

Tweakable Blockciphers with Asymptotically Optimal Security

Rodolphe Lampe^{1*} and Yannick Seurin^{2**}

¹ University of Versailles, France

² ANSSI, Paris, France

`rodolphe.lampe@gmail.com,yannick.seurin@m4x.org`

Abstract. We consider tweakable blockciphers with beyond the birthday bound security. Landecker, Shrimpton, and Terashima (CRYPTO 2012) gave the first construction with security up to $\mathcal{O}(2^{2n/3})$ adversarial queries (n denotes the block size in bits of the underlying blockcipher), and for which changing the tweak does not require changing the keys for blockcipher calls. In this paper, we extend this construction, which consists of two rounds of a previous proposal by Liskov, Rivest, and Wagner (CRYPTO 2002), by considering larger numbers of rounds $r > 2$. We show that asymptotically, as r increases, the resulting tweakable blockcipher approaches security up to the information bound, namely $\mathcal{O}(2^n)$ queries. Our analysis makes use of a coupling argument, and carries some similarities with the analysis of the iterated Even-Mansour cipher by Lampe, Patarin, and Seurin (ASIACRYPT 2012).

Keywords: tweakable blockcipher, beyond birthday bound, coupling, message authentication code

1 Introduction

Tweakable Blockciphers. Tweakable blockciphers (TBC), introduced by Liskov, Rivest, and Wagner [12], are families of (efficiently invertible) permutations indexed by two functionally distinct parameters: the key (as usual for a blockcipher) and the *tweak*. Phrased differently, a TBC is a family of blockciphers indexed by a tweak. The tweak is usually seen as a public parameter bringing more versatility to the blockcipher, and in particular is assumed to be under control of the attacker when defining security for a TBC.

There are very few constructions of blockciphers which are tweakable “by-design”. The notable examples are the Hasty Pudding cipher [21], Mercy [3], and Threefish, the blockcipher underlying the Skein hash function [6]. See also Goldberg *et al.* [7] who considered how to incorporate a tweak in a Feistel structure.

* This author is partially supported by the French Direction Générale de l’Armement.

** This author is partially supported by the French National Agency of Research: ANR-11-INS-011.

Most of the time however, proposed constructions start from an existing blockcipher (which is assumed to be a secure strong pseudorandom permutation) and build on top of it (in a black-box way) a new family of permutations admitting a tweak. An important property of a TBC is that changing the tweak should be very efficient (this is required for example for applications such as disk or database encryption). Most of the time, changing the key in a blockcipher is a costly operation. Hence, TBC designs where a change in the tweak implies a change in the keys used for calls to the underlying blockcipher tend to be avoided.

Simple constructions of tweakable blockciphers, such as the two proposals made in the original paper by Liskov *et al.* [12], or the XE and XEX constructions by Rogaway [20], are usually proven secure up to the so-called *birthday bound* (BB), *i.e.* up to $\mathcal{O}(2^{n/2})$ adversarial queries for a blockcipher with n -bit block length. The first proposal with beyond BB security was made by Mine-matsu [16], however the construction suffers from a restricted tweak length and requires rekeying the blockcipher when changing the tweak. More recently, Landecker, Shrimpton, and Terashima [11] considered chaining two rounds of the second proposal by [12] (called LRW2 in [11]), which works as follows: given a blockcipher E with keyspace \mathcal{K} and an ε -AXU₂ family of functions \mathcal{H} , the TBC constructed from E through the LRW2 construction has key space $\mathcal{K} \times \mathcal{H}$, and given a key $k \in \mathcal{K}$ and a function $h \in \mathcal{H}$, the encryption of x with tweak t is given by $\tilde{E}_{k,h}(t, x) = h(t) \oplus E_k(x \oplus h(t))$. Landecker *et al.* named CLRW2 the construction resulting from the chaining of two LRW2 constructions, namely:

$$\tilde{E}_{(k_1, k_2), (h_1, h_2)}(t, x) = h_2(t) \oplus E_{k_2}(h_1(t) \oplus E_{k_1}(x \oplus h_1(t)) \oplus h_2(t)) .$$

They proved that the resulting TBC is secure (against adaptive chosen-plaintexts and ciphertexts attacks) up to $\mathcal{O}(2^{2n/3})$ queries. Moreover it admits arbitrary tweaks (by choosing a suitable family \mathcal{H}) and does not require rekeying the blockcipher E when changing the tweak, hence resulting in a very interesting design.

Contributions of this Work. In this paper, we extend the work of Landecker *et al.* [11] by considering longer chains of the LRW2 construction, with the hope that security increases with the number r of rounds (see Figure 1 for an idea of the construction). We simply call this the CLRW construction with r rounds (r -CLRW for short). And indeed, we show that asymptotically as r goes to $+\infty$, the r -CLRW TBC achieves security up to $\mathcal{O}(2^{(1-\varepsilon)n})$ adversarial queries. More precisely, we show the following:

- first, against non-adaptive chosen-plaintexts (NCPA) adversaries, r -CLRW achieves security up to $\mathcal{O}(2^{rn/(r+1)})$ queries;
- then, we prove a general “two weak make one strong” composition theorem for TBCs stating that, given two TBCs \tilde{E} and \tilde{E}' secure against (information-theoretic) NCPA adversaries, the composition $\tilde{E}'^{-1} \circ \tilde{E}$ is secure against adaptive chosen-plaintexts and ciphertexts (CCA) adversaries (care must

be taken in how the tweak is handled when composing). We then use this theorem to prove that r -CLRW achieves security up to $\mathcal{O}(2^{rn/(r+2)})$ queries against CCA adversaries (in other words, it is a strong tweakable pseudo-random permutation up to this number of queries).

Our proof technique for the first part (NCPA adversaries) of the proof uses a coupling argument. The coupling technique is a very useful tool for upper bounding the statistical distance of the distribution of the outputs of an iterated structure to the uniform distribution, and was previously used in cryptography in [17,18,8]. More specifically, our analysis carries some similarities with the analysis of the iterated Even-Mansour cipher by Lampe, Patarin, and Seurin [10], with important differences though. The iterated Even-Mansour cipher [5,2] (also called *key-alternating cipher*) is the construction of a blockcipher in the random permutation model defined as follows: given r public permutations P_1, \dots, P_r on $\{0, 1\}^n$, encryption of x is computed as:

$$y = k_r \oplus P_r(k_{r-1} \oplus P_{r-1}(\dots P_1(k_0 \oplus x) \dots)) ,$$

where k_0, \dots, k_r are $r + 1$ keys of n bits.³ This construction was shown to be secure (against CCA adversaries) up to $\mathcal{O}(2^{n/2})$ queries for $r = 1$ in [5], and up to $\mathcal{O}(2^{2n/3})$ queries for $r = 2$ in [2]. Later, Lampe *et al.* [10] showed, using a coupling argument, that the construction is secure up to $\mathcal{O}(2^{rn/(r+1)})$ queries against NCPA adversaries, and up to $\mathcal{O}(2^{rn/(r+2)})$ queries against CCA adversaries.

Though these results sound similar to ours, the two settings are quite different. Namely, in the Even-Mansour setting, internal permutations P_1, \dots, P_r are publicly accessible by the adversary, whereas in the CLRW setting, E_{k_1}, \dots, E_{k_r} remain “hidden” in the construction. On the other hand, in the Even-Mansour setting, keys are drawn at random at the beginning of the security experiment and fixed afterwards, whereas in the CLRW setting, values $h_i(t)$ (which may be seen as the analog of keys in the iterated Even-Mansour cipher) can be “refreshed” by the adversary through the tweak t . Yet, in both settings, problems that have to be handled in the security proof are collisions at the input of the internal permutations (but the way the adversary provokes such events in both settings is quite different).

Application to MAC and Authenticated Encryption. In [11], the authors defined a nonce-based MAC construction from a TBC called TBC-MAC2 (this is a variant of a previous proposal by [12] called TBC-MAC). This construction preserves the security of the underlying TBC. When instantiated with r -CLRW,

³ We remark that the iterated Even-Mansour cipher can be modified to use only r keys (k_1, \dots, k_r) as follows: the encryption of x is computed as the composition of r rounds of the single-key construction $x \mapsto k_i \oplus P_i(x \oplus k_i)$. The resulting construction is then the strict analog of r -CLRW. Moreover results of [10] carry over to this construction.

this directly yields a secure MAC (*i.e.* a secure PRF) up to $\mathcal{O}(2^{rn/(r+2)})$ queries.⁴ MAC schemes with security beyond the birthday bound are quite rare, and two notable examples have been given by Yasuda [24,25]. Dodis and Steinberger [4] also gave an example with security close to $\mathcal{O}(2^n)$ queries. Their construction is more complex, but relies only on the weaker assumption that the underlying blockcipher is unpredictable.

Besides MAC schemes, the r -CLRW construction can also be used to obtain an authenticated encryption scheme with security close to the information bound: for example, the OCB1 construction by Rogaway [20] gives an authenticated encryption scheme from a TBC with a tight security bound.

Open Problems. We conjecture that our NCPA bound in fact also holds for CCA adversaries, *i.e.* that the r -CLRW construction is secure up to $\mathcal{O}(2^{rn/(r+1)})$ queries *against CCA adversaries*. We think that this is probably the main open problem regarding the construction since for r small, this makes a meaningful gap in the bound. For example, we prove security up to $\mathcal{O}(2^{3n/4})$ queries against CCA adversaries only for 6 rounds, but we conjecture that this already holds for 3 rounds. We note that the corresponding problem is equally open for the iterated Even-Mansour cipher. In a recent preprint [23], Steinberger showed that the iterated Even-Mansour cipher with 3 rounds is secure up to $\mathcal{O}(2^{3n/4})$ queries against CCA adversaries. We are currently unable to transfer his proof technique to the r -CLRW construction for $r = 3$.

We also stress that we view our security proofs more as a *feasibility* result than a practical one. Indeed, as soon as r is more than 4 or maybe 5, the key size and the number of blockcipher calls of the resulting construction will become too large to be reasonably practical. We think however that it is interesting to see that a relatively simple construction enables to approach the information bound. Moreover, improvements may come which will make the construction more efficient or even practical for larger values of r .

Organization. We define the notation and give some useful definitions in Section 2. Then, in Section 3, we prove our security result for r -CLRW against NCPA adversaries. Finally, in Section 4, we prove our composition theorem for tweakable blockciphers and apply it to characterize the security of r -CLRW against CCA adversaries.

2 Preliminaries

2.1 Notation and Security Definitions

The set of integers i such that $a \leq i \leq b$ will be denoted $[a; b]$. When S is a non-empty finite set, we write $s \leftarrow_{\mathfrak{S}} S$ to mean that a value is sampled uniformly

⁴ The security of TBC-MAC2 relies on the security of the underlying TBC against *adaptive* CPA adversaries. We do not have a better bound for r -CLRW against such adversaries than against adaptive CCA ones.

at random from S and assigned to s . By $\mathcal{A}^{\mathcal{O}_1, \mathcal{O}_2, \dots}(x, y, \dots) \Rightarrow z$ we denote the operation of running the (possibly probabilistic) algorithm \mathcal{A} on inputs x, y, \dots with access to oracles $\mathcal{O}_1, \mathcal{O}_2, \dots$ (possibly none), and letting z be the output.

For a set \mathcal{D} , we note $\text{Perm}(\mathcal{D})$ the set of permutations of \mathcal{D} , and we use $\text{Perm}(n)$ to denote the set of permutations of $\mathcal{D} = \{0, 1\}^n$. For two sets \mathcal{D} and \mathcal{K} , we denote $\text{BC}(\mathcal{K}, \mathcal{D})$ the set of blockciphers with domain \mathcal{D} and key space \mathcal{K} , *i.e.* the set of functions $E : \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{D}$ such that for all $k \in \mathcal{K}$, $E_k := E(k, \cdot) \in \text{Perm}(\mathcal{D})$. For three sets \mathcal{D} , \mathcal{K} , and \mathcal{T} , we denote $\text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{D})$ the set of tweakable blockciphers with domain \mathcal{D} , key space \mathcal{K} , and tweak space \mathcal{T} , *i.e.* the set of functions $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{D} \rightarrow \mathcal{D}$ such that for each tweak $t \in \mathcal{T}$, $\tilde{E}(\cdot, t, \cdot) \in \text{BC}(\mathcal{K}, \mathcal{D})$. We will use $\tilde{E}_k(\cdot, \cdot)$ as a shorthand for $\tilde{E}(k, \cdot, \cdot)$. We denote $\text{BC}(\mathcal{K}, n)$ (resp. $\text{TBC}(\mathcal{K}, \mathcal{T}, n)$) the set of blockciphers (resp. tweakable blockciphers) with domain $\mathcal{D} = \{0, 1\}^n$. The *perfect cipher* over \mathcal{D} is defined as the (inefficient) blockcipher whose key space is $\text{Perm}(\mathcal{D})$. In the following, when the domain is clear ($\mathcal{D} = \{0, 1\}^n$ most of the time), we will simply denote E^* the perfect cipher over \mathcal{D} . Sampling a random key for E^* simply means sampling a random permutation over \mathcal{D} .

Fix an integer $q \leq |\mathcal{D}|$. Given a tuple $t = (t_1, \dots, t_q) \in \mathcal{T}^q$, we will denote $\Omega_t \subset \mathcal{D}^q$ the set of possible inputs $x = (x_1, \dots, x_q) \in \mathcal{D}^q$ such that all pairs (t_i, x_i) are pairwise distinct:

$$\Omega_t = \{x := (x_1, \dots, x_q) \in \mathcal{D}^q : (x_i, t_i) \neq (x_j, t_j), \forall i \neq j\} .$$

Let $\mathcal{D}, \mathcal{K}, \mathcal{T}$ be sets, $E \in \text{BC}(\mathcal{K}, \mathcal{D})$ a blockcipher and $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{D})$ a tweakable blockcipher. An adversary \mathcal{A} is said to be non-adaptive if it chooses all its queries (possibly randomly) before issuing the first one, and adaptive otherwise. For any q, τ , we define the following advantages (where, depending on the security experiment, one has $k \leftarrow_{\S} \mathcal{K}$, $\pi \leftarrow_{\S} \text{Perm}(\mathcal{D})$, or $\tilde{\pi} \leftarrow_{\S} \text{BC}(\mathcal{T}, \mathcal{D})$):

$$\text{Adv}_E^{\text{n CPA}}(q, \tau) = \max_{\mathcal{A}} \left| \Pr \left[\mathcal{A}^{E_k(\cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\pi(\cdot)} \Rightarrow 1 \right] \right|$$

$$\text{Adv}_E^{\text{CCA}}(q, \tau) = \max_{\mathcal{A}} \left| \Pr \left[\mathcal{A}^{E_k(\cdot), E_k^{-1}(\cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\pi(\cdot), \pi^{-1}(\cdot)} \Rightarrow 1 \right] \right|$$

$$\text{Adv}_{\tilde{E}}^{\widetilde{\text{n CPA}}}(q, \tau) = \max_{\mathcal{A}} \left| \Pr \left[\mathcal{A}^{\tilde{E}_k(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\tilde{\pi}(\cdot, \cdot)} \Rightarrow 1 \right] \right|$$

$$\text{Adv}_{\tilde{E}}^{\widetilde{\text{CCA}}}(q, \tau) = \max_{\mathcal{A}} \left| \Pr \left[\mathcal{A}^{\tilde{E}_k(\cdot, \cdot), \tilde{E}_k^{-1}(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\tilde{\pi}(\cdot, \cdot), \tilde{\pi}^{-1}(\cdot, \cdot)} \Rightarrow 1 \right] \right| ,$$

where for n CPA and $\widetilde{\text{n CPA}}$ (resp. CCA and $\widetilde{\text{CCA}}$) the max is taken over non-adaptive (resp. adaptive) adversaries making at most q oracle queries and running in time at most τ . The probabilities are over the random coins of \mathcal{A} and the random draw of k , π or $\tilde{\pi}$. In the following, we will refer to $\tilde{\pi}$ as a *tweakable permutation* (though this object is syntactically equivalent to a blockcipher) since it takes the tweak as first input rather than the key.

Definition 1. Let S be an arbitrary set. A family of functions \mathcal{H} from S to $\{0, 1\}^n$ is said to be ε -almost-2-XOR-universal (ε -AXU₂) if for all distinct $x, x' \in S$ and all $y \in \{0, 1\}^n$, one has $\Pr [h \leftarrow_{\S} \mathcal{H} : h(x) \oplus h(x') = y] \leq \varepsilon$.

Note that there exists very efficient and well-studied constructions of ε -AXU₂ function families with $\varepsilon \simeq 2^{-n}$ [22], with short descriptions (*i.e.* keys). We will stick to the convention of using a notation where the key is implicit in the remaining of the paper.

2.2 Statistical Distance and Coupling

Given a finite event space Ω and two probability distributions μ and ν defined on Ω , the *statistical distance* (or total variation distance) between μ and ν , denoted $\|\mu - \nu\|$ is defined as:

$$\|\mu - \nu\| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)| .$$

The following definitions can easily be seen equivalent:

$$\|\mu - \nu\| = \max_{S \subset \Omega} \{\mu(S) - \nu(S)\} = \max_{S \subset \Omega} \{\nu(S) - \mu(S)\} = \max_{S \subset \Omega} \{|\mu(S) - \nu(S)|\} .$$

A *coupling* of μ and ν is a distribution λ on $\Omega \times \Omega$ such that for all $x \in \Omega$, $\sum_{y \in \Omega} \lambda(x, y) = \mu(x)$ and for all $y \in \Omega$, $\sum_{x \in \Omega} \lambda(x, y) = \nu(y)$. In other words, λ is a joint distribution whose marginal distributions are resp. μ and ν . The fundamental result of the coupling technique is the following one. For completeness, we provide the proof in Appendix A.

Lemma 1 (Coupling Lemma). *Let μ and ν be probability distributions on a finite event space Ω , let λ be a coupling of μ and ν , and let $(X, Y) \sim \lambda$ (*i.e.* (X, Y) is a random variable sampled according to distribution λ). Then $\|\mu - \nu\| \leq \Pr[X \neq Y]$.*

2.3 Description of the r -CLRW Construction

We use and adapt the notation of [11]. Let \mathcal{K} be a set and $E \in \text{BC}(\mathcal{K}, n)$ a blockcipher. Let \mathcal{T} be a set and \mathcal{H} a set of functions from \mathcal{T} to $\{0, 1\}^n$. We define the tweakable blockcipher $\text{LRW}^{E, \mathcal{H}}$ with domain $\{0, 1\}^n$, key space $\tilde{\mathcal{K}} = \mathcal{K} \times \mathcal{H}$, and tweak space \mathcal{T} as follows. For any $(k_1, h_1) \in \mathcal{K} \times \mathcal{H}$, $t \in \mathcal{T}$, and $x \in \{0, 1\}^n$, let:

$$\text{LRW}^{E, \mathcal{H}}((k_1, h_1), t, x) = E_{k_1}(x \oplus h_1(t)) \oplus h_1(t) .$$

We also denote $\text{LRW}_{k_1, h_1}^{E, \mathcal{H}} := \text{LRW}^{E, \mathcal{H}}((k_1, h_1), \cdot, \cdot)$ the mapping taking as input $(t, x) \in \mathcal{T} \times \{0, 1\}^n$ and returning $y \in \{0, 1\}^n$.

This construction was called the LRW2 construction in [11], being the second construction proposed by Liskov *et al.* in [12] to build a tweakable blockcipher. We simply call it the LRW construction in this paper. In [11], the authors proposed to chain two LRW constructions to increase the security beyond the birthday bound, and called the resulting construction CLRW2. In this paper, we will consider chaining r LRW constructions with $r > 2$ to obtain security asymptotically close to the information bound.

Let r be a positive integer. We define the tweakable blockcipher $\text{CLR}W^{r,E,\mathcal{H}}$ with domain $\{0,1\}^n$, key space $\tilde{\mathcal{K}} = \mathcal{K}^r \times \mathcal{H}^r$, and tweak space \mathcal{T} as follows. For any $k = (k_1, \dots, k_r) \in \mathcal{K}^r$, $h = (h_1, \dots, h_r) \in \mathcal{H}^r$, $t \in \mathcal{T}$, and $x \in \{0,1\}^n$, let $\text{CLR}W^{r,E,\mathcal{H}}((k,h),t,x)$ be defined as the value y_r obtained recursively as:

$$\begin{cases} y_0 = x \\ y_i = \text{LR}W^{E,\mathcal{H}}((k_i, h_i), t, y_{i-1}) \quad \text{for } 1 \leq i \leq r . \end{cases}$$

We also denote $\text{CLR}W_{k,h}^{r,E,\mathcal{H}} := \text{CLR}W^{r,E,\mathcal{H}}((k,h), \cdot, \cdot)$ the mapping taking as input $(t,x) \in \mathcal{T} \times \{0,1\}^n$ and returning $y \in \{0,1\}^n$. The construction is depicted on Figure 1.




Fig. 1. The $\text{CLR}W^{r,E,\mathcal{H}}$ tweakable blockcipher construction.

Thereafter, we will need the more ideal construction where the blockcipher E is replaced by the perfect cipher E^* over $\{0,1\}^n$ as defined in Section 2.1. The resulting (inefficient) TBC will be denoted $\text{CLR}W^{r,E^*,\mathcal{H}}$, and for every $\pi = (\pi_1, \dots, \pi_r) \in \text{Perm}(n)^r$ and $h = (h_1, \dots, h_r) \in \mathcal{H}^r$, we denote $\text{CLR}W_{\pi,h}^{r,E^*,\mathcal{H}}$ the function defined as $\text{CLR}W_{k,h}^{r,E,\mathcal{H}}$ above, where calls to E_{k_i} are replaced by calls to π_i .

3 Security Analysis for Non-Adaptive Adversaries

In this section, we first deal with non-adaptive chosen-plaintext (NCPA) adversaries. Using a coupling argument, we will prove the following theorem.

Theorem 1. *Let \mathcal{K}, \mathcal{T} be sets, $E \in \text{BC}(\mathcal{K}, n)$ be a blockcipher, and \mathcal{H} be an ε -AXU₂ family of functions from \mathcal{T} to $\{0,1\}^n$. Then one has:*

$$\text{Adv}_{\text{CLR}W^{r,E,\mathcal{H}}}^{\text{nCPA}}(q, \tau) \leq r \cdot \text{Adv}_E^{\text{nCPA}}(q, \tau + rqT) + \frac{q^{r+1}}{r+1} (2\varepsilon)^r ,$$

where T is the time to compute E or E^{-1} .

Using an ε -AXU₂ function family with $\varepsilon \simeq 2^{-n}$, one can see that the construction ensures security up to $\mathcal{O}(2^{rn/(r+1)})$ queries (assuming E is sufficiently secure against NCPA adversaries). The remaining of the section is devoted to the proof of Theorem 1.

3.1 An Hybrid Argument

As a first step in the proof, we replace the blockcipher E used in the CLRW construction by the perfect cipher E^* , *i.e.* we replace calls to E_{k_i} for random and independent keys k_i by calls to uniformly random permutations π_i . If we assume that the blockcipher is a (strong) pseudorandom permutation, the construction using E is only slightly less secure than the construction using E^* , as is captured by the following lemma (we also treat the case of CCA adversaries).

Lemma 2. *For any q, τ , one has:*

$$\begin{aligned} \mathbf{Adv}_{\text{CLRW}^r, E, \mathcal{H}}^{\widetilde{\text{n CPA}}} (q, \tau) &\leq r \cdot \mathbf{Adv}_E^{\text{n CPA}} (q, \tau + rqT) + \mathbf{Adv}_{\text{CLRW}^r, E^*, \mathcal{H}}^{\widetilde{\text{n CPA}}} (q, \tau) \\ \mathbf{Adv}_{\text{CLRW}^r, E, \mathcal{H}}^{\widetilde{\text{CCA}}} (q, \tau) &\leq r \cdot \mathbf{Adv}_E^{\text{CCA}} (q, \tau + rqT) + \mathbf{Adv}_{\text{CLRW}^r, E^*, \mathcal{H}}^{\widetilde{\text{CCA}}} (q, \tau) , \end{aligned}$$

where T is the time to compute E or E^{-1} .

Proof. This is a classical hybrid argument. We only prove the NCPA case, the CCA case is similar. Let \mathcal{A} be a NCPA adversary trying to distinguish $\text{CLRW}^r, E, \mathcal{H}$ from a random tweakable permutation. For each $i \in [1; r]$, consider the following adversary \mathcal{A}_i trying to distinguish E from a random permutation. \mathcal{A}_i runs \mathcal{A} , answering its queries as follows: it computes the r -CLRW construction where the first $i-1$ permutations are uniformly random permutations, the i -th permutation is computed by querying \mathcal{A}_i 's own oracle, and the last $r-i$ permutations correspond to E with randomly drawn keys. Note that \mathcal{A}_i is non-adaptive, makes at most q queries to its own oracle, and runs in time at most $\tau + rqT$. Denote \mathcal{O}_i the oracle defined as the r -CLRW construction where the first i permutations are uniformly random, and the last $r-i$ permutations are E with uniformly random keys. Then, when \mathcal{A}_i is interacting with a random permutation, it answers \mathcal{A} 's queries as \mathcal{O}_{i+1} , whereas when it is interacting with E_k it implements \mathcal{O}_i . Moreover $\mathcal{O}_0 = \text{CLRW}^r, E, \mathcal{H}$ and $\mathcal{O}_r = \text{CLRW}^r, E^*, \mathcal{H}$. By the triangular inequality:

$$\begin{aligned} \left| \Pr \left[\mathcal{A}^{\text{CLRW}^r, E, \mathcal{H}} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\widetilde{\pi}} \Rightarrow 1 \right] \right| &\leq \\ &\sum_{i=0}^{r-1} \left| \Pr \left[\mathcal{A}^{\mathcal{O}_i} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\mathcal{O}_{i+1}} \Rightarrow 1 \right] \right| + \\ &\left| \Pr \left[\mathcal{A}^{\text{CLRW}^r, E^*, \mathcal{H}} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\widetilde{\pi}} \Rightarrow 1 \right] \right| . \end{aligned}$$

The lemma follows by noting that the r first terms are exactly the advantages of adversaries \mathcal{A}_i , which are all upper bounded by $\mathbf{Adv}_E^{\text{n CPA}} (q, \tau + rqT)$. \square

Hence to study the security of $\text{CLRW}^r, E, \mathcal{H}$, we have to study the security of $\text{CLRW}^r, E^*, \mathcal{H}$. This is what we do in the remaining of the proof.

3.2 NCPA Advantage and Statistical Distance

A classical result states that the advantage of a (computationally unbounded) NCPA adversary in distinguishing two systems S_1 and S_2 with at most q queries is upper bounded by the max over any q inputs of the statistical distance between the outputs of the two systems. The two systems we consider here are $\text{CLR}W^{r,E^*,\mathcal{H}}$ and a random tweakable permutation $\tilde{\pi}$, and we want to upper bound the statistical distance between the outputs of $\text{CLR}W^{r,E^*,\mathcal{H}}$ and the outputs of a random tweakable permutation for any q queries to these two systems.

Thereafter, we will consider a NCPA-distinguisher \mathcal{A} which chooses all its (plaintexts) queries in advance. We will denote the i -th query (t_i, x_i) . We denote μ_q the distribution of the outputs when the distinguisher is accessing the $\text{CLR}W^{r,E^*,\mathcal{H}}$ construction (the distribution is defined by the random draw of $\pi = (\pi_1, \dots, \pi_r) \in \text{Perm}(n)^r$ and $h = (h_1, \dots, h_r) \in \mathcal{H}^r$) and μ_0 the distribution of the outputs when the distinguisher is accessing a uniformly random tweakable permutation $\tilde{\pi}$. Hence, for any τ (since this holds even for computationally unbounded adversaries):

$$\text{Adv}_{\text{CLR}W^{r,E^*,\mathcal{H}}}^{\widetilde{\text{nCPA}}}(q, \tau) \leq \|\mu_q - \mu_0\| .$$

In the following, we will denote $\tau = +\infty$ in the advantage when it applies to computationally unbounded adversaries.

3.3 Dividing the Problem into q Simpler Problems

We now give another way to describe the distribution μ_0 . For any $t \in \mathcal{T}$, we do the following experiment: let I_t be the set of indexes $i \in [1; q]$ such that $t_i = t$ and let $(u_i)_{i \in I_t}$ be uniformly random pairwise distinct elements. We claim that the distribution of the outputs of $(t_1, u_1), \dots, (t_q, u_q)$ by any (not necessarily uniform) random tweakable permutation $\tilde{\pi}'$ whose distribution is independent of the distribution of (u_i) is μ_0 , *i.e.* the distribution of the outputs of (t_i, x_i) by a uniformly random tweakable permutation $\tilde{\pi}$. Indeed, for every t , the values $(\tilde{\pi}(t_i, x_i))_{i \in I_t}$ and $(\tilde{\pi}'(t_i, u_i))_{i \in I_t}$ are both uniformly random and pairwise distinct.

Now that we gave this new description of μ_0 , we can split the computation of $\|\mu_q - \mu_0\|$ in q simpler computations. The idea is to construct a distribution μ_ℓ for every $\ell \leq q$ such that μ_ℓ is the distribution of the outputs of a random instance of $\text{CLR}W^{r,E^*,\mathcal{H}}$ queried with (t_i, x_i) for $i = 1, \dots, \ell$ and the last $q - \ell$ queries keep the same tweak t_i as in adversarial queries, but their last coordinate is uniformly random among unqueried values. More precisely, for each $\ell \in [0; q]$, let $((t_1, z_1), \dots, (t_q, z_q))$ be a tuple of queries such that $z_i = x_i$ for $i \leq \ell$, and z_i is uniformly random in $\{0, 1\}^n \setminus \{z_j \mid t_j = t_i, j < i\}$ for $i > \ell$. This means that the first ℓ queries are the adversary's queries and the remaining z_i are chosen uniformly at random among the possible values (all queries have to be pairwise distinct). Denoting μ_ℓ the distribution of the tuple of q outputs when a random

instance of $\text{CLRWR}^{r,E^*,\mathcal{H}}$ receives inputs $((t_1, z_1), \dots, (t_q, z_q))$, we have:

$$\text{Adv}_{\text{CLRWR}^{r,E^*,\mathcal{H}}}^{\widetilde{\text{nCpa}}}(q, \tau = +\infty) \leq \|\mu_q - \mu_0\| \leq \sum_{\ell=0}^{q-1} \|\mu_{\ell+1} - \mu_\ell\| . \quad (1)$$

3.4 Coupling of $\mu_{\ell+1}$ and μ_ℓ

Restricting to the first $\ell + 1$ queries

It remains to upper bound the statistical distance between $\mu_{\ell+1}$ and μ_ℓ , for each $\ell \in [0; q - 1]$. For this, we will construct a suitable coupling of the two distributions. Note that we only have to consider the first $\ell + 1$ elements of the two tuples of outputs since for both distributions, the i -th inputs for $i > \ell + 1$ are sampled uniformly at random. In other words,

$$\|\mu_{\ell+1} - \mu_\ell\| = \|\mu'_{\ell+1} - \mu'_\ell\| , \quad (2)$$

where $\mu'_{\ell+1}$ and μ'_ℓ are the respective distributions of the $\ell + 1$ first outputs of the r -CLRWR construction.

Construction of μ'_ℓ and $\mu'_{\ell+1}$

To define the coupling of $\mu'_{\ell+1}$ and μ'_ℓ , we consider a random $\text{CLRWR}_{\pi,h}^{r,E^*,\mathcal{H}}$ (i.e. $\pi = (\pi_1, \dots, \pi_r)$ and $h = (h_1, \dots, h_r)$ are uniformly random in respectively $\text{Perm}(n)^r$ and \mathcal{H}^r) that receives inputs (t_j, x_j) for $j = 1, \dots, \ell + 1$ so that the outputs are distributed according to $\mu'_{\ell+1}$, and we consider another random $\text{CLRWR}_{\pi',h'}^{r,E^*,\mathcal{H}}$ (i.e. $\pi' = (\pi'_1, \dots, \pi'_r)$ and $h' = (h'_1, \dots, h'_r)$ are uniformly random in respectively $\text{Perm}(n)^r$ and \mathcal{H}^r) with inputs (t_j, z_j) for $j = 1, \dots, \ell + 1$ with $z_j = x_j$ for every $j \leq \ell$ and $z_{\ell+1}$ being uniformly random in $\{0, 1\}^n \setminus \{x_j \mid t_j = t_{\ell+1}, j < \ell + 1\}$, so that the outputs are distributed according to μ'_ℓ .

Notation

For every $j \leq \ell + 1$ and every $i \in [0; r]$, we note x_j^i and z_j^i the values defined by induction:

$$\begin{cases} x_j^0 = x_j, z_j^0 = z_j \\ x_j^{i+1} = h_i(t_j) \oplus \pi_i(x_j^i \oplus h_i(t_j)) \\ z_j^{i+1} = h'_i(t_j) \oplus \pi'_i(z_j^i \oplus h'_i(t_j)) . \end{cases} \quad (3)$$

In order to apply the Coupling Lemma (Lemma 1), we have to find how to correlate (π, h) and (π', h') so that the outputs of both systems $(x_1^r, \dots, x_{\ell+1}^r)$ and $(z_1^r, \dots, z_{\ell+1}^r)$ are equal with high probability. We choose (π, h) uniformly at random and we construct (π', h') as a function of (π, h) . We have to pay attention that the distribution of (π', h') remains uniform in order for $(z_1^r, \dots, z_{\ell+1}^r)$ to be distributed according to μ'_ℓ .

Coupling of the first ℓ queries

For every $j \leq \ell$, the j -th queries x_j^0 and z_j^0 are equal by definition. Considering the system (3), we set $h' = h$ and $\pi'_i(x_j^i \oplus h_i(t_j)) = \pi_i(x_j^i \oplus h_i(t_j))$ for every $j \leq \ell$ and $i \leq r$. This implies that the first ℓ outputs (x_1^r, \dots, x_ℓ^r) and (z_1^r, \dots, z_ℓ^r) are equal.

Coupling of the $(\ell + 1)$ -th query

For every $i \in [0; r - 1]$ we define the coupling for the $\ell + 1$ -th query as follows:

- (1) if there exists $j \leq \ell$ such that $z_{\ell+1}^i \oplus h_i(t_{\ell+1}) = z_j^i \oplus h_i(t_j)$ then $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1}))$ is already defined. Unless we have coupled $z_{\ell+1}^i$ and $x_{\ell+1}^i$ in a previous round, we cannot couple $z_{\ell+1}^{i+1}$ and $x_{\ell+1}^{i+1}$ at this round.
- (2) else, if $z_{\ell+1}^i \oplus h_i(t_{\ell+1}) \neq z_j^i \oplus h_i(t_j)$ for all $j \leq \ell$, then:
 - (a) if there exists $j \leq \ell$ such that $x_{\ell+1}^i \oplus h_i(t_{\ell+1}) = x_j^i \oplus h_i(t_j)$ then we choose $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1}))$ uniformly at random in $\{0, 1\}^n \setminus \{\pi'_i(z_j^i \oplus h_i(t_j)), j \leq \ell\}$. We cannot couple $z_{\ell+1}^{i+1}$ and $x_{\ell+1}^{i+1}$ at this round.
 - (b) else, we define $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1})) = \pi_i(x_{\ell+1}^i \oplus h_i(t_{\ell+1}))$. This implies that $z_{\ell+1}^{i+1} = x_{\ell+1}^{i+1}$.

Note that once $z_{\ell+1}^{i+1} = x_{\ell+1}^{i+1}$, $z_{\ell+1}^{i'} = x_{\ell+1}^{i'}$ for any subsequent round $i' \geq i + 1$, in particular for $i' = r$, so that the coupling is successful.

Verification that (π', h') is uniformly random

We set $h' = h$ and h is uniformly random so h' is uniformly random. During the coupling of the first ℓ queries, we set $\pi'_i(x_j^i \oplus h_i(t_j)) = \pi_i(x_j^i \oplus h_i(t_j))$ for every $j \leq \ell$ and $i \leq r$ and $\pi_i(x_j^i \oplus h_i(t_j))$ is uniformly random among possible values so $\pi'_i(x_j^i \oplus h_i(t_j))$ is uniformly random among possible values. Rule (1) says that if there is a collision with a previous input of π'_i , we cannot choose the value of $\pi'_i(z_j^i \oplus h_i(t_j))$ so this does not change anything to the distribution of π'_i . When conditions of rule (2)(a) are met, we have for some $j \leq \ell$:

$$\begin{cases} \pi_i(x_{\ell+1}^i \oplus h_i(t_{\ell+1})) = \pi_i(x_j^i \oplus h_i(t_j)) = \pi'_i(z_j^i \oplus h_i(t_j)) \\ z_{\ell+1}^i \oplus h_i(t_{\ell+1}) \neq z_j^i \oplus h_i(t_j) \end{cases},$$

which implies that $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1})) \neq \pi_i(x_j^i \oplus h_i(t_j))$. This means that the coupling is impossible and we choose $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1}))$ uniformly at random among possible values to keep π'_i uniformly distributed. Finally, when conditions of rule (2)(b) are met, we have no problem to couple: $\pi_i(x_{\ell+1}^i \oplus h_i(t_{\ell+1}))$ and $\pi'_i(z_{\ell+1}^i \oplus h_i(t_{\ell+1}))$ are both uniformly random among possible values. In conclusion, permutations π'_i are uniformly random and independent as wanted, so that $(z_1^r, \dots, z_{\ell+1}^r)$ is distributed according to μ'_ℓ .

Failure Probability of the Coupling

It remains to upper bound the probability that the coupling fails, *i.e.*

$$(z_1^r, \dots, z_{\ell+1}^r) \neq (x_1^r, \dots, x_{\ell+1}^r) .$$

For every $i \in [0; r-1]$, we denote \mathbf{fail}^i the event that it exists $j \leq \ell$ such that $z_{\ell+1}^i \oplus h_i(t_{\ell+1}) = z_j^i \oplus h_i(t_j)$ or $x_{\ell+1}^i \oplus h_i(t_{\ell+1}) = x_j^i \oplus h_i(t_j)$. This is the event of failing to couple at round i . Then we have:

$$\begin{aligned} \Pr [\mathbf{fail}^i] &\leq \sum_{j \leq \ell} \Pr \left[z_{\ell+1}^i \oplus h_i(t_{\ell+1}) = z_j^i \oplus h_i(t_j) \right. \\ &\quad \left. \text{or } x_{\ell+1}^i \oplus h_i(t_{\ell+1}) = x_j^i \oplus h_i(t_j) \right] \\ &= \sum_{j \leq \ell} \Pr \left[h_i(t_j) \oplus h_i(t_{\ell+1}) = z_j^i \oplus z_{\ell+1}^i \right. \\ &\quad \left. \text{or } h_i(t_j) \oplus h_i(t_{\ell+1}) = x_j^i \oplus x_{\ell+1}^i \right] \\ &\leq \sum_{j \leq \ell} 2\varepsilon = 2\ell\varepsilon , \end{aligned}$$

where the second inequality comes from the ε -AXU₂ property of \mathcal{H} (note that when $t_{\ell+1} = t_j$, necessarily $z_{\ell+1}^i \neq z_j^i$ and $x_{\ell+1}^i \neq x_j^i$ since all queries must be distinct, so that the probability is zero). Since the functions h_i are independent, we have:

$$\Pr \left[\bigcap_{i=0}^{r-1} \mathbf{fail}^i \right] \leq (2\ell\varepsilon)^r . \quad (4)$$

Using the Coupling Lemma and the fact that $z_j^r = x_j^r$ for all $j \leq \ell$, we have:

$$\|\mu'_{\ell+1} - \mu'_\ell\| \leq \Pr [(z_1^r, \dots, z_{\ell+1}^r) \neq (x_1^r, \dots, x_{\ell+1}^r)] \leq \Pr [z_{\ell+1}^r \neq x_{\ell+1}^r] . \quad (5)$$

If we succeed to couple the last query at some round $i \leq r-1$, we know that $z_{\ell+1}^{i'}$ and $x_{\ell+1}^{i'}$ remain equal in the subsequent rounds so that:

$$\Pr [z_{\ell+1}^r \neq x_{\ell+1}^r] \leq \Pr \left[\bigcap_{i=0}^{r-1} \mathbf{fail}^i \right] . \quad (6)$$

Using (4), (5) and (6), we have:

$$\|\mu'_{\ell+1} - \mu'_\ell\| \leq (2\ell\varepsilon)^r . \quad (7)$$

Finally, using (1), (2) and (7), we obtain:

$$\begin{aligned}
\text{Adv}_{\text{CLRWR}^r, E^*, \mathcal{H}}^{\widetilde{\text{nCPA}}}(q, \tau = +\infty) &\leq \sum_{\ell=0}^{q-1} \|\mu'_{\ell+1} - \mu'_\ell\| \\
&\leq \sum_{\ell=0}^{q-1} (2\ell\varepsilon)^r \\
&\leq \int_0^q (2\ell\varepsilon)^r d\ell \\
&= \frac{q^{r+1}}{r+1} (2\varepsilon)^r .
\end{aligned}$$

Theorem 1 then follows from the above inequality combined with Lemma 2.

4 Security Analysis for Adaptive Adversaries

In this section, we first prove a general composition theorem for tweakable blockciphers similar to the “two weak make one strong” theorem for the composition of usual blockciphers. This theorem roughly states that composing two blockciphers secure against NCPA adversaries yields a blockcipher secure against CCA adversaries [14,15]. We prove exactly the same result for TBCs, but we stress that the exact way the tweak is used in composition is important: namely, the *same* tweak must be used in both ciphers. We state this theorem in the information-theoretic setting (*i.e.* for computationally unbounded adversaries) since we will then apply it to the $\text{CLRWR}^r, E^*, \mathcal{H}$ construction which has information-theoretic security. Corresponding theorems in the computational setting are usually much harder to obtain. Our proof technique is an extension of the “H Coefficients” technique of Patarin [19] to tweakable blockciphers. One could probably use the formalism of random systems [13] to obtain a tight bound in the computational setting as in [15], however subtle problems have been recently found in this proof technique [9] so that we prefer the more simple and straightforward statistical approach. We then apply this result to prove the security of r -CLRW against CCA adversaries up to $\mathcal{O}(2^{rn/(r+2)})$ queries.

4.1 Definitions and Preliminary Results

Fix $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{D})$. For any $t = (t_1, \dots, t_q) \in \mathcal{T}^q$, and any $x = (x_1, \dots, x_q) \in \Omega_t$, we denote $\nu_{(t,x)}$ the distribution on Ω_t induced by \tilde{E} and $\nu_{(t,x)}^*$ the distribution induced by a random tweakable permutation on inputs (t_i, x_i) , namely for $y = (y_1, \dots, y_q) \in \Omega_t$:

$$\begin{cases} \nu_{(t,x)}(y) = \Pr [k \leftarrow_{\S} \mathcal{K} : \tilde{E}_k(t_i, x_i) = y_i, \forall i \leq q] \\ \nu_{(t,x)}^*(y) = \Pr [\tilde{\pi} \leftarrow_{\S} \text{BC}(\mathcal{T}, \mathcal{D}) : \tilde{\pi}(t_i, x_i) = y_i, \forall i \leq q] . \end{cases}$$

Note that $\nu_{(t,x)}^*$ is uniform over Ω_t (the exact cardinality of Ω_t depends on t). For any $\alpha \in [0, 1]$, we note $S_{\alpha, (t,x)}$ the set of $y \in \Omega_t$ satisfying $\nu_{(t,x)}(y) \geq (1 - \alpha)\nu_{(t,x)}^*(y)$.

We start by proving two lemmas which will be useful for our main result. The first one says that if, for every $t = (t_1, \dots, t_q)$, $x = (x_1, \dots, x_q)$, and $y = (y_1, \dots, y_q)$, the probability that \tilde{E}_k maps (t_i, x_i) to y_i for all i is close to the corresponding probability for a random tweakable permutation, then the advantage of any adversary in distinguishing \tilde{E} from a random tweakable permutation with q queries is small.

Lemma 3. *Fix $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{D})$, and $q \leq |\mathcal{D}|$. If there exists $\alpha \in [0, 1]$ such that, for all $t \in \mathcal{T}^q$ and for all $x \in \Omega_t$, $\nu_{(t,x)}^*(S_{\alpha, (t,x)}) = 1$, then*

$$\mathbf{Adv}_{\tilde{E}}^{\text{cca}}(q, \tau = +\infty) \leq \alpha .$$

Proof. Consider a computationally unbounded CCA attacker \mathcal{A} making q queries to an oracle \mathcal{O} acting like \tilde{E} or like a random tweakable permutation $\tilde{\pi}$. We assume *wlog* that \mathcal{A} is deterministic. We note $\delta = (\delta_1, \dots, \delta_q) \in \mathcal{D}^q$ the transcript of the attack, defined as follows. If \mathcal{A} makes a direct query (t_1, x_1) and receives an answer y_1 , one has $\delta_1 = y_1$ and then, the attacker continues his attack and receives the next answers $\delta_2, \dots, \delta_q$. If the attacker makes an inverse query (t_i, y_i) then δ_i is the answer x_i . For each transcript δ , we denote $t(\delta), x(\delta)$ and $y(\delta)$ the corresponding values of $t_1, \dots, t_q, x_1, \dots, x_q, y_1, \dots, y_q$. We denote Σ the set of transcripts δ such that the attacker outputs 1. If the oracle is acting like \tilde{E} then the probability that the attacker outputs 1 is exactly

$$\sum_{\delta \in \Sigma} \nu_{t(\delta), x(\delta)}(y(\delta)) .$$

If the oracle is acting like a random tweakable permutation $\tilde{\pi}$ then the probability that the attacker outputs 1 is exactly

$$\sum_{\delta \in \Sigma} \nu_{t(\delta), x(\delta)}^*(y(\delta)) .$$

We deduce that the advantage of \mathcal{A} equals:

$$\left| \sum_{\delta \in \Sigma} \left(\nu_{t(\delta), x(\delta)}(y(\delta)) - \nu_{t(\delta), x(\delta)}^*(y(\delta)) \right) \right| . \quad (8)$$

Since for every $t \in \mathcal{T}^q$, $x \in \Omega_t$, and $y \in \Omega_t$, one has $\nu_{(t,x)}(y) \geq (1 - \alpha)\nu_{(t,x)}^*(y)$, it follows that:

$$\begin{aligned} \sum_{\delta \in \Sigma} \left(\nu_{t(\delta), x(\delta)}(y(\delta)) - \nu_{t(\delta), x(\delta)}^*(y(\delta)) \right) &\geq -\alpha \\ \text{and } \sum_{\delta \notin \Sigma} \left(\nu_{t(\delta), x(\delta)}(y(\delta)) - \nu_{t(\delta), x(\delta)}^*(y(\delta)) \right) &\geq -\alpha . \end{aligned} \quad (9)$$

Finally, it is easy to verify that

$$\begin{aligned} \sum_{\delta \in \Sigma} \left(\nu_{t(\delta), x(\delta)}(y(\delta)) - \nu_{t(\delta), x(\delta)}^*(y(\delta)) \right) = \\ - \sum_{\delta \notin \Sigma} \left(\nu_{t(\delta), x(\delta)}(y(\delta)) - \nu_{t(\delta), x(\delta)}^*(y(\delta)) \right) \end{aligned} \quad (10)$$

because

$$\sum_{\delta} \nu_{t(\delta), x(\delta)}(y(\delta)) = \sum_{\delta} \nu_{t(\delta), x(\delta)}^*(y(\delta)) = 1 .$$

Using (8), (9) and (10), we deduce that the advantage of \mathcal{A} is upper bounded by α . \square

The second lemma says that if the advantage of the best NCPA adversary is small, then, for all t, x , almost all y are such that the probability of sending (t, x) to y for a random \tilde{E}_k is close to the probability of sending (t, x) to y for a random tweakable permutation.

Lemma 4. Fix $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{D})$ and $q \leq |\mathcal{D}|$. If there exists $\beta \in [0, 1]$ such that

$$\text{Adv}_{\tilde{E}}^{\text{nCPA}}(q, \tau = +\infty) \leq \beta ,$$

then, for all $t \in \mathcal{T}^q$ and $x \in \Omega_t$, one has:

$$\nu_{(t,x)}^* \left(S_{\sqrt{\beta}, (t,x)} \right) \geq 1 - \sqrt{\beta} .$$

Proof. By contrapositive, suppose there exists $t = (t_1, \dots, t_q) \in \mathcal{T}^q$ and $x = (x_1, \dots, x_q) \in \Omega_t$ such that

$$\nu_{(t,x)}^* \left(S_{\sqrt{\beta}, (t,x)} \right) < 1 - \sqrt{\beta} .$$

Consider the adversary which queries $(t_1, x_1), \dots, (t_q, x_q)$ and outputs 0 if the answers $y = (y_1, \dots, y_q)$ are such that $y \in S_{\sqrt{\beta}, (t,x)}$ and 1 otherwise. His advantage is exactly

$$\left| \sum_{y \notin S_{\sqrt{\beta}, (t,x)}} \nu_{(t,x)}(y) - \nu_{(t,x)}^*(y) \right| ,$$

and $y \notin S_{\sqrt{\beta}, (t,x)}$ means, by definition, that $\nu_{(t,x)}(y) < (1 - \sqrt{\beta})\nu_{(t,x)}^*(y)$, so that the advantage of this adversary is strictly greater than:

$$\sqrt{\beta} \times \left(1 - \nu_{(t,x)}^* \left(S_{\sqrt{\beta}, (t,x)} \right) \right) > \beta ,$$

hence the result. \square

4.2 A Composition Theorem for Tweakable Blockciphers

Given two TBCs sharing the same set of tweaks and the same domain $\tilde{E}_1 \in \text{TBC}(\mathcal{K}_1, \mathcal{T}, \mathcal{D})$ and $\tilde{E}_2 \in \text{TBC}(\mathcal{K}_2, \mathcal{T}, \mathcal{D})$, we define the tweakable blockcipher $\tilde{E}_2 \circ \tilde{E}_1 \in \text{TBC}(\mathcal{K}_1 \times \mathcal{K}_2, \mathcal{T}, \mathcal{D})$ as:

$$\begin{aligned} \forall (t, x) \in \mathcal{D} \times \mathcal{T}, (k_1, k_2) \in \mathcal{K}_1 \times \mathcal{K}_2, \\ \tilde{E}_2 \circ \tilde{E}_1((k_1, k_2), t, x) := \tilde{E}_2(k_2, t, \tilde{E}_1(k_1, t, x)) . \end{aligned}$$

Theorem 2. *Let $\tilde{E}_1 \in \text{TBC}(\mathcal{K}_1, \mathcal{T}, \mathcal{D})$ and $\tilde{E}_2 \in \text{TBC}(\mathcal{K}_2, \mathcal{T}, \mathcal{D})$ be two TBCs satisfying:*

$$\text{Adv}_{\tilde{E}_1}^{\widetilde{\text{n CPA}}} (q, \tau = +\infty) \leq \beta_1 \text{ and } \text{Adv}_{\tilde{E}_2}^{\widetilde{\text{n CPA}}} (q, \tau = +\infty) \leq \beta_2 .$$

Then:

$$\text{Adv}_{\tilde{E}_2^{-1} \circ \tilde{E}_1}^{\widetilde{\text{CCA}}} (q, \tau = +\infty) \leq 2(\sqrt{\beta_1} + \sqrt{\beta_2}) .$$

Proof. We denote ν^1 , ν^2 , and ν^3 the distributions associated respectively to \tilde{E}_1 , \tilde{E}_2 and $\tilde{E}_2^{-1} \circ \tilde{E}_1$. For every $t \in \mathcal{T}^q$, $x \in \Omega_t$, and $\alpha \in [0, 1]$, we denote $S_{\alpha, (t, x)}^{\tilde{E}_i}$ the set $S_{\alpha, (t, x)}$ corresponding to \tilde{E}_i , $i = 1, 2$.

By Lemma 4, for all $t \in \mathcal{T}^q$, $x \in \Omega_t$, and $y \in \Omega_t$, we have:

$$\nu_{(t, x)}^* \left(S_{\sqrt{\beta_1}, (t, x)}^{\tilde{E}_1} \right) \geq 1 - \sqrt{\beta_1} \quad \text{and} \quad \nu_{(t, y)}^* \left(S_{\sqrt{\beta_2}, (t, y)}^{\tilde{E}_2} \right) \geq 1 - \sqrt{\beta_2} . \quad (11)$$

Furthermore, for all $(k_1, k_2) \in \mathcal{K}_1 \times \mathcal{K}_2$, $\tilde{E}_2^{-1} \circ \tilde{E}_1((k_1, k_2), \cdot, \cdot)$ maps (t, x) to y if and only if for all $i \leq q$, $\tilde{E}_1(k_1, t_i, x_i) = \tilde{E}_2(k_2, t_i, y_i)$. Denoting $S' = S_{\sqrt{\beta_1}, (t, x)}^{\tilde{E}_1} \cap S_{\sqrt{\beta_2}, (t, y)}^{\tilde{E}_2}$, one has, for any $y \in \Omega_t$:

$$\begin{aligned} \nu_{(t, x)}^3(y) &= \sum_{z \in \Omega_t} \nu_{(t, x)}^1(z) \cdot \nu_{(t, y)}^2(z) \\ &\geq \sum_{z \in S'} \nu_{(t, x)}^1(z) \cdot \nu_{(t, y)}^2(z) \\ &\geq \sum_{z \in S'} \left(1 - \sqrt{\beta_1}\right) \nu_{(t, x)}^*(z) \cdot \left(1 - \sqrt{\beta_2}\right) \nu_{(t, y)}^*(z) \\ &\geq \left(1 - \sqrt{\beta_1}\right) \left(1 - \sqrt{\beta_2}\right) \frac{|S'|}{|\Omega_t|^2} \\ &= \left(1 - \sqrt{\beta_1}\right) \left(1 - \sqrt{\beta_2}\right) \nu_{(t, x)}^*(S') \nu_{(t, x)}^*(y) . \end{aligned}$$

By definition of S' and using Eq. 11, one has $\nu_{(t, x)}^*(S') \geq (1 - \sqrt{\beta_1} - \sqrt{\beta_2})$ (note that ν^* in fact only depends on t), so that:

$$\nu_{(t, x)}^3(y) \geq (1 - 2(\sqrt{\beta_1} + \sqrt{\beta_2})) \nu_{(t, x)}^*(y) .$$

Since this holds for any t , x , and y , the theorem follows by applying Lemma 3 with $\alpha = 2(\sqrt{\beta_1} + \sqrt{\beta_2})$. \square

4.3 Application to the r -CLRW Construction

Finally, we apply the previous result to prove the security of r -CLRW against CCA adversaries.

Theorem 3. *Let \mathcal{K}, \mathcal{T} be sets, $E \in \text{BC}(\mathcal{K}, n)$ be a blockcipher, and \mathcal{H} be an ε -AXU₂ family of functions from \mathcal{T} to $\{0, 1\}^n$. Then for any even integer r , one has:*

$$\text{Adv}_{\text{CLRW}^r, E, \mathcal{H}}^{\widetilde{\text{cca}}}(q, \tau) \leq r \cdot \text{Adv}_E^{\text{cca}}(q, \tau + rqT) + \frac{4\sqrt{2}}{\sqrt{r+2}} q^{(r+2)/4} (2\varepsilon)^{r/4} ,$$

where T is the time to compute E or E^{-1} .

Proof. Noting that the inverse of a $r/2$ -CLRW construction is again a $r/2$ -CLRW construction, we can apply Theorem 2 to get:

$$\text{Adv}_{\text{CLRW}^r, E^*, \mathcal{H}}^{\widetilde{\text{cca}}}(q, \tau = +\infty) \leq 4\sqrt{\alpha} ,$$

where

$$\alpha := \text{Adv}_{\text{CLRW}^{r/2}, E^*, \mathcal{H}}^{\widetilde{\text{ncpa}}}(q, \tau = +\infty) \leq \frac{q^{r/2+1}}{r/2+1} (2\varepsilon)^{r/2}$$

by the results of Section 3. The theorem then follows from Lemma 2. \square

Again, using an ε -AXU₂ function family with $\varepsilon \simeq 2^{-n}$, the construction achieves security against CCA adversaries up to $\mathcal{O}(2^{rn/(r+2)})$ queries.

References

1. D. J. Aldous. Random walks on finite groups and rapidly mixing Markov chains. In *Séminaire de Probabilités XVII*, volume 986 of *Lecture Notes in Mathematics*, pages 243–297. Springer, 1983.
2. A. Bogdanov, L. R. Knudsen, G. Leander, F.-X. Standaert, J. P. Steinberger, and E. Tischhauser. Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations - (Extended Abstract). In D. Pointcheval and T. Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 45–62. Springer, 2012.
3. P. Crowley. Mercy: A Fast Large Block Cipher for Disk Sector Encryption. In B. Schneier, editor, *Fast Software Encryption - FSE 2000*, volume 1978 of *Lecture Notes in Computer Science*, pages 49–63. Springer, 2000.
4. Y. Dodis and J. P. Steinberger. Domain Extension for MACs Beyond the Birthday Barrier. In K. G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 323–342. Springer, 2011.
5. S. Even and Y. Mansour. A Construction of a Cipher from a Single Pseudorandom Permutation. *Journal of Cryptology*, 10(3):151–162, 1997.
6. N. Ferguson, S. Lucks, B. Schneier, D. Whiting, M. Bellare, T. Kohno, J. Callas, and J. Walker. The Skein Hash Function Family. SHA3 Submission to NIST (Round 3), 2010.

7. D. Goldenberg, S. Hohenberger, M. Liskov, E. C. Schwartz, and H. Seyalioglu. On Tweaking Luby-Rackoff Blockciphers. In K. Kurosawa, editor, *Advances in Cryptology - ASIACRYPT 2007*, volume 4833 of *Lecture Notes in Computer Science*, pages 342–356. Springer, 2007.
8. V. T. Hoang and P. Rogaway. On Generalized Feistel Networks. In T. Rabin, editor, *Advances in Cryptology - CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 613–630. Springer, 2010.
9. D. Jetchev, O. Özen, and M. Stam. Understanding adaptivity: Random systems revisited. In X. Wang and K. Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 313–330. Springer, 2012.
10. R. Lampe, J. Patarin, and Y. Seurin. An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher. In X. Wang and K. Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 278–295. Springer, 2012.
11. W. Landecker, T. Shrimpton, and R. S. Terashima. Tweakable Blockciphers with Beyond Birthday-Bound Security. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology - CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 14–30. Springer, 2012.
12. M. Liskov, R. L. Rivest, and D. Wagner. Tweakable Block Ciphers. In M. Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2002.
13. U. M. Maurer. Indistinguishability of Random Systems. In L. R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. Springer, 2002.
14. U. M. Maurer and K. Pietrzak. Composition of Random Systems: When Two Weak Make One Strong. In M. Naor, editor, *Theory of Cryptography Conference - TCC 2004*, volume 2951 of *Lecture Notes in Computer Science*, pages 410–427. Springer, 2004.
15. U. M. Maurer, K. Pietrzak, and R. Renner. Indistinguishability Amplification. In A. Menezes, editor, *Advances in Cryptology - CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2007.
16. K. Minematsu. Beyond-Birthday-Bound Security Based on Tweakable Block Cipher. In O. Dunkelman, editor, *Fast Software Encryption - FSE 2009*, volume 5665 of *Lecture Notes in Computer Science*, pages 308–326. Springer, 2009.
17. I. Mironov. (Not So) Random Shuffles of RC4. In M. Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 304–319. Springer, 2002.
18. B. Morris, P. Rogaway, and T. Stegers. How to Encipher Messages on a Small Domain. In S. Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 286–302. Springer, 2009.
19. J. Patarin. New Results on Pseudorandom Permutation Generators Based on the DES Scheme. In J. Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 301–312. Springer, 1991.
20. P. Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In P. J. Lee, editor, *Advances in Cryptology - ASIACRYPT 2004*, volume 3329 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2004.
21. R. Schroepfel. The Hasty Pudding Cipher. AES submission to NIST, 1998.

22. V. Shoup. On Fast and Provably Secure Message Authentication Based on Universal Hashing. In N. Kobitz, editor, *Advances in Cryptology - CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 313–328. Springer, 1996.
23. J. Steinberger. Improved Security Bounds for Key-Alternating Ciphers via Hellinger Distance. IACR Cryptology ePrint Archive Report 2012/481, 2012. Available at <http://eprint.iacr.org/2012/481>.
24. K. Yasuda. The Sum of CBC MACs Is a Secure PRF. In J. Pieprzyk, editor, *Topics in Cryptology - CT-RSA 2010*, volume 5985 of *Lecture Notes in Computer Science*, pages 366–381. Springer, 2010.
25. K. Yasuda. A New Variant of PMAC: Beyond the Birthday Bound. In P. Rogaway, editor, *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 596–609. Springer, 2011.

A Proof of the Coupling Lemma

The original statement and proof of the Coupling Lemma is due to Aldous [1]. Here we follow closely a proof by Vigoda.⁵

Let λ be a coupling of μ and ν , and $(X, Y) \sim \lambda$. By definition, we have that for any $z \in \omega$, $\lambda(z, z) \leq \min\{\mu(z), \nu(z)\}$. Moreover, $\Pr[X = Y] = \sum_{z \in \Omega} \lambda(z, z)$. Hence we have:

$$\Pr[X = Y] \leq \sum_{z \in \Omega} \min\{\mu(z), \nu(z)\} .$$

Therefore:

$$\begin{aligned} \Pr[X \neq Y] &\geq 1 - \sum_{z \in \Omega} \min\{\mu(z), \nu(z)\} \\ &= \sum_{z \in \Omega} (\mu(z) - \min\{\mu(z), \nu(z)\}) \\ &= \sum_{\substack{z \in \Omega \\ \mu(z) \geq \nu(z)}} (\mu(z) - \nu(z)) \\ &= \max_{S \subset \Omega} \{\mu(S) - \nu(S)\} \\ &= \|\mu - \nu\| . \end{aligned}$$

⁵ Available from www.cc.gatech.edu/~vigoda/MCMC_Course/MC-basics.pdf.