathcal{K}' = \mathcal{K}$ and $\Phi = \{k \mapsto k\}$, the definition of related-key security boils down to the classical definition: $\mathbf{Adv}_{\{k \mapsto k\}, E}^{\mathrm{srkprp}}(\mathcal{D}) = \mathbf{Adv}_E^{\mathrm{sprp}}(\mathcal{D})$. Another famous set of related-key deriving functions is $\Phi_\oplus = \{k \mapsto k \oplus \delta \mid \delta \in \mathcal{K}\}$. The set may also include more involved functions, e.g., ones that internally rely on evaluations of E as well [2]. Throughout, for any set Φ, we assume that it never contains two identical functions, and we denote by $|\Phi|$ the number of functions in the set.

2.2 Tweakable Blockciphers

A tweakable blockcipher $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ is a mapping such that for every $k \in \mathcal{K}$ and every tweak $t \in \mathcal{T}$, the function $\widetilde{E}_k(t, \cdot) = \widetilde{E}(k, t, \cdot)$ is a permutation on $\mathcal{M}$. Like before, its inverse is denoted as $\widetilde{E}_k^{-1}(\cdot, \cdot)$. Let $\widetilde{\mathsf{P}}(\mathcal{T}, \mathcal{M})$ consist of all functions $\widetilde{\pi} : \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ such that for all $t \in \mathcal{T}$, $\widetilde{\pi}(t, \cdot) \in \mathsf{P}(\mathcal{M})$. We define the *standard-model* strong tweakable-PRP security of $\widetilde{E}$ as

$$\mathbf{Adv}_{\widetilde{E}}^{\mathrm{s\text{-}}\widetilde{\mathrm{sprp}}}(\mathcal{D}) = \left| \mathbf{Pr}\left(\mathcal{D}^{\widetilde{E}_k^{\pm}} = 1 \right) - \mathbf{Pr}\left(\mathcal{D}^{\widetilde{\pi}^{\pm}} = 1 \right) \right| ,$$

where probabilities are over $k \xleftarrow{\$} \mathcal{K}$ and $\widetilde{\pi} \xleftarrow{\$} \widetilde{\mathsf{P}}(\mathcal{T}, \mathcal{M})$, and the random coins of $\mathcal{D}$. As before, $\mathcal{D}$ is typically bounded to operate in τ time and q queries to its oracle.

This definition applies to an arbitrary tweakable cipher $\widetilde{E}$. The q queries are solely made to $\widetilde{E}_k^{\pm}$ or $\widetilde{\pi}^{\pm}$, and the time τ can be spent at the distinguisher's discretion. Suppose $\widetilde{E}$ uses a blockcipher E as underlying primitive. If we denote by τ_E the uniform time needed for one evaluation of E, the distinguisher can evaluate this underlying cipher at most $r := \tau / \tau_E$ times. Assuming this block-cipher E does not show underlying weaknesses, we can consider an abstraction of the model and consider the distinguisher to be information-theoretic and to have query access to E and $\widetilde{E}_k^{\pm}$. The approach is also known as the ideal model [9,14,54]. More formally, we define the *ideal-model* strong tweakable-PRP security of $\widetilde{E}$ based on E as

$$\mathbf{Adv}_{\widetilde{E}}^{\mathrm{i\text{-}}\widetilde{\mathrm{sprp}}}(\mathcal{D}) = \left| \mathbf{Pr}\left(\mathcal{D}^{\widetilde{E}_k^{\pm}, E^{\pm}} = 1 \right) - \mathbf{Pr}\left(\mathcal{D}^{\widetilde{\pi}^{\pm}, E^{\pm}} = 1 \right) \right| ,$$

where the probabilities are over $k \xleftarrow{\$} \mathcal{K}$, $E \xleftarrow{\$} \mathsf{BC}(\mathcal{K}, \mathcal{M})$, and $\widetilde{\pi} \xleftarrow{\$} \widetilde{\mathsf{P}}(\mathcal{T}, \mathcal{M})$, and the random coins of $\mathcal{D}$. Distinguisher $\mathcal{D}$ is typically bounded to make q queries to its first (construction) oracle and r queries to its second (primitive) oracle.

3 Formalization of Reduction and Optimality

Formalization of Reduction. In order to formally argue about reductionist security of tweakable blockciphers to classical blockciphers, we first settle our definition of a reductionist proof.

Definition 1. *Let $\widetilde{E}$ be a tweakable blockcipher that internally uses a dedicated blockcipher E. We say that the strong tweakable-PRP security of $\widetilde{E}$ reduces to the strong related-key PRP security of E if for any* s-$\widetilde{\mathrm{sprp}}$ *distinguisher $\mathcal{D}$ there exists an* rk-sprp *distinguisher $\mathcal{D}'$ with comparable resources such that*

$$\mathbf{Adv}_{\widetilde{E}}^{\mathrm{s\text{-}}\widetilde{\mathrm{sprp}}}(\mathcal{D}) \leq \delta \cdot \mathbf{Adv}_{\Phi, E}^{\mathrm{srkprp}}(\mathcal{D}') + \varepsilon ,$$

where Φ is some set of related-key deriving functions depending on the design of $\widetilde{E}$, δ a small constant, and ε is a term negligible in the security parameter of $\widetilde{E}$.

All existing standard-model security proofs on tweakable blockciphers from classical blockciphers [15, 34–36, 42, 43, 48, 50] derive a reductionist bound of the form of Definition 1. Even stronger, all of these results *implicitly* rely on a generic standard-to-ideal reduction which we formalize in below lemma.

Lemma 1 (Generic Standard-to-Ideal Reduction). *Let $\widetilde{E}$ be a tweakable blockcipher that internally uses a dedicated blockcipher E. Assume that $\widetilde{E}$ makes ρ calls to its underlying E and let Φ denote the set of all related-key deriving functions under which E is evaluated. For any* s-$\widetilde{\mathrm{sprp}}$ *distinguisher $\mathcal{D}$,*

$$\mathbf{Adv}_{\widetilde{E}}^{\text{s-}\widetilde{\mathrm{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi,E}^{\text{srkprp}}(\mathcal{D}') + \mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\mathrm{sprp}}}(\mathcal{D}''),$$

where $\mathcal{D}'$ is a distinguisher making at most $\rho \cdot q$ queries and running in time τ, and $\mathcal{D}''$ is an information-theoretic distinguisher making at most q queries to its construction oracle and 0 queries to its primitive oracle.

Proof. The proof follows a simple hybrid argument: first replace the underlying blockcipher evaluations by a random blockcipher $rE \xleftarrow{\$} \mathsf{BC}(\mathcal{K}, \mathcal{M})$. This step costs us $\mathbf{Adv}_{\Phi,E}^{\text{srkprp}}(\mathcal{D}')$. For the remaining analysis of $\mathbf{Adv}_{\widetilde{E}}^{\text{s-}\widetilde{\mathrm{sprp}}}(\mathcal{D})$ with E replaced with secret rE: the distinguisher has no access to $\mathrm{rk}[r\widetilde{E}]_k$ as it does not know k nor rE. Therefore, we can safely assume it has unbounded computational power, and transform it to an information-theoretic adversary that is not allowed to query the underlying primitive. Hence, we obtain the term $\mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\mathrm{sprp}}}(\mathcal{D}'')$ where $\mathcal{D}''$ has resources $(q, 0)$. □

We remark that in Definition 1 and Lemma 1, the set of related-key deriving functions Φ depends on the tweakable blockcipher. In many cases, Φ just consists of the identity function, $\Phi = \{k \mapsto k\}$, in which case the related-key security boils down to the classical strong PRP security. This is for example the case for LRW1 and LRW2, cf. Theorem 1 in Section 5. An example of a more elaborate set of key-deriving functions is $\Phi_\oplus$, cf. Theorem 3 in Section 5.

We furthermore remark that Lemma 1 consists of a somewhat pessimistic bounding: the distinguishers $\mathcal{D}'$ and $\mathcal{D}''$ are in fact constructed from $\mathcal{D}$, and a more accurate bounding would be of the form

$$\mathbf{Adv}_{\widetilde{E}}^{\text{s-}\widetilde{\mathrm{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi,E}^{\text{srkprp}}(\mathcal{D}'[\mathcal{D}]) + \mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\mathrm{sprp}}}(\mathcal{D}''[\mathcal{D}]).$$

In the context of Lemma 1, one would usually maximize both sides of the inequality over all possible distinguishers $\mathcal{D}, \mathcal{D}', \mathcal{D}''$, while in the more accurate bounding one would simply maximize both sides over $\mathcal{D}$. In other words, the bound of Lemma 1 gives a slightly more pessimistic result, but nevertheless, it exactly covers the reduction that is implicitly used in the proofs of [15, 34–36, 42, 43, 48, 50].

Beyond this list of tweakable blockcipher results, the reduction of Lemma 1 in fact finds implicit use in myriad other blockcipher based cryptographic designs, including various authenticated encryption schemes [1, 3, 11, 21, 28, 33, 37, 44, 50, 51] and message authentication codes [4, 13, 16, 24, 29, 30, 41, 47, 57–59]. We are not aware of any security result of a construction based on a standard-model blockcipher that does not follow this reduction but that uses a structurally different approach.

Optimality. We additionally define what we mean with an optimally secure $\widetilde{E}$.

Definition 2. *Let $\widetilde{E}$ be a tweakable blockcipher that internally uses a dedicated blockcipher E. We say that it is* optimally *standard/ideal-model secure if for any distinguisher $\mathcal{D}$ making q queries to its construction oracle and r evaluations of the primitive (where in the standard model, $r = \tau/\tau_E$):*

$$\mathbf{Adv}^{\mathrm{s/i}\text{-}\widetilde{\mathrm{sprp}}}_{\widetilde{E}}(\mathcal{D}) \leq \frac{const \cdot \max\{q, r\}}{\min\{|\mathcal{K}|, |\mathcal{M}|\}} \, ,$$

for some small constant const.

The term $r/|\mathcal{K}|$ corresponds to recovering the key for $\widetilde{E}$; apart from that, the bound is rather arbitrary and conservative to maintain generality. We refer to Bellare and Rogaway [10, Section 3.6] for an informal justification of the bound. We refer to Bernstein and Lange [12] for an interesting discussion on the heuristic existence of hard-to-find attackers.

4 Lower Bound on the Strong Related-Key PRP Security

We will derive a lower bound on the strong related-key PRP security of an arbitrary blockcipher E for any set of key-deriving functions Φ, demonstrating that it can always be distinguished from a random blockcipher up to approximately the birthday bound (apart from various technicalities). Earlier lower bounds, for instance by Bellare and Kohno [7], targeted *specific* sets Φ, but it turns out that the problem gets significantly harder if an *arbitrary* set of key-deriving functions is considered. This is in part attributed to the fact that the lower bound would depend on certain structural properties of Φ.

For a set of key-deriving functions Φ and a key $k \in \mathcal{K}$, we write $\Phi(k) = \{\phi(k) \mid \phi \in \Phi\}$. We denote by $\mathbf{Ex}\left(|\Phi(k)|\right)$ the expected size of the set $\Phi(k)$, where the randomness is taken over the choice of $k \xleftarrow{\$} \mathcal{K}$.

Proposition 1. *Consider a blockcipher $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$, and denote by τ_E the uniform time needed for one evaluation of E. Let Φ be a set of related-key deriving functions. There exists a distinguisher $\mathcal{D}$ making q queries and operating in about τ time, such that*

$$\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi, E}(\mathcal{D}) \geq \max_{\Phi' \subseteq \Phi, |\Phi'| = q'} \frac{\mathbf{Ex}\left(|\Phi'(k)|\right) \cdot r'}{2|\mathcal{K}|} - \frac{1}{|\mathcal{M}| - 1} \, ,$$

where $q' = \min\{q - 1, |\Phi|\}$ and $r' = \tau/\tau_E - 1$, which are required to satisfy $q' \cdot r' \leq |\mathcal{K}|$.

Proof. Let $k \xleftarrow{\$} \mathcal{K}$ be the secret key used to instantiate the distinguisher's oracle.

Let $\Phi' = \{\phi_1, \ldots, \phi_{q'}\} \subseteq \Phi$ be any subset of Φ of size q'. We construct distinguisher $\mathcal{D}_{\Phi'}$ as follows. Denote its oracle by $\mathcal{O}_k \in \{\mathrm{rk}[E]_k, \mathrm{rk}[rE]_k\}$.

(i) Fix any $m \in \mathcal{M}$;

(ii) Let $\mathcal{K}' = \{l_1, \ldots, l_{r'}\} \overset{\$}{\subseteq} \mathcal{K}$ be a set of randomly drawn key values;

(iii) For $i = 1, \ldots, q'$, query $c_i \leftarrow \mathcal{O}_k(\phi_i, m)$;

(iv) For $j = 1, \ldots, r'$, evaluate $y_j \leftarrow E(l_j, m)$;

(v) If for some i, j we have $c_i = y_j$:

    – Fix any $m' \in \mathcal{M}\backslash\{m\}$;

    – Query $c_i' \leftarrow \mathcal{O}_k(\phi_i, m')$ and evaluate $y_j' \leftarrow E(l_j, m')$;

    – If $c_i' = y_j'$, return 1;

(vi) Return 0.

Remains to bound the success probability of $\mathcal{D}_{\Phi'}$. Recall that

$$\mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}_{\Phi'}) \geq \mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[E]_k^{\pm}} = 1\right) - \mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[rE]_k^{\pm}} = 1\right), \qquad (6)$$

and we will analyze these probabilities separately.

Starting with first probability of (6), if $\phi_i(k) = l_j$ for some (i, j), then we necessarily have $c_i = y_j$ and $c_i' = y_j'$. Therefore,

$$\mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[E]_k^{\pm}} = 1\right) \geq \mathbf{Pr}\left(\exists l \in \Phi'(k) \; : \; l \in \mathcal{K}'\right)$$
$$= \sum_{\mathcal{L} \subseteq \mathcal{K}} \mathbf{Pr}\left(\exists l \in \mathcal{L} \; : \; l \in \mathcal{K}' \mid \Phi'(k) = \mathcal{L}\right) \mathbf{Pr}\left(\Phi'(k) = \mathcal{L}\right). \quad (7)$$

Note that two independent sources of randomness are involved: the drawing of the key $k \overset{\$}{\leftarrow} \mathcal{K}$ and the generation of random subset $\mathcal{K}' \overset{\$}{\subseteq} \mathcal{K}$. We proceed with the first probability of (7) for any fixed $\mathcal{L}$ of size at most q'. Via the inclusion-exclusion principle, Bonferroni's inequality states

$$\mathbf{Pr}\left(\exists l \in \mathcal{L} \; : \; l \in \mathcal{K}' \mid \Phi'(k) = \mathcal{L}\right)$$
$$= \sum_{\beta=1}^{q'} (-1)^{\beta-1} \sum_{\substack{\mathcal{L}' \subseteq \mathcal{L} \\ |\mathcal{L}'|=\beta}} \mathbf{Pr}\left(\forall l \in \mathcal{L}' \; : \; l \in \mathcal{K}' \mid \Phi'(k) = \mathcal{L}\right)$$
$$\geq \sum_{l \in \mathcal{L}} \mathbf{Pr}\left(l \in \mathcal{K}' \mid \Phi'(k) = \mathcal{L}\right) - \sum_{\substack{l,l' \in \mathcal{L} \\ l \neq l'}} \mathbf{Pr}\left(l, l' \in \mathcal{K}' \mid \Phi'(k) = \mathcal{L}\right) \qquad (8)$$
$$= \sum_{l \in \mathcal{L}} \frac{r}{|\mathcal{K}|} - \sum_{\substack{l,l' \in \mathcal{L} \\ l \neq l'}} \frac{\binom{r}{2}}{\binom{|\mathcal{K}|}{2}} = \frac{|\mathcal{L}| \cdot r'}{|\mathcal{K}|} - \frac{\binom{|\mathcal{L}|}{2}\binom{r'}{2}}{\binom{|\mathcal{K}|}{2}} \geq \frac{|\mathcal{L}| \cdot r'}{2|\mathcal{K}|},$$

as $q', r' \geq 1$ and $q' \cdot r' \leq |\mathcal{K}|$. This gives for (7):

$$\mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[E]_k^{\pm}} = 1\right) \geq \sum_{\mathcal{L} \subseteq \mathcal{K}} \frac{|\mathcal{L}| \cdot r'}{2|\mathcal{K}|} \mathbf{Pr}\left(\Phi'(k) = \mathcal{L}\right) = \frac{\mathbf{Ex}\left(|\Phi'(k)|\right) \cdot r'}{2|\mathcal{K}|}.$$

For the second probability of (6), focus on the indices (i, j) for which the if-clause is evaluated. We have

$$\mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[rE]_k^{\pm}} = 1\right) \leq \mathbf{Pr}\left(c_i' = y_j' \mid c_i = y_j\right) = \frac{1}{|\mathcal{M}| - 1},$$

using that rE is a random permutation.

We thus obtain from (6):

$$\mathbf{Adv}_{\Phi, E}^{\mathrm{srkprp}}(\mathcal{D}_{\Phi'}) \geq \frac{\mathbf{Ex}\left(|\Phi'(k)|\right) \cdot r'}{2|\mathcal{K}|} - \frac{1}{|\mathcal{M}| - 1}.$$

Note that this bound holds for every choice of Φ'. The claim of Proposition 1 is satisfied for $\mathcal{D} = \mathcal{D}_{\Phi''}$, where

$$\Phi'' = \underset{\Phi' \subseteq \Phi, |\Phi'| = q'}{\mathrm{argmax}} \; \mathbf{Ex}\left(|\Phi'(k)|\right). \qquad \square$$

We remark that the bounding of $\mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[E]_k^{\pm}} = 1\right)$ could be improved (i) by involving more terms of the inclusion-exclusion principle in (8), and (ii) for specific sets of key-deriving functions Φ, by choosing Φ' and $\mathcal{K}'$ more smartly. For instance, for $\Phi = \Phi_\oplus$, the bound reads

$$\mathbf{Pr}\left(\mathcal{D}_{\Phi'}^{\mathrm{rk}[E]_k^{\pm}} = 1\right) \geq \frac{q' \cdot r'}{2|\mathcal{K}|},$$

because $\mathbf{Ex}\left(|\Phi'(k)|\right) = q'$ for $\Phi' \subseteq \Phi_\oplus$ of size q'. It is a straightforward exercise to verify that, for a smart choice of Φ' and $\mathcal{K}'$, the probability can be pulled up to $\frac{q' \cdot r'}{|\mathcal{K}|}$. Nevertheless, the bound of Proposition 1 suffices for our purposes.

We furthermore remark that the attack of Bellare and Kohno [7] for $\Phi = \Phi_\oplus \cup \Phi_+$ is better than the one resulting from Proposition 1. In fact, their attack exploits potential collisions in Φ, rather than preimages. A generalization of Proposition 1 to cover attacks of this kind is beyond the scope of this paper. Nevertheless, we think that it is an interesting problem to derive a generalized *tight* attack on any E and for any Φ, or at least a generalized attack that covers Proposition 1, the attack of Bellare and Kohno, and more.

5 Examples

We discuss two state-of-the-art examples: one from Liskov et al. [36], and one from Mennink [38].

5.1 Liskov et al.'s Scheme

In their original work [36], Liskov et al. introduced two tweakable blockcipher constructions, both achieving approximately $2^{n/2}$ security. We consider the construction that is based on two keys: $k \in \{0, 1\}^n$ and h coming from a universal hash function family H (see also Figure 1):

$$\mathsf{LRW2}([k, h], t, m) = E(k, m \oplus z) \oplus z, \text{ where } z = h(t).$$
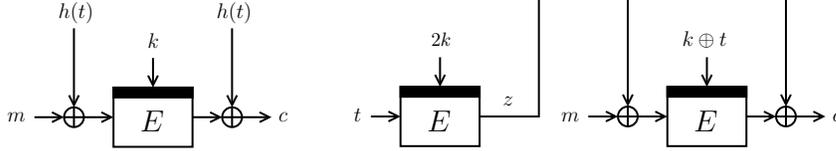
Fig. 1: Tweakable blockcipher LRW2   Fig. 2: Tweakable blockcipher Men2

Follow-up results analyzed the security of a cascade of more than one independent LRW2's [34, 35, 48]; the currently outlined example directly generalizes to these results.

Theorem 1 (Liskov et al. [36], Minematsu [42]). *Let $n \geq 1$, and let H be an ε-almost 2-XOR-universal hash function family.[4] Let $\mathcal{D}$ be a distinguisher making at most q construction queries and running in time τ. Then,*

$$\mathbf{Adv}_{\mathsf{LRW2}}^{\text{s-}\widetilde{\text{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_E^{\text{sprp}}(\mathcal{D}') + \varepsilon q^2 \,,$$

where $\mathcal{D}'$ is a distinguisher making at most q queries and running in time τ.

Note that the strong tweakable-PRP security of LRW2 reduces to the strong PRP security of E in the terminology of Definition 1. The implicit presence of the generic standard-to-ideal reduction of Lemma 1 is obvious from the bound. The term εq^2 is the security bound for LRW2 if the underlying blockcipher is replaced with an ideal secret permutation π.

5.2 Mennink's Scheme

Mennink [38,39] recently introduced two tweak-rekeyable tweakable blockciphers and analyzed them in the ideal cipher model. One of the constructions is the following (see also Figure 2):

$$\mathsf{Men2}(k, t, m) = E(k \oplus t, m \oplus z) \oplus z, \text{ where } z = E(2k, t) \,.$$

Note that we have taken the adjusted scheme from the full version [39], where the masking is done with key $2k$ instead of k. This adjustment was introduced in order to resolve a simple oversight in the proof as pointed out by Wang et al. [56]. Mennink [39] showed that this (adjusted) scheme Men2 achieves approximately 2^n security. We remark that Wang et al. generalized the approach to designing optimally secure tweakable blockciphers. The currently outlined example directly generalizes to the constructions of [56].

Theorem 2 (Mennink [38, 39]). *Let $n \geq 1$. Let $\mathcal{D}$ be a distinguisher making at most q construction queries and r primitive queries. Then,*

$$\mathbf{Adv}_{\mathsf{Men2}}^{\text{i-}\widetilde{\text{sprp}}}(\mathcal{D}) \leq \frac{q+r}{2^n} + \frac{2qr}{(2^n - q)(2^n - q - r)} \,.$$

---

[4] A hash function family $H : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is called ε-almost 2-XOR-universal if for all distinct $x, x' \in \mathcal{X}$ and $y \in \mathcal{Y}$, $\mathbf{Pr}\left(h \xleftarrow{\$} \mathcal{K} \ : \ h(x) \oplus h(x') = y\right) \leq \varepsilon$ [32, 49].

It is easy to verify that for $\max\{q, r\} \leq 2^n/4$, the advantage can be upper bounded by $4\max\{q, r\}/2^n$. Thus, Men2 is optimally ideal-model secure in terms of Definition 2. In the standard model, using the generic transformation of Lemma 1 and the definition of Φ_{Men2} from Section 1,

$$\Phi_{\mathsf{Men2}} = \Phi_\oplus \cup \{k \mapsto 2k\} = \{k \mapsto k \oplus \delta \mid \delta \in \mathcal{K}\} \cup \{k \mapsto 2k\},$$

one can obtain the following result on Men2:

Theorem 3. *Let $n \geq 1$. Let $\mathcal{D}$ be a distinguisher making at most q construction queries and running in time τ. Then,*

$$\mathbf{Adv}_{\mathsf{Men2}}^{\mathrm{s\text{-}\widetilde{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi_{\mathsf{Men2}}, E}^{\mathrm{srkprp}}(\mathcal{D}') + \frac{q}{2^n},$$

where $\mathcal{D}'$ is a distinguisher making at most $2q$ queries and running in time τ.

Proof. By Lemma 1, we have

$$\mathbf{Adv}_{\mathsf{Men2}}^{\mathrm{s\text{-}\widetilde{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi_{\mathsf{Men2}}, E}^{\mathrm{srkprp}}(\mathcal{D}') + \mathbf{Adv}_{\mathsf{Men2}}^{\mathrm{i\text{-}\widetilde{sprp}}}(\mathcal{D}''),$$

where $\mathcal{D}'$ is a distinguisher making at most $2q$ queries and runs in time τ, and $\mathcal{D}''$ an information-theoretic distinguisher making at most q queries to its construction oracle and $r = 0$ queries to its primitive oracle. By Theorem 2, we have $\mathbf{Adv}_{\mathsf{Men2}}^{\mathrm{i\text{-}\widetilde{sprp}}}(\mathcal{D}'') \leq \frac{q}{2^n}$. □

While the bound of Theorem 3 seems to improve over the one of Theorem 2, this is not the case. Indeed, by the remark *after* Proposition 1:

$$\mathbf{Adv}_{\Phi_{\mathsf{Men2}}, E}^{\mathrm{srkprp}}(\mathcal{D}') \geq \mathbf{Adv}_{\Phi_\oplus, E}^{\mathrm{srkprp}}(\mathcal{D}') \geq \frac{(2q-1)(r-1)}{|\mathcal{K}|} - \frac{1}{2^n - 1} = \Omega\left(\frac{qr}{|\mathcal{K}|}\right),$$

contradictory implying that Men2 *cannot be provably optimally standard-model secure if the standard-to-ideal reduction is used.* However, the attack of Proposition 1 to break the strong RK-security of E for related-key deriving functions $\Phi_\oplus$ *does not apply* to Men2: its in- and output of E themselves are masked via a key. A way to resolve this discrepancy would be to include the maskings *within* the definition of related-key security, say the "strong masked related-key PRP" but such a security notion would in fact be equivalent to the strong tweakable-PRP security of Men2. It would be like reducing the security of $E = \mathsf{AES}$ to the "AES-security" of E.

We note that in case one uses Men2 with two independent keys, i.e., replacing $2k$ with independent key k', a comparable reasoning to that of Theorem 3 gives bound

$$\mathbf{Adv}_{\mathsf{Men2}}^{\mathrm{s\text{-}\widetilde{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi_\oplus, E}^{\mathrm{srkprp}}(\mathcal{D}') + \mathbf{Adv}_E^{\mathrm{sprp}}(\mathcal{D}'') + \frac{q}{2^n},$$

where $\mathcal{D}'$ and $\mathcal{D}''$ are distinguishers making at most q queries and running in time τ. The same reasoning as before subsequently applies.

13

6 Generalized Tweakable Blockcipher Design

We consider a generalized tweakable blockcipher $\widetilde{E}$ based on a classical block-cipher E. It follows the generic design of valid tweakable blockciphers by Mennink [38], with two differences. First, for simplicity and sake of presentation, we separate the number of calls to E into ρ message-independent calls and σ message-dependent calls, where ρ and σ are constants independent of the security parameter n. This is without loss of generality, looking back at the formalization of [38] and the assumption that $\widetilde{E}$ processes the data m "as a whole." Second, we will explicitly use two different keys k^a and k^b: k^b is only used in the key inputs to E and k^a is only used in the masking (and indirectly in the data inputs to E).[5] We remark that our description is equivalent to the one of [38] if we set $k^a = k^b$. In the generic design we consider tweaks of size n bits. The generic construction easily generalizes to arbitrarily sized tweaks, but our impossibility result of Section 7 assumes the tweak size to be close to n.

Formally, let $n \geq 1$ and consider a blockcipher $E : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$. We consider a generic tweakable blockcipher $\widetilde{E}[\rho, \sigma] : \{0,1\}^{2n} \times \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ based on $\rho \geq 0$ message-independent precomputation calls to E and $\sigma \geq 0$ message-dependent calls to E as follows (see also Figure 3):

$$
\begin{aligned}
&\textbf{procedure } \widetilde{E}[\rho, \sigma](k^a \| k^b, t, m) \\
&\quad \textbf{for } i = 1, \ldots, \rho \textbf{ do} \\
&\qquad x_i^{\mathsf{pre}} = A_i^{\mathsf{pre}}(k^a, t, y_1^{\mathsf{pre}}, \ldots, y_{i-1}^{\mathsf{pre}}) \\
&\qquad l_i^{\mathsf{pre}} = B_i^{\mathsf{pre}}(k^b, t, y_1^{\mathsf{pre}}, \ldots, y_{i-1}^{\mathsf{pre}}) \\
&\qquad y_i^{\mathsf{pre}} = E(l_i^{\mathsf{pre}}, x_i^{\mathsf{pre}}) \\
&\quad y_0 = m \\
&\quad \textbf{for } i = 1, \ldots, \sigma \textbf{ do} \\
&\qquad x_i = A_i(k^a, t, y_1^{\mathsf{pre}}, \ldots, y_\rho^{\mathsf{pre}}, y_{i-1}) \\
&\qquad l_i = B_i(k^b, t, y_1^{\mathsf{pre}}, \ldots, y_\rho^{\mathsf{pre}}) \\
&\qquad y_i = E(l_i, x_i) \\
&\quad \textbf{return } c = A_{\sigma+1}(k^a, t, y_1^{\mathsf{pre}}, \ldots, y_\rho^{\mathsf{pre}}, y_\sigma)
\end{aligned}
$$

The functions $A_i^{\mathsf{pre}} : \{0,1\}^{(i+1)n} \to \{0,1\}^n$ and $A_i : \{0,1\}^{(\rho+3)n} \to \{0,1\}^n$ compute the data inputs to E (and are keyed via k^a), while the functions $B_i^{\mathsf{pre}} : \{0,1\}^{(i+1)n} \to \{0,1\}^n$ and $B_i : \{0,1\}^{(\rho+2)n} \to \{0,1\}^n$ compute the key inputs to E (and are keyed via k^b). To guarantee invertibility of $\widetilde{E}$, we require that for fixed $k^a, t, y_1^{\mathsf{pre}}, \ldots, y_\rho^{\mathsf{pre}}$ the functions

$$
A_i(k^a, t, y_1^{\mathsf{pre}}, \ldots, y_\rho^{\mathsf{pre}}, \cdot)
$$

are invertible for all $i = 1, \ldots, \sigma+1$. (This is also the reason that A_i does not get inputs $y_0, \ldots, y_{i-2}$.) Apart from this condition, the functions $A_i^{\mathsf{pre}}, B_i^{\mathsf{pre}}, A_i, B_i$

---

[5] Our generalized design, as well as all follow-up results, can be easily generalized to the case of $2(\rho + \sigma) + 1$ keys (one key for every processing function).
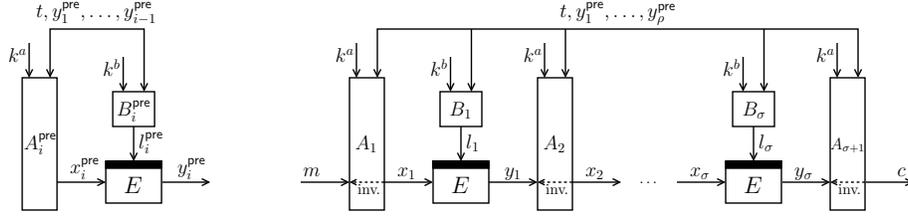
Fig. 3: Tweakable blockcipher $\widetilde{E}[\rho, \sigma]$: precomputation of y_i^{pre} (left) and processing of m (right). "inv." means that the function is invertible

can be any function, as long as they are sufficiently efficient. We put no limitation on how these functions process t; it may be split apart and processed by multiple functions separately.

Note that the message-independent precomputation calls can to a certain extent be reordered. Without loss of generality, there exists a $\rho' \leq \rho$ such that $y_1^{\mathsf{pre}}, \ldots, y_{\rho'}^{\mathsf{pre}}$ are only used as inputs to $A_i^{\mathsf{pre}}, B_i^{\mathsf{pre}}, B_i$, and that $y_{\rho'+1}^{\mathsf{pre}}, \ldots, y_{\rho}^{\mathsf{pre}}$ are also used as inputs to A_i. We define $\rho'' = \rho - \rho'$.

6.1 Key-Uniformity

In the remainder of this work we will require a technical condition on $\widetilde{E}$, which informally assures that $\widetilde{E}$ does not behave structurally different for different keys. For instance, it should not be the case that for some keys, l_1 can take only one value independent of the tweak, while for other keys, it can take 2^n values (one for every tweak). We will call this property "key-uniformity." Note that the condition slightly limits the generality of the scheme, but it is quite reasonable that a scheme should behave comparably for all keys.

For brevity, view the functions B_i^{pre} for $i = 1, \ldots, \rho$ as mappings $(k^a, k^b, t) \mapsto l_i^{\mathsf{pre}}$, and the functions B_i for $i = 1, \ldots, \sigma$ as mappings $(k^a, k^b, t) \mapsto l_i$. Note that, indeed, $(y_1^{\mathsf{pre}}, \ldots, y_i^{\mathsf{pre}})$, is a function of (k^a, k^b, t) for any i.

Definition 3. *We say that $\widetilde{E}$ is c-key-uniform for some $c \geq 0$, if there exist $\lambda_1^{\mathsf{pre}}, \ldots, \lambda_\rho^{\mathsf{pre}}, \lambda_1, \ldots, \lambda_\sigma$ such that for any $k^a \| k^b \in \{0,1\}^{2n}$:*

$$\text{for } i = 1, \ldots, \rho: \qquad 2^{\lambda_i^{\mathsf{pre}} - c} \leq \left| \mathsf{rng}(B_i^{\mathsf{pre}}(k^a, k^b, \cdot)) \right| \leq 2^{\lambda_i^{\mathsf{pre}}},$$

$$\text{for } i = 1, \ldots, \sigma: \qquad 2^{\lambda_i - c} \leq \left| \mathsf{rng}(B_i(k^a, k^b, \cdot)) \right| \leq 2^{\lambda_i}.$$

An observation we will use later on is that $\widetilde{E}$ calls its underlying E with key-deriving functions $\Phi = \Phi_B^{\mathsf{pre}} \cup \Phi_B$, where:

$$
\begin{aligned}
\Phi_B^{\mathsf{pre}} &:= \{(k^a, k^b) \mapsto B_i^{\mathsf{pre}}(k^a, k^b, t) \mid i \in \{\rho' + 1, \ldots, \rho\}, t \in \{0,1\}^n\}, \\
\Phi_B &:= \{(k^a, k^b) \mapsto B_i(k^a, k^b, t) \mid i \in \{1, \ldots, \sigma\}, t \in \{0,1\}^n\}.
\end{aligned}
\tag{9}
$$

15

6.2 Examples

The generalized design represents LRW2 of Figure 1 for $\rho = 0$, $\sigma = 1$, $k^a = h$ (abusing notation), $k^b = k$, and the following processing functions:

$$A_1(h, t, m) = h(t) \oplus m \,,$$
$$B_1(k, t) = k \,,$$
$$A_2(h, t, y_1) = h(t) \oplus y_1 \,.$$

Note that LRW2 is 0-key-uniform (by putting $\lambda_1 = 0$).

The generalized design represents Men2 of Figure 2 for $\rho = \sigma = 1$, $k^b = k$ (k^a is not used), and the following processing functions:

$$A_1^{\mathsf{pre}}(t) = t \,, \qquad\qquad A_1(t, y_1^{\mathsf{pre}}, m) = y_1^{\mathsf{pre}} \oplus m \,,$$
$$B_1^{\mathsf{pre}}(k, t) = 2k \,, \qquad\qquad B_1(k, t, y_1^{\mathsf{pre}}) = k \oplus t \,,$$
$$A_2(t, y_1^{\mathsf{pre}}, y_1) = y_1^{\mathsf{pre}} \oplus y_1 \,.$$

Also Men2 is 0-key-uniform (by putting $\lambda_1^{\mathsf{pre}} = 0$ and $\lambda_1 = n$).

7 Impossibility

We will provide a heuristic argument that if the standard-to-ideal reduction of Lemma 1 is used, optimal security in the standard model by a tweak-rekeyable tweakable blockcipher as described in Section 6 is at least as hard as achieving it by a non-tweak-rekeyable one. The analysis is based on below Assumption 1.

Assumption 1. For any scheme $\widetilde{E}$ as described in Section 6 that is *non-tweak-rekeyable* (hence, l_i^{pre} and l_i are independent of t), and any $\mathcal{T} \subseteq \{0,1\}^n$ of size $|\mathcal{T}| \geq 2^{(\rho''+\sigma)n/(\rho''+\sigma+1)}$, we have

$$\mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\text{sprp}}}(\mathcal{D}) \geq \frac{q^{\rho''+\sigma+1}}{2^{(\rho''+\sigma)n}}$$

for some distinguisher $\mathcal{D}$ which only takes tweaks from $\mathcal{T}$.

The lower bound on $|\mathcal{T}|$ in Assumption 1 is argued by the observation that the bound on $\mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\text{sprp}}}(\mathcal{D})$ is void for $q \geq 2^{(\rho''+\sigma)n/(\rho''+\sigma+1)}$. In other words: any attacker against $\widetilde{E}$ will make at most approximately $q \leq 2^{(\rho''+\sigma)n/(\rho''+\sigma+1)}$ queries and thus require at most that many tweaks for its attack. The assumption is discussed in further detail in Section 8.

Theorem 4. *Let $n \geq 1$ and let $\rho, \sigma \geq 0$. Let $\widetilde{E}$ be any tweakable blockcipher as in Section 6 that is c-key-uniform for some small c. Let Φ be as in (9). If Assumption 1 holds, then*

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\text{srkprp}}(\mathcal{D}') + \max_{\mathcal{D}''} \mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\text{sprp}}}(\mathcal{D}'')$$
$$= \Omega\left(\min\left\{\frac{qr}{2^n}, \frac{q^{\rho''+\sigma+1}}{2^{(\rho''+\sigma)n}}, \frac{r^{(\rho''+\sigma)(\rho''+\sigma+1)}}{2^{((\rho''+\sigma)(\rho''+\sigma+1)-1)n}}\right\}\right) \,, \qquad (10)$$

where the first maximum is taken over all srkprp *distinguishers $\mathcal{D}'$ that make at most $(\rho'' + \sigma) \cdot q$ construction queries and at most r primitive evaluations, and the second maximum is taken over all information-theoretic* i-$\widetilde{\mathrm{sprp}}$ *distinguishers $\mathcal{D}''$ that make q construction queries and 0 primitive queries.*

We give an interpretation of Theorem 4 in Section 7.1, and its proof in Section 7.2.

7.1   Interpretation of Theorem 4

Suppose our goal is to prove security of $\widetilde{E}$ against any s-$\widetilde{\mathrm{sprp}}$ distinguisher $\mathcal{D}$, that can make q construction queries and r evaluations of the primitive. If we would opt to follow the standard-to-ideal reduction of Lemma 1, the first transition would give us an *unavoidable* bound

$$\mathbf{Adv}_{\widetilde{E}}^{\text{s-}\widetilde{\mathrm{sprp}}}(\mathcal{D}) \leq \mathbf{Adv}_{\Phi,E}^{\text{srkprp}}(\mathcal{D}') + \mathbf{Adv}_{\widetilde{E}}^{\text{i-}\widetilde{\mathrm{sprp}}}(\mathcal{D}''),$$

where $\mathcal{D}'$ is a distinguisher making at most $(\rho'' + \sigma)q$ queries[6] and making r primitive queries, and $\mathcal{D}''$ is an information-theoretic distinguisher making at most q queries to its construction oracle and 0 queries to its primitive oracle. Effectively, this step corresponds to replacing the ρ'' message-independent evaluations of E that are used by the masking functions $A_1, \ldots, A_\sigma$ and the σ message-dependent evaluations of E by a secret random related-key blockcipher $\mathrm{rk}[rE]_{k^b}$. The remaining ρ' evaluations of E in the message-independent precomputation occur *indirectly* via the related-key deriving functions.

A next step in the security analysis would be to bound both terms for the strongest possible distinguishers $\mathcal{D}'$ and $\mathcal{D}''$. However, Theorem 4 shows that we can *impossibly* prove optimal security of this bound in terms of Definition 2. The theorem can henceforth be informally captured as follows.

Corollary 1. *If $2^{\sigma n/(\sigma+1)}$ is the best one can get* without *tweak-rekeying, optimal 2^n provable security* with *tweak-rekeying via the generic standard-to-ideal reduction is impossible.*

The bound of Theorem 4 is worse than the bound of Assumption 1, an unavoidable loss to cover worst-case scenarios. The loss shows that with tweak-rekeying we can get *closer to* 2^n than without tweak-rekeying, but we can never achieve optimal security. That is, the bound of (10) cannot give $2^n/const$ security provided that ρ and σ are constant.

The result leaves aside the question of whether the generic standard-to-ideal reduction is strictly necessary. We will discuss this question in Section 8.

---

[6] We remark that the complexity of $\mathcal{D}'$ in Lemma 1 may be optimized depending on the scheme: if ρ' out of ρ calls to the underlying E are *solely* made for the purpose of computing subkeys to later blockcipher calls, then these evaluations of E will be absorbed by the set of related-key deriving functions. This is for instance the case for Min of (3), where the set of key-deriving functions will be Φ_{Min} of Section 1, and $\mathcal{D}'$ can make at most q queries.

7.2 Proof of Theorem 4

Before going to the proof of Theorem 4, we will give a high-level intuition. The core idea is to consider the two terms of (10), and to make a distinction depending on how much freedom the distinguisher has in influencing the rekeying of the σ message-dependent evaluations of E. We consider two cases:

(1) Tweaks have little to no influence on the rekeying of *each* of the blockciphers. In this case, the lower bound on $\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D}')$ (Proposition 1) will be small and we cannot argue based on this part of the bound. On the other hand, the distinguisher can select a large set of tweaks $\mathcal{T}$ for which the blockciphers will never be rekeyed. This way, $\mathcal{D}$ would simply be considering a non-tweak-rekeyable cipher, for which Assumption 1 applies;

(2) Tweaks have a significant influence on the rekeying of *some* of the blockciphers. In this case, the lower bound on $\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D}')$ (Proposition 1) will be significant, and imply the impossibility of an optimal security bound.

Combining the two cases will imply the lower bound of Theorem 4. This high-level overview omits a few technicalities. Most importantly, case (1) requires an upper bound on the influence of the tweaks while case (2) requires a lower bound. This is resolved using the c-key-uniformity of Definition 3.

Proof (Proof of Theorem 4). Let k^a, k^b be two fixed secret keys. Recall that $\widetilde{E}$ is c-key-uniform for some small c. Let

$$\lambda^* = \max\{\lambda^{\mathsf{pre}}_{\rho'+1}, \ldots, \lambda^{\mathsf{pre}}_{\rho}, \lambda_1, \ldots, \lambda_\sigma\}\,.$$

We will derive a lower bound on

$$\max_{\mathcal{D}'} \mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D}') + \max_{\mathcal{D}''} \mathbf{Adv}^{\mathrm{i\text{-}\widetilde{\mathrm{sprp}}}}_{\widetilde{E}}(\mathcal{D}'') \tag{11}$$

by making a case distinction depending on λ^*.

Case $2^{n-\lambda^*(\rho''+\sigma)} \geq 2^{(\rho''+\sigma)n/(\rho''+\sigma+1)}$. For simplicity, we bound (11) as

$$\max_{\mathcal{D}'} \mathbf{Adv}^{\mathrm{srkprp}}_{\Phi,E}(\mathcal{D}') + \max_{\mathcal{D}''} \mathbf{Adv}^{\mathrm{i\text{-}\widetilde{\mathrm{sprp}}}}_{\widetilde{E}}(\mathcal{D}'') \geq \max_{\mathcal{D}''} \mathbf{Adv}^{\mathrm{i\text{-}\widetilde{\mathrm{sprp}}}}_{\widetilde{E}}(\mathcal{D}'')\,,$$

and argue based on the i-$\widetilde{\mathrm{sprp}}$ security, where the maximum is taken over any information-theoretic $\mathcal{D}''$ that makes at most q construction queries and 0 primitive evaluations.

By maximality of λ^*, there is a set $\mathcal{T}' \subseteq \{0,1\}^n$ of size

$$|\mathcal{T}'| \geq \frac{2^n}{\prod_{i=\rho'+1}^{\rho} |\mathsf{rng}(B^{\mathsf{pre}}_i(k^a, k^b, \cdot))| \cdot \prod_{i=1}^{\sigma} |\mathsf{rng}(B_i(k^a, k^b, \cdot))|} \geq 2^{n-\lambda^*(\rho''+\sigma)}$$

such that $B^{\mathsf{pre}}_i(k^a, k^b, t) = B^{\mathsf{pre}}_i(k^a, k^b, t')$ and $B_i(k^a, k^b, t) = B_i(k^a, k^b, t')$ for all $t, t' \in \mathcal{T}'$. By Assumption 1, applied for this $\mathcal{T}'$, we obtain

$$\max_{\mathcal{D}''} \mathbf{Adv}^{\mathrm{i\text{-}\widetilde{\mathrm{sprp}}}}_{\widetilde{E}}(\mathcal{D}'') \geq \frac{q^{\rho''+\sigma+1}}{2^{(\rho''+\sigma)n}}\,. \tag{12}$$

18

Note that $\mathcal{T}'$ is key-dependent and the distinguisher from Assumption 1 does not know $\mathcal{T}'$. This is not a problem, though, as in (12) we are *maximizing* over all distinguishers: the maximum over all distinguishers equals the maximum over all distinguishers that only take tweaks from $\mathcal{T}'$, maximized over all possible sets $\mathcal{T}'$.

Case $2^{n-\lambda^*(\rho''+\sigma)} \leq 2^{(\rho''+\sigma)n/(\rho''+\sigma+1)}$. For simplicity, we bound (11) as

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}') + \max_{\mathcal{D}''} \mathbf{Adv}_{\widetilde{E}}^{\mathrm{i\text{-}\widetilde{sprp}}}(\mathcal{D}'') \geq \max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}')\,,$$

and argue based on the srkprp security, where the maximum is taken over any distinguisher $\mathcal{D}'$ that makes at most $(\rho''+\sigma)\cdot q$ construction queries and at most r primitive evaluations.

By Proposition 1,

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}') \geq \max_{\Phi' \subseteq \Phi, |\Phi'|=q'} \frac{\mathbf{Ex}\left(|\Phi'(k)|\right) \cdot r'}{2^{n+1}} - \frac{1}{2^n-1}\,,$$

where $q' = \min\{(\rho''+\sigma)q - 1, |\Phi|\}$ and $r' = r - 1$. Note that

$$\max_{\Phi' \subseteq \Phi, |\Phi'|=q'} \mathbf{Ex}\left(|\Phi'(k)|\right) \geq \min\{(\rho''+\sigma)q - 1, 2^{\lambda^*-c}\}\,.$$

This maximum is achieved for Φ' being a subset of the set of key-deriving functions for which the maximum λ^* is achieved. As $2^{\lambda^*} \geq 2^{n/((\rho''+\sigma)(\rho''+\sigma+1))}$, we derive:

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}') \geq \frac{\min\{(\rho''+\sigma)q - 1, 2^{n/((\rho''+\sigma)(\rho''+\sigma+1))-c}\} \cdot r'}{2^{n+1}} - \frac{1}{2^n-1}\,.$$

Assuming that $\frac{2^{n/((\rho''+\sigma)(\rho''+\sigma+1))}r'}{2^{n+1+c}} \leq 1$ (otherwise the term will not influence the bound), above term is lower bounded by

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}') \geq \min\left\{\frac{((\rho''+\sigma)q-1)r'}{2^{n+1}}, \frac{2^n r'^{(\rho''+\sigma)(\rho''+\sigma+1)}}{2^{(n+1+c)(\rho''+\sigma)(\rho''+\sigma+1)}}\right\} - \frac{1}{2^n-1}\,.$$

Conclusion. We get for (11):

$$\max_{\mathcal{D}'} \mathbf{Adv}_{\Phi,E}^{\mathrm{srkprp}}(\mathcal{D}') + \max_{\mathcal{D}''} \mathbf{Adv}_{\widetilde{E}}^{\mathrm{i\text{-}\widetilde{sprp}}}(\mathcal{D}'')$$
$$= \Omega\left(\min\left\{\frac{qr}{2^n}, \frac{q^{\rho''+\sigma+1}}{2^{(\rho''+\sigma)n}}, \frac{r^{(\rho''+\sigma)(\rho''+\sigma+1)}}{2^{((\rho''+\sigma)(\rho''+\sigma+1)-1)n}}\right\}\right)\,,$$

assuming that c is a small constant. This completes the proof. $\qquad\square$

19

8 Discussion

The results shine a negative light on optimal standard-model security of tweakable blockciphers and give rise to multiple questions.

What are the implications of the negative standard-model result on Men2 *of Theorem 3?* Despite what the lower bound of Theorem 3 suggests, the gap is mainly caused by the estimation in the hybrid step. More detailed, the step where E is replaced with $\mathbf{Adv}^{\mathrm{srkprp}}_{\Phi_{\mathsf{Men2}}, E}(\mathcal{D}')$ is extremely loose, and an attacker $\mathcal{D}'$ that maximizes its success probability in breaking the related-key security of E is not transformable to an attacker on Men2. Concretely, standard-model security derivations simply *cannot confirm* this.

How do the standard and ideal model compare, and what are the implications of results in both models? This question is not easy to answer. Results in the ideal cipher model are likely to be over-optimistic, while the standard-model results are too loose, mainly due to the seemingly necessary generic reduction of Lemma 1. Intuitively, the "real" security of a scheme satisfies

$$\text{ideal-model security} \ \leq \ \text{"real" security} \ \leq \ \text{standard-model security.}$$

The question is, which of the estimates is tighter? In the ideal-model versus standard-model results on Men2, Theorem 2 versus Theorem 3, the standard-model bound seems to be too loose. For different schemes, it may be the other way around. A potential approach to go is to weaken the ideal-model, an approach for instance followed by Shrimpton and Terashima [55], yet, this approach is ultimately still an ideal-model approach.

In either situation, the findings of this work contribute to a better understanding of how both models compare, and demonstrate that results in the two models should be interpreted with care. We believe that, taking these issues into account, the ideal-cipher security model is still reasonable to consider.

Is Assumption 1 reasonable? Recall that Lampe and Seurin [34] conjectured that the cascade of σ LRW2's achieves $2^{\sigma n/(\sigma+1)}$ security (for the cascade of LRW2's we have $\rho = \rho' = \rho'' = 0$). Assumption 1 suggests that this is the best possible for non-tweak-rekeyable tweakable blockciphers. Regardless of this, it is merely used as starting point: if the assumption holds, then tweak-rekeying will not help in achieving optimal security. Assumption 1 allows for some stretch: if it is not true and a slightly more secure tweakable blockcipher can be constructed, the results (and in particular Theorem 4) generalize accordingly.

The heuristic bound in Theorem 4 is better than the one of Assumption 1, which indicates that tweak-rekeyability *may* result in a better bound than non-tweak-rekeyability (but no optimal one). However, the derivation of the bound of Theorem 4 is very conservative. For instance, it relies on the superset bound $\Phi \supseteq \Phi_B$ of (9) and on a lower bound on $|\Phi_B|$, both of which are loose. Tighter bounds for Theorem 4 may be achieved if more properties of $\widetilde{E}$ are taken into account.

Can we Salvage the Generic Standard-to-Ideal Reduction? Theoretically, Theorem 4 gives a lower bound on an upper bound argued via the generic reduction of Lemma 1. This is in itself little informative, yet it shows us that *if* this classical first-step reduction is used, we cannot get optimal security. Note that we do not claim that the standard-to-ideal reduction is unavoidable, but that *if* this reduction is applied, the term of (10) is unavoidable. A way to circumvent the usage of the reduction and the strong (related-key) PRP security definition as formalized in Section 2.1 may be by using a generalized security model for blockciphers, such as the "strong masked related-key PRP security." Such a generalized security model would, however, only absorb various design properties of the tweakable blockcipher, and shift the problem instead of solving it.

References

1. Abed, F., Fluhrer, S., Forler, C., List, E., Lucks, S., McGrew, D., Wenzel, J.: Pipelineable on-line encryption. In: Fast Software Encryption 2014. Lecture Notes in Computer Science, vol. 8540, pp. 205–223. Springer, Heidelberg (2015)
2. Albrecht, M.R., Farshim, P., Paterson, K.G., Watson, G.J.: On cipher-dependent related-key attacks in the ideal-cipher model. In: Fast Software Encryption 2011. Lecture Notes in Computer Science, vol. 6733, pp. 128–145. Springer, Heidelberg (2011)
3. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Tischhauser, E., Yasuda, K.: Parallelizable and authenticated online ciphers. In: Advances in Cryptology - ASIACRYPT 2013, Part I. Lecture Notes in Computer Science, vol. 8269, pp. 424–443. Springer, Heidelberg (2013)
4. Andreeva, E., Daemen, J., Mennink, B., Van Assche, G.: Security of keyed sponge constructions using a modular proof approach. In: Fast Software Encryption 2015. Lecture Notes in Computer Science, vol. 9054, pp. 364–384. Springer, Heidelberg (2015)
5. Bellare, M.: New proofs for NMAC and HMAC: Security without collision-resistance. In: Advances in Cryptology - CRYPTO 2006. Lecture Notes in Computer Science, vol. 4117, pp. 602–619. Springer, Heidelberg (2006)
6. Bellare, M., Canetti, R., Krawczyk, H.: Keying hash functions for message authentication. In: Advances in Cryptology - CRYPTO '96. Lecture Notes in Computer Science, vol. 1109, pp. 1–15. Springer, Heidelberg (1996)
7. Bellare, M., Kohno, T.: A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In: Advances in Cryptology - EUROCRYPT 2003. Lecture Notes in Computer Science, vol. 2656, pp. 491–506. Springer, Heidelberg (2003)
8. Bellare, M., Krovetz, T., Rogaway, P.: Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. In: Advances in Cryptology - EUROCRYPT '98. Lecture Notes in Computer Science, vol. 1403, pp. 266–280. Springer, Heidelberg (1998)

9. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: ACM Conference on Computer and Communications Security. pp. 62–73. ACM, New York (1993)

10. Bellare, M., Rogaway, P.: Introduction to Modern Cryptography, Pseudorandom Functions (2005), `http://cseweb.ucsd.edu/~mihir/cse207/classnotes.html`

11. Bellare, M., Rogaway, P., Wagner, D.: The EAX mode of operation. In: Fast Software Encryption 2004. Lecture Notes in Computer Science, vol. 3017, pp. 389–407. Springer, Heidelberg (2004)

12. Bernstein, D.J., Lange, T.: Non-uniform cracks in the concrete: The power of free precomputation. In: Advances in Cryptology - ASIACRYPT 2013, Part II. Lecture Notes in Computer Science, vol. 8270, pp. 321–340. Springer, Heidelberg (2013)

13. Black, J., Rogaway, P.: CBC MACs for arbitrary-length messages: The three-key constructions. In: Advances in Cryptology - CRYPTO 2000. Lecture Notes in Computer Science, vol. 1880, pp. 197–215. Springer, Heidelberg (2000)

14. Black, J., Rogaway, P., Shrimpton, T.: Black-box analysis of the block-cipher-based hash-function constructions from PGV. In: Advances in Cryptology - CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 320–335. Springer, Heidelberg (2002)

15. Chakraborty, D., Sarkar, P.: A general construction of tweakable block ciphers and different modes of operations. In: Conference on Information Security and Cryptology - Inscrypt 2006. Lecture Notes in Computer Science, vol. 4318, pp. 88–102. Springer, Heidelberg (2006)

16. Chang, D., Dworkin, M., Hong, S., Kelsey, J., Nandi, M.: A keyed sponge construction with pseudorandomness in the standard model. NIST's 3rd SHA-3 Candidate Conference 2012 (2012)

17. Cogliati, B., Lampe, R., Seurin, Y.: Tweaking Even-Mansour ciphers. In: Advances in Cryptology - CRYPTO 2015, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 493–517. Springer, Heidelberg (2015)

18. Cogliati, B., Seurin, Y.: Beyond-birthday-bound security for tweakable Even-Mansour ciphers with linear tweak and key mixing. In: Advances in Cryptology - ASIACRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9453, pp. 134–158. Springer, Heidelberg (2015)

19. Cogliati, B., Seurin, Y.: On the provable security of the iterated Even-Mansour cipher against related-key and chosen-key attacks. In: Advances in Cryptology - EUROCRYPT 2015, Part I. Lecture Notes in Computer Science, vol. 9056, pp. 584–613. Springer, Heidelberg (2015)

20. Crowley, P.: Mercy: A fast large block cipher for disk sector encryption. In: Fast Software Encryption 2000. Lecture Notes in Computer Science, vol. 1978, pp. 49–63. Springer, Heidelberg (2001)

21. Datta, N., Nandi, M.: ELmE: A misuse resistant parallel authenticated encryption. In: Australasian Conference on Information Security and Privacy - ACISP 2014. Lecture Notes in Computer Science, vol. 8544, pp. 306–321. Springer, Heidelberg (2014)

22. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: The Skein Hash Function Family (2010), submission to NIST's SHA-3 competition

23. Gaži, P., Pietrzak, K., Rybár, M.: The exact PRF-security of NMAC and HMAC. In: Advances in Cryptology - CRYPTO 2014, Part I. Lecture Notes in Computer Science, vol. 8616, pp. 113–130. Springer, Heidelberg (2014)

24. Gaži, P., Pietrzak, K., Tessaro, S.: The exact PRF security of truncation: Tight bounds for keyed sponges and truncated CBC. In: Advances in Cryptology - CRYPTO 2015, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 368–387. Springer, Heidelberg (2015)

25. Gaži, P., Pietrzak, K., Tessaro, S.: Generic security of NMAC and HMAC with input whitening. In: Advances in Cryptology - ASIACRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9453, pp. 85–109. Springer, Heidelberg (2015)

26. Goldenberg, D., Hohenberger, S., Liskov, M., Crump Schwartz, E., Seyalioglu, H.: On tweaking Luby-Rackoff blockciphers. In: Advances in Cryptology - ASIACRYPT 2007. Lecture Notes in Computer Science, vol. 4833, pp. 342–356. Springer, Heidelberg (2007)

27. Granger, R., Jovanovic, P., Mennink, B., Neves, S.: Improved masking for tweakable blockciphers with applications to authenticated encryption. In: Advances in Cryptology - EUROCRYPT 2016, Part I. Lecture Notes in Computer Science, vol. 9665, pp. 263–293. Springer, Heidelberg (2016)

28. Hoang, V., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Advances in Cryptology - EUROCRYPT 2015, Part I. Lecture Notes in Computer Science, vol. 9056, pp. 15–44. Springer, Heidelberg (2015)

29. Iwata, T., Kurosawa, K.: OMAC: One-key CBC MAC. In: Fast Software Encryption 2003. Lecture Notes in Computer Science, vol. 2887, pp. 129–153. Springer, Heidelberg (2003)

30. Iwata, T., Kurosawa, K.: Stronger security bounds for OMAC, TMAC, and XCBC. In: Progress in Cryptology - INDOCRYPT 2003. Lecture Notes in Computer Science, vol. 2904, pp. 402–415. Springer, Heidelberg (2003)

31. Jean, J., Nikolić, I., Peyrin, T.: Tweaks and keys for block ciphers: the TWEAKEY framework. In: Advances in Cryptology - ASIACRYPT 2014, Part II. Lecture Notes in Computer Science, vol. 8874, pp. 274–288. Springer, Heidelberg (2014)

32. Krawczyk, H.: LFSR-based hashing and authentication. In: Advances in Cryptology - CRYPTO '94. Lecture Notes in Computer Science, vol. 839, pp. 129–139. Springer, Heidelberg (1994)

33. Krovetz, T., Rogaway, P.: The software performance of authenticated-encryption modes. In: Fast Software Encryption 2011. Lecture Notes in Computer Science, vol. 6733, pp. 306–327. Springer, Heidelberg (2011)

34. Lampe, R., Seurin, Y.: Tweakable blockciphers with asymptotically optimal security. In: Fast Software Encryption 2013. Lecture Notes in Computer Science, vol. 8424, pp. 133–151. Springer, Heidelberg (2013)

35. Landecker, W., Shrimpton, T., Terashima, R.S.: Tweakable blockciphers with beyond birthday-bound security. In: Advances in Cryptology - CRYPTO 2012. Lecture Notes in Computer Science, vol. 7417, pp. 14–30. Springer, Heidelberg (2012)

36. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable block ciphers. In: Advances in Cryptology - CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 31–46. Springer, Heidelberg (2002)

37. McGrew, D.A., Viega, J.: The security and performance of the Galois/Counter Mode (GCM) of operation. In: Progress in Cryptology - INDOCRYPT 2004. Lecture Notes in Computer Science, vol. 3348, pp. 343–355. Springer, Heidelberg (2004)

38. Mennink, B.: Optimally secure tweakable blockciphers. In: Fast Software Encryption 2015. Lecture Notes in Computer Science, vol. 9054, pp. 428–448. Springer, Heidelberg (2015)

39. Mennink, B.: Optimally secure tweakable blockciphers. Cryptology ePrint Archive, Report 2015/363 (2015), full version of [38]

40. Mennink, B.: XPX: Generalized tweakable Even-Mansour with improved security guarantees. In: Advances in Cryptology - CRYPTO 2016, Part I. Lecture Notes in Computer Science, vol. 9814, pp. 64–94. Springer, Heidelberg (2016)

41. Mennink, B., Reyhanitabar, R., Vizár, D.: Security of full-state keyed and duplex sponge: Applications to authenticated encryption. In: Advances in Cryptology - ASIACRYPT 2015, Part II. Lecture Notes in Computer Science, vol. 9453, pp. 465–489. Springer, Heidelberg (2015)

42. Minematsu, K.: Improved security analysis of XEX and LRW modes. In: Selected Areas in Cryptography 2006. Lecture Notes in Computer Science, vol. 4356, pp. 96–113. Springer, Heidelberg (2007)

43. Minematsu, K.: Beyond-birthday-bound security based on tweakable block cipher. In: Fast Software Encryption 2009. Lecture Notes in Computer Science, vol. 5665, pp. 308–326. Springer, Heidelberg (2009)

44. Minematsu, K.: Parallelizable rate-1 authenticated encryption from pseudorandom functions. In: Advances in Cryptology - EUROCRYPT 2014. Lecture Notes in Computer Science, vol. 8441, pp. 275–292. Springer, Heidelberg (2014)

45. Minematsu, K., Iwata, T.: Tweak-length extension for tweakable block ciphers. In: IMA International Conference 2015. Lecture Notes in Computer Science, vol. 9496, pp. 77–93. Springer, Heidelberg (2015)

46. Mitsuda, A., Iwata, T.: Tweakable pseudorandom permutation from generalized Feistel structure. In: Provable Security 2008. Lecture Notes in Computer Science, vol. 5324, pp. 22–37. Springer, Heidelberg (2008)

47. Mouha, N., Mennink, B., Van Herrewege, A., Watanabe, D., Preneel, B., Verbauwhede, I.: Chaskey: An efficient MAC algorithm for 32-bit microcontrollers. In: Selected Areas in Cryptography 2014. Lecture Notes in Computer Science, vol. 8781, pp. 306–323. Springer, Heidelberg (2014)

48. Procter, G.: A note on the CLRW2 tweakable block cipher construction. Cryptology ePrint Archive, Report 2014/111 (2014)

49. Rogaway, P.: Bucket hashing and its application to fast message authentication. In: Advances in Cryptology - CRYPTO '95. Lecture Notes in Computer Science, vol. 963, pp. 29–42. Springer, Heidelberg (1995)

50. Rogaway, P.: Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In: Advances in Cryptology - ASIACRYPT 2004. Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer, Heidelberg (2004)

51. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: ACM Conference on Computer and Communications Security. pp. 196–205. ACM, New York (2001)

52. Sasaki, Y., Todo, Y., Aoki, K., Naito, Y., Sugawara, T., Murakami, Y., Matsui, M., Hirose, S.: Minalpher v1 (2014), submission to CAESAR competition

53. Schroeppel, R.: The Hasty Pudding Cipher (1998), submission to NIST's AES competition

54. Shannon, C.: Communication theory of secrecy systems. Bell Systems Technical Journal 28(4), 656–715 (1949)

55. Shrimpton, T., Terashima, R.: Salvaging weak security bounds for blockcipher-based constructions. In: Advances in Cryptology - ASIACRYPT 2016, Part I. Lecture Notes in Computer Science, vol. 10031, pp. 429–454. Springer, Heidelberg (2016)

56. Wang, L., Guo, J., Zhang, G., Zhao, J., Gu, D.: How to build fully secure tweakable blockcipher from classical blockcipher. In: Advances in Cryptology - ASIACRYPT 2016, Part I. Lecture Notes in Computer Science, vol. 10031, pp. 455–483. Springer, Heidelberg (2016)

57. Yasuda, K.: The sum of CBC MACs is a secure PRF. In: CT-RSA 2010. Lecture Notes in Computer Science, vol. 5985, pp. 366–381. Springer, Heidelberg (2010)

58. Yasuda, K.: A new variant of PMAC: Beyond the birthday bound. In: Advances in Cryptology - CRYPTO 2011. Lecture Notes in Computer Science, vol. 6841, pp. 596–609. Springer, Heidelberg (2011)

59. Zhang, L., Wu, W., Sui, H., Wang, P.: 3kf9: Enhancing 3GPP-MAC beyond the birthday bound. In: Advances in Cryptology - ASIACRYPT 2012. Lecture Notes in Computer Science, vol. 7658, pp. 296–312. Springer, Heidelberg (2012)