Towards Tight Security of Cascaded LRW2

Bart Mennink

Digital Security Group, Radboud University, Nijmegen, The Netherlands
b.mennink@cs.ru.nl

Abstract. The Cascaded LRW2 tweakable block cipher was introduced by Landecker et al. at CRYPTO 2012, and proven secure up to $2^{2n/3}$ queries. There has not been any attack on the construction faster than the generic attack in 2^n queries. In this work we initiate the quest towards a tight bound. We first present a distinguishing attack in $2n^{1/2}2^{3n/4}$ queries against a generalized version of the scheme. The attack is supported with an experimental verification and a formal success probability analysis. We subsequently discuss non-trivial bottlenecks in proving tight security, most importantly the distinguisher's freedom in choosing the tweak values. Finally, we prove that if every tweak value occurs at most $2^{n/4}$ times, Cascaded LRW2 is secure up to $2^{3n/4}$ queries.

Keywords: LRW2, Cascaded LRW2, tweakable block cipher, tightness.

1 Introduction

A block cipher is a family of permutations that is indexed via a secret key. While block ciphers are omnipresent in cryptographic permutations, they inherently lack flexibility and many applications of block ciphers are either implicitly or explicitly designed from a tweakable block cipher: a function $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ that is a family of permutations indexed by secret key $k \in \mathcal{K}$ and public tweak $t \in \mathcal{T}$. Tweakable block ciphers were formalized by Liskov, Rivest, and Wagner [19] and find a broad range of applications, most notably in the direction of authenticated encryption (such as OCB [15,32,33], COPA [1], AEZ [11], and Deoxys [13,29]) and in XTS disk encryption [9].

This work centers around a generic tweakable block cipher design that was introduced in Liskov et al.'s original paper [19]. It internally uses a block cipher E, and is defined as follows:

$$\mathrm{LRW2}((k,h),t,m) = E(k, m \oplus h(t)) \oplus h(t)\,, \tag{1}$$

where k is a block cipher key and h an XOR universal hash function. The construction is strongly related with Rogaway's XEX [32] (in turn used in OCB1, OCB2, OCB3, and XTS disk encryption), and extensions by Chakraborty and Sarkar [3], Minematsu [21], and Granger et al. [10]. The LRW2 tweakable block cipher is proven to achieve security up to approximately $2^{n/2}$ queries. This bound is tight: for any two queries $(t,m),(t',m')$ with $m \oplus h(t) = m' \oplus h(t')$, the corresponding ciphertexts satisfy $c \oplus c' = h(t) \oplus h(t') = m \oplus m'$, and such a collision can be found in approximately $2^{n/2}$ queries.

A notable approach towards beyond birthday bound secure tweakable block ciphers is by Landecker et al. [17], who suggested to cascade two independent evaluations of LRW2:

$$\begin{aligned}\mathrm{CLRW2}((k_1,k_2,h_1,h_2),t,m) &= \mathrm{LRW2}((k_2,h_2),t,\mathrm{LRW2}((k_1,h_1),t,m))\,,\\ &= E_{k_2}(E_{k_1}(m\oplus h_1(t))\oplus h_1(t)\oplus h_2(t))\oplus h_2(t)\,,\end{aligned}$$

where k_1, k_2 are two block cipher keys and h_1, h_2 XOR universal hash functions. They proved that this construction is indistinguishable from random up to approximately $2^{2n/3}$ queries. This proof was very technical, and Procter [30] pointed out that it was, in fact, flawed. The proof was subsequently fixed by both Landecker et al. and Procter, but it does not generalize to higher security, either for the construction as is or for a generalization to multiple cascades. So far, there has never been any attack justifying tightness of the bound; the best attack so far is a generic one in 2^n queries.

The state of affairs stands in sharp contrast with that of two rounds of Tweakable Even-Mansour, LRW2's sibling based on public permutations [6]:

$$\mathrm{CTEM}((h_1,h_2),t,m) = p_2(p_1(m\oplus h_1(t))\oplus h_1(t)\oplus h_2(t))\oplus h_2(t)\,,$$

where p_1, p_2 are two permutations and h_1, h_2 uniform and XOR universal hash functions. Cogliati et al. [6] proved that CTEM is indistinguishable from random up to approximately $2^{2n/3}$ queries, and this bound is tight: keeping the tweak constant reduces the scheme to a key alternating cipher for which Bogdanov et al. [2] derived an attack in query complexity approximately $2^{2n/3}$. This attack uses availability of the public permutations and is therefore not applicable to CLRW2.

1.1 Attack on Generalized Cascaded LRW2

We consider a generalized version of Cascaded LRW2, for brevity called "GCL:"

$$\mathrm{GCL}^{f_1,f_2,f_3}((k_1,k_2,k_f),t,m) = E(k_2,E(k_1,m\oplus f_1(t))\oplus f_2(t))\oplus f_3(t)\,, \qquad (2)$$

where k_1, k_2 are two block cipher keys and k_f a key to the masking functions (f_1, f_2, f_3) (for ease of presentation, the key input to the f_i's is left implicit throughout). $\mathrm{GCL}^{f_1,f_2,f_3}$ is depicted in Figure 1. If h_1, h_2 are two XOR universal hash functions, then $\mathrm{GCL}^{h_1,h_1\oplus h_2,h_2}$ matches CLRW2 (where we set $k_f = (h_1, h_2)$).

We derive a generic attack against $\mathrm{GCL}^{f_1,f_2,f_3}$ with arbitrary masking in $2n^{1/2}2^{3n/4}$ evaluations. The information-theoretic attack is given in Section 3 and relies on a boomerang-style observation on the mode, based on the observation that if there exist four queries where the first and second collide on the input to E_{k_1}, the second and third on the output of E_{k_2}, and the third and fourth on the input to E_{k_1}, then the first and fourth collide at the output of E_{k_2} with probability 1 if the tweak values are selected delicately.

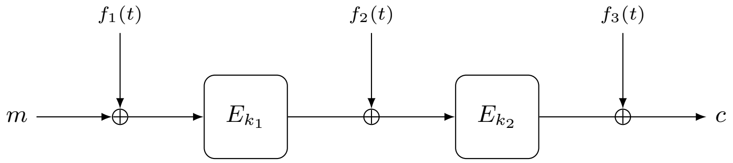

Fig. 1: Depiction of $\mathrm{GCL}^{f_1,f_2,f_3}$.

In support of its correctness, the attack is backed up with a formal success probability computation in Section 3.3 as well as an implementation in Section 3.4. The formal success analysis demonstrates that for $n \geq 27$, the distinguisher's success probability is at least $1/2$. The small-scale implementation demonstrates that for $\mathrm{GCL}^{f_1,f_2,f_3}$ based on random permutations on $n = 16, 20, 24$ bits, the special collisions as searched for in the attack indeed appear more often than usual. The gap between the accuracy in n of the experimental verification and the security proof is caused by the fact that some loose probability bounds had to be used in the rather conservative proof.

The attack is independent of the masking functions f_1, f_2, f_3. It implies that $\mathrm{GCL}^{f_1,f_2,f_3}$ cannot achieve optimal security, regardless of the choice of masking. The attack particularly applies to CLRW2, therewith improving the best known attack to date.

1.2 Towards Tight Security?

In Section 4 we approach the attack from a more theoretical perspective, and describe the main limitations in proving security of $\mathrm{GCL}^{f_1,f_2,f_3}$ beyond $2^{2n/3}$. The quasi-formal discussion relies on equating executions of $\mathrm{GCL}^{f_1,f_2,f_3}$ with a bipartite graph, and by drawing a parallel with Patarin's mirror theory [20, 22, 26, 28] we indicate various issues in trying to prove security beyond $2^{2n/3}$. The most notable one of these, namely the potential existence of four queries which alternatively collide on the input of E_{k_1} or output of E_{k_2} is precisely the one exploited in our attack in $2n^{1/2}2^{3n/4}$ queries. We also pinpoint where and how the current gap between a security lower bound of $2^{2n/3}$ and an attack upper bound of $2^{3n/4}$ arises. Most importantly, as the distinguisher can *freely choose* the value of the tweak for every query, it can set a certain distinguishing event with a significant probability.

1.3 Improved Security of Cascaded LRW2 Under Tweak Limits

In Section 5 we use these insights obtained in our quest towards tight security. We return to CLRW2, or equivalently $\mathrm{GCL}^{h_1,h_1\oplus h_2,h_2}$, and prove that if (i) h_1 and h_2 are 4-wise independent XOR universal hash functions and (ii) every tweak value occurs at most $q^{1/3}$ times, where q is the total amount of queries, then Cascaded LRW2 is secure up to $2^{3n/4}$ queries. In Section 2.2 we describe two possibilities of designing 4-wise independent XOR universal hash functions. The condition on the occurrence of the tweak seems restrictive, but many modes of operation based on a tweakable block cipher query their primitives for tweaks that are constituted of a nonce or random number concatenated with a counter value [10, 12, 15, 29]: in a nonce-respecting setting, every nonce appears at most $1 + q_f$ times, where q_f is the amount of forgery attempts.

The proof relies on Patarin's mirror theory up to the first recursion, i.e., up to $3n/4$-bit security. It shares ideas with the analysis of Mennink and Neves [20] on Encrypted Davies-Meyer [7], namely that an evaluation (t, m, c) of CLRW2 can be rewritten as a sum of permutations "in the middle." Adversarial power to choose tweak values, however, precludes optimal security, and security up to $2^{3n/4}$ is the best possible bound.

1.4 Longer Cascades?

Lampe and Seurin [16] suggested the cascade of $\rho \geq 1$ evaluations of LRW2, and proved that for even ρ this construction is secure up to approximately $2^{\rho n/(\rho+2)}$ queries. Lee et al. [18] proved that if the universal hash functions are replaced by random functions, security up to $2^{\rho n/(\rho+1)}$ is achieved. It is generally conjectured that the security of the cascade of ρ LRW2's is $2^{\rho n/(\rho+1)}$ [16–18], but also for this larger cascade, nothing is known on the attack side, besides the trivial attack in 2^n queries. Unfortunately, it does not seem possible to generalize the attack of Section 3 nor the security proof of Section 5 to larger cascades. As before, it is noteworthy that a cascade of $\rho \geq 1$ evaluations of TEM can be attacked in approximately $2^{\rho n/(\rho+1)}$ queries [2].

2 Preliminaries

For $n \in \mathbb{N}$, $\{0,1\}^n$ denotes the set of bit strings of length n, and $\mathsf{perm}(n)$ the set of all permutations on $\{0,1\}^n$. Extending notation, for $\kappa \in \mathbb{N}$, we denote by $\mathsf{iperm}(\kappa, n)$ the set of all "indexed permutations," families of permutations $p_k \in \mathsf{perm}(n)$, indexed by $k \in \{0,1\}^\kappa$. We additionally denote by $\mathsf{iperm}(\kappa, \tau, n)$ for $\tau \in \mathbb{N}$ the set of all indexed permutations where the index consists of two elements $(k, t) \in \{0,1\}^\kappa \times \{0,1\}^\tau$. For $m, n \in \mathbb{N}$ such that $m \geq n$, the falling factorial is defined as $(m)_n = m(m-1)\cdots(m-n+1) = m!/(m-n)!$. For $n \in \mathbb{N}$ and $m \in \{0, \ldots, 2^{n-1}\}$, we denote by $\langle m \rangle_n$ the encoding of m as an n-bit string. If $\mathcal{X}$ is a finite set, $x \xleftarrow{\$} \mathcal{X}$ denotes the event of uniformly randomly drawing x from $\mathcal{X}$.

2.1 Block Ciphers and Tweakable Block Ciphers

A block cipher with key size κ and state size n is a function $E \in \mathsf{iperm}(\kappa, n)$. For fixed key $k \in \{0,1\}^\kappa$ we denote $E_k(\cdot) = E(k, \cdot)$, and its inverse is denoted $E_k^{-1}(\cdot)$. A tweakable block cipher with key size κ, tweak size τ, and state size n is a function $\widetilde{E} \in \mathsf{iperm}(\kappa, \tau, n)$. For fixed key $k \in \{0,1\}^\kappa$ and $t \in \{0,1\}^\tau$ we denote $\widetilde{E}_k(t, \cdot) = \widetilde{E}(k, t, \cdot)$, and its inverse is denoted $\widetilde{E}_k^{-1}(t, \cdot)$.

Let $\kappa, n \in \mathbb{N}$ and let $E \in \mathsf{iperm}(\kappa, n)$ be a block cipher. The advantage of a distinguisher $\mathcal{D}$ in breaking the SPRP (strong pseudorandom permutation) security of E is defined as

$$\mathbf{Adv}_{E}^{\mathrm{sprp}}(\mathcal{D}) = \mathbf{Pr}\left(\mathcal{D}^{E_k^{\pm}} = 1\right) - \mathbf{Pr}\left(\mathcal{D}^{p^{\pm}} = 1\right) , \tag{3}$$

where the probabilities are taken over the random drawing of $k \xleftarrow{\$} \{0,1\}^{\kappa}$, $p \xleftarrow{\$} \mathsf{perm}(n)$, and the randomness used by $\mathcal{D}$. The resources that $\mathcal{D}$ may use are typically expressed in terms of query complexity (to the oracle) and time complexity (for offline computations).

As block ciphers are a special case of tweakable block ciphers with tweak space of size 1 ($\tau = 0$), the security definition straightforwardly generalizes to the latter. Let $\kappa, \tau, n \in \mathbb{N}$ and let $\widetilde{E} \in \mathsf{iperm}(\kappa, \tau, n)$ be a tweakable block cipher. The advantage of a distinguisher $\mathcal{D}$ in breaking the STPRP (strong tweakable pseudorandom permutation) security of $\widetilde{E}$ is defined as

$$\mathbf{Adv}_{\widetilde{E}}^{\mathrm{stprp}}(\mathcal{D}) = \mathbf{Pr}\left(\mathcal{D}^{\widetilde{E}_k^{\pm}} = 1\right) - \mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}^{\pm}} = 1\right) , \tag{4}$$

where the probabilities are taken over the random drawing of $k \xleftarrow{\$} \{0,1\}^{\kappa}$, $\widetilde{p} \xleftarrow{\$} \mathsf{iperm}(\tau, n)$, and the randomness used by $\mathcal{D}$. The resources that $\mathcal{D}$ may use are typically bounded as before.

2.2 XOR Universal Hash Functions

We use the notion of ℓ-wise independent XOR universal hash functions, a slight adaptation of the original definition of Wegman and Carter [34]. For two non-empty sets $\mathcal{X}, \mathcal{Y}$, a hash function family $H = \{h : \mathcal{X} \to \mathcal{Y}\}$ is called ℓ-wise independent almost XOR universal up to bound ε, denoted ε-AXU$_\ell$, if for any $j \in \{2, \ldots, \ell\}$, any distinct $x_1, \ldots, x_j \in \mathcal{X}$ and (not necessarily distinct) $y_2, \ldots, y_j \in \mathcal{Y}$,

$$\mathbf{Pr}\left(h \xleftarrow{\$} H \;:\; h(x_1) \oplus h(x_2) = y_2 \,,\; \ldots \,,\; h(x_1) \oplus h(x_j) = y_j\right) \leq \varepsilon^{j-1} \,.$$

For $\mathcal{X} = \mathcal{Y} = \{0,1\}^n$, a 2^{-n}-AXU$_2$ hash function family can be defined using finite field multiplication with respect to some irreducible polynomial to represent the field, i.e., $h(x) := h \otimes x$. It is not ε-AXU$_\ell$ for $\ell > 2$. Defining the hash function family as

$$\boldsymbol{h}(x) := \bigoplus_{i=1}^{\ell-1} h_i \otimes x^i$$

for $\boldsymbol{h} = (h_1, \ldots, h_{\ell-1})$ gives a 2^{-n}-AXU$_\ell$ hash function family for any $\ell \geq 2$. One can alternatively obtain a $(2^n - (\ell-1))^{-1}$-AXU$_\ell$ by defining the hash function family using an ideal cipher or a family of random permutations.

3 Generic Attack

We present a generic attack against $\mathrm{GCL}^{f_1,f_2,f_3}$ in $2n^{1/2}2^{3n/4}$ queries. The attack is generic in nature, it does not exploit any weaknesses in the underlying cipher, and as such we simply assume that $E \xleftarrow{\$} \mathsf{iperm}(\kappa, n)$ is an ideal cipher. It is fair to assume that the success probability of the attack simply *improves* if E is less than ideal, except for degenerate cases, e.g., if E_{k_1} and E_{k_2} are almost perfect nonlinear permutations (APNPs, cf., [8,23,24]). Throughout the attack, we simply denote $p_1 = E_{k_1}$ and $p_2 = E_{k_2}$ for brevity.

An informal rationale of our attack is given in Section 3.1, and the formal distinguisher in Section 3.2. Its advantage is lower bounded in Section 3.3, and the analysis is backed up with experimental verification in Section 3.4.

3.1 Informal Rationale of Attack

Suppose a distinguisher obtains four queries (t, m_1, c_1), (t', m_2', c_2'), (t, m_3, c_3), and (t', m_4', c_4') of $\mathrm{GCL}^{f_1,f_2,f_3}$ such that

$$\begin{aligned} m_1 \oplus f_1(t) &= m_2' \oplus f_1(t')\,, \\ c_2' \oplus f_3(t') &= c_3 \oplus f_3(t)\,, \\ m_3 \oplus f_1(t) &= m_4' \oplus f_1(t')\,. \end{aligned} \tag{5}$$

In other words, the first and second query collide at the input to E_{k_1}, the second and third at the output of E_{k_2}, and the third and fourth at the input to E_{k_1}. As the four queries are performed using only two tweak values, each occurring twice, we have $f_2(t) \oplus f_2(t') \oplus f_2(t) \oplus f_2(t') = 0$, and from a simple inspection of the scheme (see also Figure 2) one can conclude that, necessarily,

$$c_1 \oplus f_3(t) = c_4' \oplus f_3(t')\,. \tag{6}$$

Stated differently, under the assumption that (5) is satisfied, (6) is implied, and therefore the four equations combine to

$$\begin{aligned} m_1 \oplus m_2' = m_3 \oplus m_4' &= f_1(t) \oplus f_1(t')\,, \\ c_2' \oplus c_3 = c_1 \oplus c_4' &= f_3(t) \oplus f_3(t')\,. \end{aligned}$$

Unfortunately, the distinguisher does not know $f_1(t) \oplus f_1(t')$ and $f_3(t) \oplus f_3(t')$, but if we ignore these two values in above equations, we obtain

$$\begin{aligned} m_1 \oplus m_2' &= m_3 \oplus m_4'\,, \\ c_2' \oplus c_3 &= c_1 \oplus c_4'\,, \end{aligned} \tag{7}$$

which *necessarily* holds if $m_1 \oplus m_2' = f_1(t) \oplus f_1(t')$ and $c_2' \oplus c_3 = f_3(t) \oplus f_3(t')$, but may hold by accident as well. Stated differently, if for some $d \in \{0,1\}^n$, there are about 2^n choices for the four queries such that

$$m_1 \oplus m_2' = m_3 \oplus m_4' = d\,, \tag{8}$$

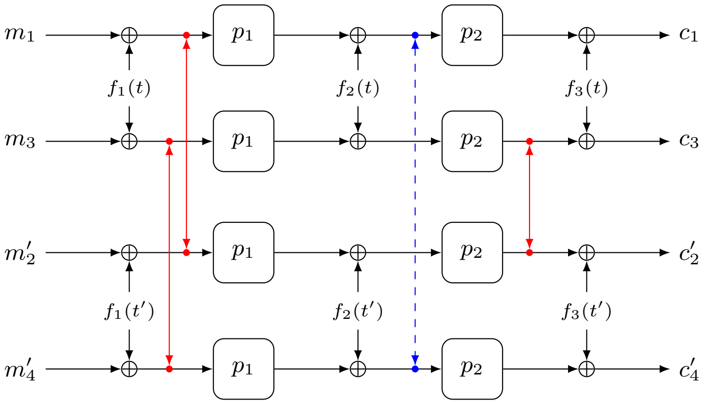

Fig. 2: Attack idea: the red (solid) collisions are targeted, the blue (dashed) one is implied by the red ones.

the expected number of solutions to (7) is close to 2 if $d = f_1(t) \oplus f_1(t')$ but close to 1 if $d \neq f_1(t) \oplus f_1(t')$. For an ideal permutation, the expected number of solutions is always close to 1 for any $d \in \{0,1\}^n$. By making approximately $2^{3n/4}$ queries, the distinguisher can ensure that there are about 2^n solutions to (8) for all d, including $d = f_1(t) \oplus f_1(t')$.

This almost allows for a distinguishing attack, but not quite: as the distinguisher does not actually know $f_1(t) \oplus f_1(t')$, it must simply hope that for some d there is a significant difference, but d may take 2^n values and false positives are likely to occur. By extending the number of queries slightly, i.e., by making about $n^{1/2} \cdot 2^{3n/4}$ queries, the case of $f_1(t) \oplus f_1(t')$ will stand out.

We remark that the attack is effectively an XOR subkey recovery attack, as the distinguisher learns $f_1(t) \oplus f_1(t')$ and $f_3(t) \oplus f_3(t')$. In case of Cascaded LRW2, where $f_1 = h_1$, $f_2 = h_1 \oplus h_2$, and $f_3 = h_2$ for two XOR universal hash functions h_1, h_2, this immediately gives $f_2(t) \oplus f_2(t')$, and potentially more, depending on the specific hash functions.

3.2 Formal Description of Distinguisher

Let $\epsilon = \log_2(n)/2$ (assumed to be integral), and consider the following distinguisher $\mathcal{D}$ making $q = 2^{3n/4+\epsilon}$ queries.

(i) Fix arbitrary distinct $t, t' \in \{0,1\}^\tau$;
(ii) For $i = 0, \dots, 2^{3n/4+\epsilon} - 1$, put $m_i = 0^{n/4-\epsilon} \| \langle i \rangle_{3n/4+\epsilon}$ and query (t, m_i) to obtain c_i;
(iii) For $i = 0, \dots, 2^{3n/4+\epsilon} - 1$, put $m'_i = \langle i \rangle_{3n/4+\epsilon} \| 0^{n/4-\epsilon}$ and query (t', m'_i) to obtain c'_i;
(iv) For $d \in \{0,1\}^n$, define $I_d = \{(i,j) \mid m_i \oplus m'_j = d\}$. Note that $|I_d| = 2^{n/2+2\epsilon}$ for all $d \in \{0,1\}^n$, and define $q' := 2^{n/2+2\epsilon}$;
(v) For all $d \in \{0,1\}^n$ do:

- Define $N_d = 0$;
- For all distinct $(i,j),(k,l) \in I_d$: if $c_i \oplus c'_l = c'_j \oplus c_k$, put $N_d = N_d + 1$;

(vi) Briefly looking forward, for a random tweakable block cipher we have $\mathbf{Ex}\left(N_d\right) = \binom{q'}{2}/(2^n-1)$ for any $d \in \{0,1\}^n$, whereas for $\mathrm{GCL}^{f_1,f_2,f_3}$, $\mathbf{Ex}\left(N_{f_1(t)\oplus f_1(t')}\right) \geq 2\binom{q'}{2}/2^n$. Inspired by this, define

$$\beta := \frac{3}{2}\binom{q'}{2}/2^n\,.$$

If there exists a $d \in \{0,1\}^n$ such that $N_d \geq \beta$, output 1. Otherwise, output 0.

3.3 Analysis of Distinguisher Advantage

A formal analysis confirms that the distinguisher succeeds with non-negligible probability.

Theorem 1. *Let $\kappa, \tau, n \in \mathbb{N}$ with $n \geq 16$, let $E \xleftarrow{\$} \mathsf{iperm}(\kappa, n)$, denote the size of the key space of (f_1, f_2, f_3) by κ_f, and consider $\mathrm{GCL}^{f_1,f_2,f_3} : \{0,1\}^{2\kappa} \times \{0,1\}^{\kappa_f} \times \{0,1\}^{\tau} \times \{0,1\}^n \to \{0,1\}^n$. Distinguisher $\mathcal{D}$ of Section 3.2 with query complexity $2n^{1/2} \cdot 2^{3n/4}$ has advantage*

$$\mathbf{Adv}^{\mathrm{stprp}}_{\mathrm{GCL}^{f_1,f_2,f_3}}(\mathcal{D}) \geq 1 - \frac{32}{n^2} - \frac{80}{n2^{n/2}} - 5 \cdot 2^n \left(\frac{10}{n}\right)^{3/100\cdot n^2} - \frac{n^7}{2^{3n/2}}\,. \tag{9}$$

One can verify that the lower bound of (9) is at least $1/2$ for $n \geq 27$. This theorem is not the core contribution of the article (which is Theorem 2), and its proof is given Appendix A.

Note that the attack is de facto a TPRP-attack, only requiring forward access to the scheme. In addition, it is information-theoretical: the distinguisher's complexity is solely measured in its number of queries. The offline complexity is around $2^{3n/2}$.

3.4 Experimental Verification

We have implemented the distinguisher of Section 3.2 on a small scale, for $n = 16, 20, 24$ and with p_1, p_2, f_1, f_2, f_3 instantiated as independent uniform random permutations, noting that a uniform random permutation is a $(2^n - 1)^{-1}$-AXU$_2$ hash function (see Section 2.2). In each case, two distinct tweaks t, t' are evaluated for $q = 2^{3n/4+\epsilon}$ queries, with $\epsilon = 0, 1, 2$ (note that $2 \lesssim \log_2(n)/2$ for $n = 16, 20, 24$). The average values N_d for both the real and ideal world and both $d = f_1(t) \oplus f_1(t')$ and random d are summarized in Table 1. The computations confirm soundness of the rationale of Section 3.1 and the expected values of Section 3.2. In more detail, the expected values given in Section 3.2 suggest that $N_d \approx 2^{4\epsilon}$ for $d = f_1(t) \oplus f_1(t')$ in the real world and $N_d \approx 2^{4\epsilon-1}$ in any other

Table 1: Number of elements in N_d for the real and ideal world, for $d = f_1(t) \oplus f_1(t')$ and for random d. For the cases $n = 16, 20$, the numbers are averaged over 32 attacks; for $n = 24$ the numbers are averaged over 8 attacks.

			N_d in real world for $d =$		N_d in ideal world for $d =$	
n	ϵ	q	$f_1(t) \oplus f_1(t')$	random	$f_1(t) \oplus f_1(t')$	random
16	0	$1 \cdot 2^{12}$	0.843750	0.437500	0.343750	0.687500
	1	$2 \cdot 2^{12}$	16.343750	6.656250	7.625000	8.500000
	2	$4 \cdot 2^{12}$	256.593750	129.781250	127.093750	127.375000
20	0	$1 \cdot 2^{15}$	0.968750	0.500000	0.687500	0.593750
	1	$2 \cdot 2^{15}$	17.156250	7.593750	8.343750	8.187500
	2	$4 \cdot 2^{15}$	265.531250	133.312500	125.625000	128.750000
24	0	$1 \cdot 2^{18}$	1.125000	0.875000	0.250000	0.125000
	1	$2 \cdot 2^{18}$	16.375000	7.625000	8.375000	7.125000
	2	$4 \cdot 2^{18}$	246.750000	131.375000	120.625000	129.875000

case (real or ideal world), and the statistics in Table 1 reasonably accurately match these numbers.

Note that, in particular, for $\epsilon = 0$ the value $N_{f_1(t) \oplus f_1(t')}$ already shows a small peak in the real world (for each of $n = 16, 20, 24$), but outliers in N_d for $d \neq f_1(t) \oplus f_1(t')$ are hidden by the statistics. For increasing ϵ, the gap becomes more significant and the success probability increases.

4 Towards Tight Security?

Consider a simplification of $\mathrm{GCL}^{f_1,f_2,f_3}$ with its two block ciphers replaced by random permutations p_1, p_2 (this is a typical hybrid argument in security proofs performed at the cost of $2\mathbf{Adv}_E^{\mathrm{sprp}}(\mathcal{D}')$ for some distinguisher $\mathcal{D}'$). For simplicity, assume that f_2 is injective (the scheme turns out to be significantly weakened if f_2 is non-injective). For an evaluation $\mathrm{GCL}^{f_1,f_2,f_3}(t, m) = c$, denote

$$x = p_1(m \oplus f_1(t)) \,,$$
$$y = p_2^{-1}(c \oplus f_3(t)) \,,$$

in such a way that $x \oplus y = f_2(t)$.

Intuitively, one may think of a proof going “fine” if there is always some randomness available. For example, consider just a single forward query (t, m) to $\mathrm{GCL}^{f_1,f_2,f_3}$. The value $m \oplus f_1(t)$ has never been evaluated by p_1, hence the value x will look uniformly randomly drawn from $\{0,1\}^n$; the value y satisfies $y = x \oplus f_2(t)$, and also y has never been evaluated by p_2 so the value $c \oplus f_3(t)$ is uniformly randomly drawn from $\{0,1\}^n$.

A more complicated case appears if there exist two distinct queries (m_1, t_1) and (m_2, t_2) such that $m_1 \oplus f_1(t_1) = m_2 \oplus f_1(t_2)$. The first query is handled as before, rendering fresh x_1 and $c_1 \oplus f_3(t_1)$. The second query satisfies $m_1 \oplus f_1(t_1) = m_2 \oplus f_1(t_2)$, meaning that $x_2 = x_1$. However, as the two queries are distinct, this equation implies that $t_1 \neq t_2$. As f_2 is injective, we subsequently have $f_2(t_1) \neq f_2(t_2)$ and thus $y_2 \neq y_1$. The evaluation of p_2 on y_2 yields a value uniformly drawn from $\{0,1\}^n \backslash \{c_1 \oplus f_3(t_1)\}$.

Likewise, two queries could also collide at the right side, i.e., $c_1 \oplus f_3(t_1) = c_2 \oplus f_3(t_2)$. It is unlikely, though, that two queries collide at *both* the left and right side, at least if f_1 and f_3 are two randomized functions (as is the case in CLRW2), and we will ignore this case. If more than two queries are involved, one could visualize queries as a bipartite graph $G = (U, V, E)$. $U = \{0,1\}^n$ corresponds to the input values to p_1, $V = \{0,1\}^n$ to the output values of p_2, and for every query tuple (t_i, m_i, c_i), the edge $(m_i \oplus f_1(t_i), c_i \oplus f_3(t_i))$ with label $f_2(t_i)$ from U to V is added to E. An example graph G is depicted in Figure 3.

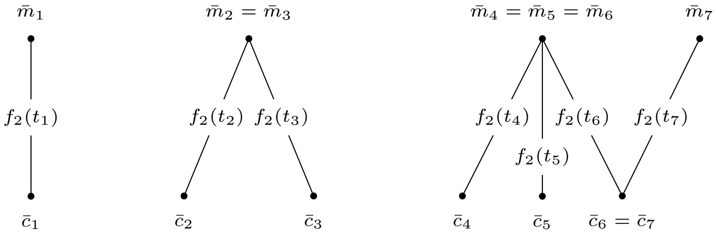

Fig. 3: Example of a bipartite graph G representing seven evaluations of $\mathrm{GCL}^{f_1,f_2,f_3}$. For brevity, we denote $\bar{m}_i = m_i \oplus f_1(t_i)$ and $\bar{c}_i = c_i \oplus f_3(t_i)$. Graph view rotated for economical reasons.

What the above comprises is an informal introduction to a potential use of Patarin's mirror theory [20,22,26,28], a powerful approach towards counting the number of solutions to a system of equations of the form $x \oplus y = \lambda$, where λ is known. If, in above graph, two queries touch on the left, i.e., $m_1 \oplus f_1(t_1) = m_2 \oplus f_1(t_2)$, they share the same $x_1 = x_2$ but have different y_1, y_2.

Unfortunately, the mirror theory does not turn out to be particularly suited here, most importantly as it is tailored towards comparing systems to random functions and we aim to compare our scheme to a family of permutations. Yet, closer inspection of the theory reveals that it puts two conditions on the graph that are "reasonably easily" violated:

(i) The graph should not contain a path of even length whose labels sum to 0;
(ii) The graph should not contain a circle.

The first condition prevents that there are two different inputs to p_1 with the same output (or two different outputs of p_2 with the same input). The second condition prevents that there exists a query with "no randomness." Both conditions are harmful for any possible even length, in the sense that Patarin's mirror theorem cannot be applied.

The attack of Section 3 relies on the fact that condition (i) can be violated easier than expected. Note that there cannot exist a path of length 2 whose labels sum to 0 (as f_2 is injective). A path of length 4 whose labels sum to 0 requires the existence of four queries $(t_1, m_1, c_1), \ldots, (t_4, m_4, c_4)$ such that

$$\begin{aligned} m_1 \oplus f_1(t_1) &= m_2 \oplus f_1(t_2)\,, \\ c_2 \oplus f_3(t_2) &= c_3 \oplus f_3(t_3)\,, \\ m_3 \oplus f_1(t_3) &= m_4 \oplus f_1(t_4)\,, \\ f_2(t_1) \oplus f_2(t_2) \oplus f_2(t_3) \oplus f_2(t_4) &= 0\,. \end{aligned} \tag{10}$$

As the four queries are distinct, the path may only appear if $t_1 \neq t_2 \neq t_3 \neq t_4$. However, it may be that $t_1 = t_3$ and $t_2 = t_4$, and this is how the attack of Section 3 exploits a path: in this case, the fourth equation of (10) is satisfied by design and the remaining three can be rewritten as

$$\begin{aligned} m_1 \oplus m_2 = m_3 \oplus m_4 &= f_1(t_1) \oplus f_1(t_2)\,, \\ c_2 \oplus c_3 &= f_3(t_1) \oplus f_3(t_2)\,. \end{aligned} \tag{11}$$

The attack of Section 3 relies on the additional fact that if these conditions are met, then the condition

$$c_4 \oplus f_3(t_2) = c_1 \oplus f_3(t_1) \tag{12}$$

holds with probability 1 in the real world (i.e., there is a circle as depicted in Figure 4, violating condition (ii)), but with negligible probability in the ideal world. This property (that (11) implies (12)) gives a *clean and well-verifiable distinguishing event.*

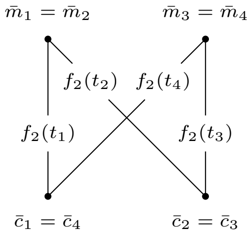

Fig. 4: A circle in bipartite graph G with $f_2(t_1) \oplus f_2(t_2) \oplus f_2(t_3) \oplus f_2(t_4) = 0$, as exploited in the attack of Section 3. We use the same convention as in Figure 3.

A distinguisher can choose the m_i's smartly to make sure that $m_1 \oplus m_2 = m_3 \oplus m_4$ is satisfied. Consider a distinguisher that makes queries for at most two tweaks t, t', each queried q times, say for queries $(m_0, c_0), \ldots, (m_{q-1}, c_{q-1})$ and $(m'_0, c'_0), \ldots, (m'_{q-1}, c'_{q-1})$. Inspired by Section 3, denote

$$I_d = \{(i, j) \in \{0, \ldots, q-1\}^2 \mid m_i \oplus m'_j = d\}\,.$$

The probability that there exist four queries $(i,j) \neq (i',j')$ that comply with the equations of (11), denoted X, is

$$\begin{aligned}
\mathbf{Pr}\left(X\right) &= \sum_{d\in\{0,1\}^n} \mathbf{Pr}\left(X \mid f_1(t_1) \oplus f_1(t_2) = d\right) \cdot \mathbf{Pr}\left(f_1(t_1) \oplus f_1(t_2) = d\right) \\
&\approx \sum_{d\in\{0,1\}^n} \frac{\binom{|I_d|}{2}}{2^n} \cdot \mathbf{Pr}\left(f_1(t_1) \oplus f_1(t_2) = d\right) \\
&\approx \sum_{d\in\{0,1\}^n} \frac{\binom{|I_d|}{2}}{2^n} \cdot \frac{1}{2^n}\,, \qquad (13)
\end{aligned}$$

where the first approximation assumes independence of events and that the c_i's are generated using a random function (for simplicity of reasoning), and the second approximation assumes that f_1 is close to a 2^{-n}-AXU$_2$ hash function. The two extremes in selecting the m_i's are the following:

- Choose the m_i's and m'_i's such that $|I_d| = q$ for q values of d and $|I_d| = 0$ for the remaining $2^n - q$ values. This is achieved by setting $m_i = m'_i = 0^{n-\log_2(q)} \| \langle i \rangle_{\log_2(q)}$ for $i = 0,\ldots,q-1$. In this case, we obtain for (13):
$$(13) = q \cdot \binom{q}{2} / 2^{2n} \approx q^3/2^{2n}\,;$$
- Choose the m_i's and m'_i's such that $|I_d| = q^2/2^n$ for all values of d, i.e., I_d is equally large for all d. This is achieved by setting $m_i = 0^{n-\log_2(q)} \| \langle i \rangle_{\log_2(q)}$ and $m'_i = \langle i \rangle_{\log_2(q)} \| 0^{n-\log_2(q)}$ for $i = 0,\ldots,q-1$ (as in the attack of Section 3). In this case, we obtain for (13):
$$(13) = 2^n \cdot \binom{q^2/2^n}{2} / 2^{2n} \approx q^4/2^{3n}\,.$$

A security analysis, i.e., an upper bound on the distinguisher's success probability, would have to take into account any possible distinguisher, and it therefore seems such analysis caps at around $q^3/2^{2n}$. Yet, if the attack of Section 3 would have been based on the former strategy instead of the latter, it would have succeeded *only* if $|I_{f_1(t_1)\oplus f_1(t_2)}| \neq 0$, and the attack should have been evaluated $2^n/q$ times to succeed (resulting in total complexity of about 2^n). By making $2^{3n/4}$ queries, the distinguisher makes sure that $|I_d|$ is equally large for all d's and that way spreads its chances, but unfortunately, we see little opportunities in improving the attack.

It is important to remark that the attack of Section 3 and the discussion on the distinguishing event (11) consider the case where the distinguisher can *choose* the tweak values. This implies that an improved security bound can be achieved if the maximum number of queries for each tweak is fixed.

We explicitly remark that this limitation is *not a necessary condition.* In particular, above reasoning is informal and only included for intuitive reasons, and we cannot draw any formal conclusion from it. However, even for this limited scenario, improved security of CLRW2 is still a non-trivial open problem. We elaborate on the possibility of releasing the tweak usage limitation in Section 5.7.

A final condition that the mirror theory puts on the graph, in addition to (i) and (ii) above, is the following:

(iii) The graph should not contain an excessively large tree.

This is a merely technical requirement to make the proof argument of the mirror theory work, and it is not clear how a violation of condition (iii) may break the scheme. That said, also condition (iii) can be easily violated, depending on the mixing functions in use. For example, if $f_1(t) = f_1 \otimes t$ (i.e., the example AXU_2 hash function of Section 2), a collision of the form

$$m_1 \oplus f_1(m_1) = m_2 \oplus f_1(m_2)\,,$$

for $m_1, m_2 \neq 0$ implies that also

$$m_2 \oplus f_1(m_2) = m_1^{-1}m_2^2 \oplus f_1(m_1^{-1}m_2^2) = \cdots = m_1^{-\lambda}m_2^{\lambda+1} \oplus f_1(m_1^{-\lambda}m_2^{\lambda+1})\,,$$

for any $\lambda \geq 0$, potentially rendering an excessively large tree. The issue can be resolved by resorting to 4-wise independent XOR universal hash functions (see Section 2.2).

5 Improved Security of Cascaded LRW2 Under Tweak Limits

Based on the two conclusions from Section 4, we prove that if h_1 and h_2 are two 4-wise independent XOR universal hash functions and every tweak occurs at most $q^{1/3}$ times, the Cascaded LRW2 construction $\mathrm{GCL}^{h_1,h_1\oplus h_2,h_2}$ of (2) achieves security up to complexity approximately $2^{3n/4}$.

Theorem 2. *Let $\kappa, \tau, n \in \mathbb{N}$, let $E \in \mathsf{iperm}(\kappa, n)$, H be an ε-AXU_4 hash function family, and consider $\mathrm{GCL}^{h_1,h_1\oplus h_2,h_2} : \{0,1\}^{2\kappa} \times H^2 \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$. Let $\gamma \in \mathbb{N}$ such that $2 \leq \gamma \leq q/4$ be a threshold. For any distinguisher $\mathcal{D}$ with query complexity at most $q \leq 2^n/1600$ that queries each tweak at most γ times, there exists a distinguisher $\mathcal{D}'$ that makes at most q queries such that*

$$\mathbf{Adv}^{\mathrm{stprp}}_{\mathrm{GCL}^{h_1,h_1\oplus h_2,h_2}}(\mathcal{D}) \leq 6\binom{q}{4}2^n\varepsilon^4 + \binom{q}{2}(2\gamma+1)\varepsilon^2 + \frac{(\gamma+3)q}{2^n} + 2\mathbf{Adv}^{\mathrm{sprp}}_E(\mathcal{D}')\,. \tag{14}$$

Putting $\gamma = q^{1/3}$, the bound of (14) yields security up to $q \leq 2^{3n/4}$ queries. The limitation γ on the number of tweak repeats sounds restrictive, but it is not. In practical applications [10, 12, 29], the tweak is constituted of a random value concatenated with a counter.

The proof of Theorem 2 is based on Patarin's mirror theory [22, 26, 28], which found popularization in the work of Mennink and Neves on Encrypted Davies-Meyer and its dual [20]. Although the mirror theory is quite simple to understand and apply, its proof is heavy and the recursive argument underneath it is debated by some. In this work, however, we will only use the mirror theory up to $3n/4$-bit security, i.e., rely on the first recursion in the mirror theory proof only.

The security proof is comparable to that of EDM [20], and in particular also relies on the observation that any evaluation of $c = \mathrm{GCL}^{h_1, h_1 \oplus h_2, h_2}(\mathbf{k}, t, m)$ for $\mathbf{k} = (k_1, k_2, h_1, h_2)$ can be rewritten as

$$E_{k_1}(m \oplus h_1(t)) \oplus E_{k_2}^{-1}(c \oplus h_2(t)) = h_1(t) \oplus h_2(t) \,. \tag{15}$$

Differences in the analysis occur due to the possibility of the adversary to choose the tweak and the fact that the tweak occurs in all three parts of the equation (input to E_{k_1}, to $E_{k_2}^{-1}$, and in the right hand side $h_1(t) \oplus h_2(t)$). These differences cause that only security up to $2^{3n/4}$ is achievable. However, the differences compared with the analysis in [20] mostly affect description of oracle views and analysis of bad views; the application of the mirror theory is fairly the same. Therefore, we discard much of the details on mirror theory from the proof and include it in Appendix B; the proof is fully intelligible without this appendix.

The proof is given in Sections 5.1-5.6. We discuss the possibility of releasing the limitation γ on the tweak usage in Section 5.7.

5.1 H-Coefficient Technique

We will use Patarin's H-coefficient technique [25,27], for which we follow the description by Chen and Steinberger [5]. Consider two oracles $\mathcal{O}$ and $\mathcal{P}$ with identical interfaces, and a deterministic distinguisher $\mathcal{D}$ with query complexity q and unbounded computational power that tries to distinguish both oracles. Denote its success probability by $\Delta_{\mathcal{D}}(\mathcal{O}\,;\,\mathcal{P})$. Let $X_{\mathcal{O}}$ denote the probability distribution of views when $\mathcal{D}$ is interacting with $\mathcal{O}$, and similarly $X_{\mathcal{P}}$ the distribution of views for interaction with $\mathcal{P}$. A view ν is called "attainable" if $\mathbf{Pr}\left(X_{\mathcal{P}} = \nu\right) > 0$, and denote by $\mathcal{V}$ the set of all attainable views. The H-coefficient technique states the following:

Lemma 1 (H-coefficient technique). *Let $\mathcal{D}$ be a deterministic distinguisher, and consider a partition $\mathcal{V} = \mathcal{V}_{\mathrm{bad}} \cup \mathcal{V}_{\mathrm{good}}$ of the set of attainable views. Let $\delta, \epsilon \in [0,1]$ be such that $\mathbf{Pr}\left(X_{\mathcal{P}} \in \mathcal{V}_{\mathrm{bad}}\right) \leq \delta$, and $\dfrac{\mathbf{Pr}\left(X_{\mathcal{O}} = \nu\right)}{\mathbf{Pr}\left(X_{\mathcal{P}} = \nu\right)} \geq 1 - \epsilon$ for all $\nu \in \mathcal{V}_{\mathrm{good}}$. Then, the distinguishing advantage satisfies $\Delta_{\mathcal{D}}(\mathcal{O}\,;\,\mathcal{P}) \leq \delta + \epsilon$.*

A proof of the technique is given among others in [4,5,20].

For view $\nu = \{(x_1, y_1), \ldots, (x_q, y_q)\}$ consisting of q input/output tuples, an oracle $\mathcal{O}$ is said to *extend* ν, denoted $\mathcal{O} \vdash \nu$, if $\mathcal{O}(x_i) = y_i$ for all $i = \{1, \ldots, q\}$.

5.2 General Setting and Views

Let $\widetilde{p} \xleftarrow{\$} \mathsf{iperm}(\tau, n)$, $\mathbf{k} \xleftarrow{\$} \{0,1\}^{2\kappa} \times H^2$, and $p_1, p_2 \xleftarrow{\$} \mathsf{perm}(n)$. Consider any distinguisher $\mathcal{D}$ whose goal is to distinguish $\mathrm{GCL}_{\mathbf{k}}^{h_1, h_1 \oplus h_2, h_2}$ from $\widetilde{p}$.

As a first step, we replace (E_{k_1}, E_{k_2}) by (p_1, p_2^{-1}) at the cost of $2\mathbf{Adv}_E^{\mathrm{sprp}}(\mathcal{D}')$, where $\mathcal{D}'$ is some distinguisher with the same query complexity q as $\mathcal{D}$. (Note that we replaced E_{k_2} by *the inverse* of p_2 for simplicity of further analysis.) Denote the resulting scheme with F for brevity; it remains to bound the advantage of $\mathcal{D}$ in distinguishing $\mathcal{O} = F$ (the real world) from $\mathcal{P} = \widetilde{p}$ (the ideal world). As of now, we give the distinguisher unbounded computational power, and its complexity will only be measured by the number of oracle queries it makes. Without loss of generality, we can consider it to be deterministic, and will apply the H-coefficient technique of Lemma 1.

$\mathcal{D}$ makes q construction queries which are recorded in view $\nu' = \{(t_1, m_1, c_1), \dots, (t_q, m_q, c_q)\}$. After $\mathcal{D}$'s interaction with its oracle, but before it outputs its decision bit, its oracle will reveal the subkeys h_1, h_2. In the real world, these are the XOR universal hash functions used in F, whereas in the ideal world these are dummy functions randomly drawn from H. We denote the complete view by

$$\nu = (\nu', h_1, h_2)\,. \tag{16}$$

Without loss of generality, we assume that $\mathcal{D}$ never repeats queries, and hence that $(t_i, m_i) \neq (t_j, m_j)$ and $(t_i, c_i) \neq (t_j, c_j)$ for any $i \neq j$.

5.3 Attainable Index Mappings

In the real world $\mathcal{O}$, each tuple $(t_i, m_i, c_i) \in \nu'$ corresponds to an evaluation of F and satisfies

$$p_1(m_i \oplus h_1(t_i)) \oplus p_2(c_i \oplus h_2(t_i)) = h_1(t_i) \oplus h_2(t_i)\,,$$

where we recall that E_{k_2} was replaced with p_2^{-1}. Writing $P_{a_i} := p_1(m_i \oplus h_1(t_i))$ and $P_{b_i} := p_2(c_i \oplus h_2(t_i))$, view ν defines the following q equations:

$$\begin{aligned} P_{a_1} \oplus P_{b_1} &= h_1(t_1) \oplus h_2(t_1)\,,\\ P_{a_2} \oplus P_{b_2} &= h_1(t_2) \oplus h_2(t_2)\,,\\ &\vdots\\ P_{a_q} \oplus P_{b_q} &= h_1(t_q) \oplus h_2(t_q)\,. \end{aligned} \tag{17}$$

Here, some of the unknowns may be equal to each other. We have that $P_{a_i} \neq P_{a_j}$ if and only if $m_i \oplus h_1(t_i) \neq m_j \oplus h_1(t_j)$, and $P_{b_i} \neq P_{b_j}$ if and only if $c_i \oplus h_2(t_i) \neq c_j \oplus h_2(t_j)$. No condition a priori holds for P_{a_i} versus P_{b_j}, as these are defined by independent permutations. We have

$$r = |\{m_i \oplus h_1(t_i) \mid i \in \{1, \dots, q\}\}| + |\{c_i \oplus h_2(t_i) \mid i \in \{1, \dots, q\}\}| \tag{18}$$

unknowns.

5.4 Bad Views

Inspired by the discussion in Section 4, we associate a bipartite graph $G(\nu) = (U, V, E(\nu))$ with the view ν. $U = \{0,1\}^n$ corresponds to the input values to p_1, $V = \{0,1\}^n$ to the output values of p_2^{-1}, and for every $(t_i, m_i, c_i) \in \nu'$, the edge $(m_i \oplus h_1(t_i), c_i \oplus h_2(t_i))$ with label $h_1(t_i) \oplus h_2(t_i)$ from U to V is added to $E(\nu)$. The example graph of Figure 3 still applies, be it with $f_1 = h_1$, $f_2 = h_1 \oplus h_2$, and $f_3 = h_2$.

In Section 4, we already informally discussed what problems could occur in such a graph, i.e., what properties would make the mirror theory inapplicable: it should not contain a path of even length whose labels sum to 0, a circle, or an excessively large tree. The latter is informal, it is often based on a pre-defined threshold on the maximum size of the tree. As our security analysis will cap on $3n/4$-bit security anyway, we can keep it simple, and put as one of the bad events that $G(\nu)$ should not contain a subgraph of ≥ 4 edges. This would imply the non-existence of an excessively large tree, as well as circles and paths of length ≥ 4. We still have to rule out the existence of a path of length 2 whose labels sum to 0 and a circle of length 2.

Formally, we say that a view ν is a *bad view* if its corresponding tree $G(\nu)$ contains

(i) a path of length 2 whose labels sum to 0;
(ii) a circle of length 2;
(iii) a subgraph of ≥ 4 edges.

5.5 Probability of Bad Views (δ)

By Lemma 1, we have to analyze the probability that a view generated in the ideal world is bad, and the analysis will rely on the fact that h_1 and h_2 are 4-wise independent universal hash functions. We have

$$\mathbf{Pr}\left(X_{\widetilde{p}} \in \mathcal{V}_{\mathrm{bad}}\right) \leq \mathbf{Pr}\left(\mathrm{path}\right) + \mathbf{Pr}\left(\mathrm{circle}\right) + \mathbf{Pr}\left(\mathrm{subgraph}\right) , \tag{19}$$

where the sizes of the path, circle, and subgraph, are left implicit.

(i) a path. Consider any two distinct queries $(t_i, m_i, c_i), (t_j, m_j, c_j)$. They yield a 0-label-sum path if either

$$m_i \oplus h_1(t_i) = m_j \oplus h_1(t_j) \text{ and } h_1(t_i) \oplus h_2(t_i) = h_1(t_j) \oplus h_2(t_j) ,$$

or

$$c_i \oplus h_2(t_i) = c_j \oplus h_2(t_j) \text{ and } h_1(t_i) \oplus h_2(t_i) = h_1(t_j) \oplus h_2(t_j) .$$

If $t_i = t_j$, then necessarily $m_i \neq m_j$ and $c_i \neq c_j$ (as the two queries are distinct) and the conditions happen with probability 0. Otherwise, as h_1 and h_2 are ε-AXU$_4$, both conditions happen with probability at most ε^2. Thus,

$$\mathbf{Pr}\left(\mathrm{path}\right) \leq 2\binom{q}{2}\varepsilon^2 . \tag{20}$$

(ii) a circle. Consider any two distinct queries $(t_i, m_i, c_i), (t_j, m_j, c_j)$. They yield a circle if

$$m_i \oplus h_1(t_i) = m_j \oplus h_1(t_j) \text{ and } c_i \oplus h_2(t_i) = c_j \oplus h_2(t_j)\,,$$

which, as before, happens with probability at most ε^2. Thus,

$$\mathbf{Pr}\,(\text{circle}) \leq \binom{q}{2}\varepsilon^2\,. \tag{21}$$

(iii) a subgraph. Consider any four distinct queries $(t_{i_1}, m_{i_1}, c_{i_1}), \ldots, (t_{i_4}, m_{i_4}, c_{i_4})$ to yield a subgraph. We can consider six possible configurations, as described in Figure 5. In these configurations, only collisions are explicitly indicated; two nodes that are different in the configuration may or may not collide. We treat all configurations independently, where we will rely on the fact that h_1 and h_2 are ε-AXU$_4$.

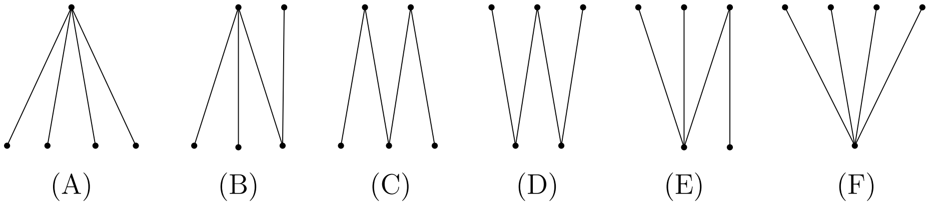

Fig. 5: Possible configurations of subgraphs of 4 edges. Upper shore is U, lower shore is V, and labels are omitted for brevity. Two nodes in the same shore may or may not be equal.

(A) Configuration (A) happens only if

$$m_{i_1} \oplus h_1(t_{i_1}) = m_{i_2} \oplus h_1(t_{i_2}) = m_{i_3} \oplus h_1(t_{i_3}) = m_{i_4} \oplus h_1(t_{i_4})\,.$$

If the tweaks are not all distinct, the condition is satisfied with probability 0. On the other hand, if $t_{i_1}, t_{i_2}, t_{i_3}, t_{i_4}$ are all distinct, the condition is satisfied with probability at most ε^3. There are at most $\binom{q}{4}$ possible choices of queries that satisfy this condition on the tweaks;

(B) Configuration (B) happens only if

$$\begin{aligned} m_{i_1} \oplus h_1(t_{i_1}) = m_{i_2} \oplus h_1(t_{i_2}) = m_{i_3} \oplus h_1(t_{i_3})\,,\\ c_{i_3} \oplus h_2(t_{i_3}) = c_{i_4} \oplus h_2(t_{i_4})\,. \end{aligned}$$

Further analysis depends on the values of the tweaks.

- If $t_{i_1}, t_{i_2}, t_{i_3}, t_{i_4}$ are all distinct, the condition is satisfied with probability at most ε^3. There are at most $\binom{q}{4}$ possible choices of queries that satisfy this condition on the tweaks;

- If $t_{i_1} = t_{i_2}$, $t_{i_1} = t_{i_3}$, $t_{i_2} = t_{i_3}$, or $t_{i_3} = t_{i_4}$, the condition is satisfied with probability 0;
- If $t_{i_1} = t_{i_4}$, but $t_{i_1}, t_{i_2}, t_{i_3}$ are all distinct, the condition is satisfied with probability at most ε^3. There are at most $\binom{q}{3} \cdot (\gamma - 1)$ possible choices of queries that satisfy this condition on the tweaks, noting that every tweak occurs at most γ times;
- If $t_{i_2} = t_{i_4}$, but $t_{i_1}, t_{i_2}, t_{i_3}$ are all distinct, a similar reasoning applies.

Overall, configuration (B) is satisfied with probability at most

$$\max\left\{\binom{q}{4}\varepsilon^3, \binom{q}{3}(\gamma-1)\varepsilon^3\right\} \le \binom{q}{4}\varepsilon^3\,,$$

for $\gamma \le q/4$;

(C) Configuration (C) happens only if

$$\begin{aligned} m_{i_1} \oplus h_1(t_{i_1}) &= m_{i_2} \oplus h_1(t_{i_2})\,, \\ c_{i_2} \oplus h_2(t_{i_2}) &= c_{i_3} \oplus h_2(t_{i_3})\,, \\ m_{i_3} \oplus h_1(t_{i_3}) &= m_{i_4} \oplus h_1(t_{i_4})\,. \end{aligned}$$

Further analysis depends on the values of the tweaks.
- If $t_{i_1}, t_{i_2}, t_{i_3}, t_{i_4}$ are all distinct, the condition is satisfied with probability at most $2^n\varepsilon^4$ (obtained by summing over all possible connections between the first and third equation, and then applying the ε-AXU$_4$ bound). There are at most $\binom{q}{4}$ possible choices of queries that satisfy this condition on the tweaks;
- If $t_{i_1} = t_{i_2}$, $t_{i_2} = t_{i_3}$, or $t_{i_3} = t_{i_4}$, the condition is satisfied with probability 0;
- If $t_{i_1} = t_{i_3}$, but $t_{i_1}, t_{i_2}, t_{i_4}$ are all distinct, the condition is satisfied with probability at most ε^3. There are at most $\binom{q}{3} \cdot (\gamma - 1)$ possible choices of queries that satisfy this condition on the tweaks, noting that every tweak occurs at most γ times;
- If $t_{i_2} = t_{i_4}$, but $t_{i_1}, t_{i_2}, t_{i_3}$ are all distinct, a similar reasoning applies;
- If $t_{i_1} = t_{i_4}$, but $t_{i_1}, t_{i_2}, t_{i_3}$ are all distinct, a similar reasoning applies;
- If $t_{i_1} = t_{i_3}$ and $t_{i_2} = t_{i_4}$ but t_{i_1}, t_{i_2} are distinct, the condition is satisfied with probability at most ε^2. There are at most $\binom{q}{2} \cdot (\gamma - 1)$ possible choices of queries that satisfy this condition on the tweaks, noting that every tweak occurs at most γ times and that there is at most one option for $(t_{i_4}, m_{i_4}, c_{i_4})$ once the other three queries are fixed.

Overall, configuration (C) is satisfied with probability at most

$$\max\left\{\binom{q}{4}2^n\varepsilon^4, \binom{q}{3}(\gamma-1)\varepsilon^3, \binom{q}{2}(\gamma-1)\varepsilon^2\right\} \le \binom{q}{4}2^n\varepsilon^4 + \binom{q}{2}(\gamma-1)\varepsilon^2\,,$$

for $\gamma \le q/4$ and $2^n\varepsilon \ge 1$;

(D) Configuration (D) is symmetrical to configuration (C);

(E) Configuration (E) is symmetrical to configuration (B);

(F) Configuration (F) is symmetrical to configuration (A).

Thus,

$$\begin{aligned}\mathbf{Pr}\left(\text{subgraph}\right) &\leq 4\binom{q}{4}\varepsilon^3 + 2\binom{q}{4}2^n\varepsilon^4 + 2\binom{q}{2}(\gamma-1)\varepsilon^2 \\ &\leq 6\binom{q}{4}2^n\varepsilon^4 + 2\binom{q}{2}(\gamma-1)\varepsilon^2\,. \end{aligned} \tag{22}$$

Conclusion for bad events. From (19) and the individual probabilities of (20), (21), and (22), we obtain

$$\begin{aligned}\mathbf{Pr}\left(X_{\widetilde{p}} \in \mathcal{V}_{\text{bad}}\right) &\leq 3\binom{q}{2}\varepsilon^2 + 6\binom{q}{4}2^n\varepsilon^4 + 2\binom{q}{2}(\gamma-1)\varepsilon^2 \\ &\leq 6\binom{q}{4}2^n\varepsilon^4 + \binom{q}{2}(2\gamma+1)\varepsilon^2\,, \end{aligned}$$

for $\gamma \geq 2$.

5.6 Ratio for Good Views (ϵ)

Consider a given view $\nu = (\nu', h_1, h_2)$ where $\nu = \{(t_1, m_1, c_1), \ldots, (t_q, m_q, c_q)\}$. Define

$$r_1 = |\{m_i \oplus h_1(t_i) \mid i \in \{1, \ldots, q\}\}|\,, \tag{23}$$

$$r_2 = |\{c_i \oplus h_2(t_i) \mid i \in \{1, \ldots, q\}\}|\,. \tag{24}$$

Note that $r_1 + r_2$ is equal to the number of unknowns in the system of equations (see (18)). For any $t \in \{0,1\}^\tau$, we denote $u_t = |\{i \in \{1, \ldots, q\} \mid t_i = t\}|$.

For the ideal world $\widetilde{p}$, we have

$$\begin{aligned}\mathbf{Pr}\left(X_{\widetilde{p}} = \nu\right) &= \mathbf{Pr}\left(\widetilde{p} \xleftarrow{\$} \mathsf{iperm}(\tau, n) \;:\; \widetilde{p} \vdash \nu'\right) \cdot \mathbf{Pr}\left((h_1, h_2) = (h_1', h_2') \xleftarrow{\$} H^2\right) \\ &= \frac{1}{\prod_{t \in \{0,1\}^\tau} (2^n)_{u_t}} \cdot \frac{1}{|H|^2}\,, \end{aligned} \tag{25}$$

where for the first probability we use that $\widetilde{p}$ is a family of permutations and for every $t \in \{0,1\}^\tau$ the view defines u_t values.

For the real world F, recall that it is built from two permutations p_1, p_2^{-1}. We have

$$\begin{aligned}\mathbf{Pr}\left(X_F = \nu\right) &= \mathbf{Pr}\left(p_1, p_2^{-1} \xleftarrow{\$} \mathsf{perm}(n) \;:\; F \vdash \nu' \mid h_1, h_2\right) \cdot \mathbf{Pr}\left((h_1, h_2) = (h_1', h_2') \xleftarrow{\$} H^2\right) \\ &= \mathbf{Pr}\left(p_1, p_2^{-1} \xleftarrow{\$} \mathsf{perm}(n) \;:\; F \vdash \nu' \mid h_1, h_2\right) \cdot \frac{1}{|H|^2}\,. \end{aligned} \tag{26}$$

As has become clear from (17), $\nu = (\nu', h_1, h_2)$ fixes exactly q equations on r_1 unknowns for p_1 and r_2 unknowns for p_2^{-1}, where the inputs to p_1 and p_2^{-1} are fixed. We rely on the following lemma that is based on Patarin's mirror theory.

Lemma 2. *Consider good view $\nu = (\nu', h_1, h_2)$ whose system of q equations (17) has no subgraph of ≥ 4 edges, has no path of length 2 whose labels sum to 0, and no circle of length 2. As long as $5^2 \cdot q \leq 2^n/64$, the number of solutions to the $r_1 + r_2$ unknowns is at least*

$$\frac{(2^n)_{r_1}(2^n-4)_{r_2}}{2^{nq}}\,.$$

The proof of Lemma 2 is omitted: it is very similar to the reasoning on EDM in [20] and follows straightforwardly from Patarin's mirror theory as reviewed in Appendix B. The side condition $5^2 \cdot q \leq 2^n/64$ is slightly different from that in [20], as we have adopted the bound from Nachef, Patarin, and Volte [22].

Every such solution defines r_1 evaluations of p_1, and r_2 evaluations of p_2, and hence the remaining probability in (26) satisfies

$$\mathbf{Pr}\left(p_1, p_2^{-1} \xleftarrow{\$} \mathsf{perm}(n) \ :\ F \vdash \nu' \mid h_1, h_2\right) \geq \frac{(2^n)_{r_1}(2^n-4)_{r_2}}{2^{nq}\cdot(2^n)_{r_1}(2^n)_{r_2}}\,.$$

We obtain for the ratio:

$$\begin{aligned}\frac{\mathbf{Pr}\left(X_F = \nu\right)}{\mathbf{Pr}\left(X_{\widetilde{p}} = \nu\right)} &\geq \frac{\prod_{t\in\{0,1\}^\tau}(2^n)_{u_t}\cdot|H|^2}{1}\cdot\frac{(2^n)_{r_1}(2^n-4)_{r_2}}{2^{nq}\cdot(2^n)_{r_1}(2^n)_{r_2}\cdot|H|^2}\\ &= \frac{\prod_{t\in\{0,1\}^\tau}(2^n)_{u_t}\cdot(2^n-4)_{r_2}}{2^{nq}\cdot(2^n)_{r_2}}\,. \end{aligned} \tag{27}$$

Using that for all t, $u_t \leq \gamma$, and that $\sum_{t\in\{0,1\}^\tau} u_t = q$:

$$\begin{aligned}(27) &\geq \frac{\prod_{t\in\{0,1\}^\tau}(2^n-(\gamma-1))^{u_t}\cdot(2^n-4)_{r_2}}{2^{nq}\cdot(2^n)_{r_2}}\\ &= \left(\frac{2^n-(\gamma-1)}{2^n}\right)^q\cdot\prod_{i=0}^{3}\left(1-\frac{r_2}{2^n-i}\right)\,. \end{aligned} \tag{28}$$

Using that $r_2 \leq q-1$, and by simple algebra for $q \leq 2^n/3$:

$$\begin{aligned}(28) &\geq 1-\left(\frac{(\gamma-1)q}{2^n}+\frac{q-1}{2^n}+\frac{q-1}{2^n-1}+\frac{q-1}{2^n-2}+\frac{q-1}{2^n-3}\right)\\ &\geq 1-\frac{(\gamma+3)q}{2^n}\,.\end{aligned}$$

We have obtained $\epsilon = \frac{(\gamma+3)q}{2^n}$, provided $5^2\cdot q \leq 2^n/64$.

5.7 Releasing Tweak Usage Limitation

The limitation on the tweak usage, namely that the distinguisher can query each tweak at most γ times, is used at two places in the proof.

The first place is the last case of configuration (C) in Section 5.5, namely the case where $t_{i_1} = t_{i_3}$ and $t_{i_2} = t_{i_4}$. For upper bounding the number of choices for the four queries without relying on parameter γ, one may take into account that $m_{i_1} \oplus m_{i_2} = m_{i_3} \oplus m_{i_4}$ is necessarily needed. This value needs to be equal to the random value $h_1(t_{i_1}) \oplus h_2(t_{i_2})$. However, we see no possibility for deriving a formal bound here.

The second place is in the application of the mirror theory in Section 5.6. Our approach to achieve improved $3n/4$-bit security relies on Patarin's mirror theory, which is specifically developed to work well if a scheme is compared with a random function. Obviously, evaluations of CLRW2 under the same tweak will always give distinct responses. In particular, if a distinguisher uses the same tweak for all queries, all responses will be distinct, and the scheme can be distinguished from a random function with probability about $\binom{q}{2}/2^n$. More generally, if every tweak is evaluated at most γ times, the scheme can be distinguished from a random function with probability at most around $\gamma q/2^n$. Resolving the γ limitation here requires improving Patarin's mirror theory or employing a different proof technique.

Acknowledgments. Bart Mennink is supported by a postdoctoral fellowship from the Netherlands Organisation for Scientific Research (NWO) under Veni grant 016.Veni.173.017. The author would like to thank Mridul Nandi and Samuel Neves, and anonymous reviewers for their comments and suggestions.

References

1. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Tischhauser, E., Yasuda, K.: Parallelizable and Authenticated Online Ciphers. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. Lecture Notes in Computer Science, vol. 8269, pp. 424–443. Springer (2013)
2. Bogdanov, A., Knudsen, L.R., Leander, G., Standaert, F., Steinberger, J.P., Tischhauser, E.: Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations - (Extended Abstract). In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. Lecture Notes in Computer Science, vol. 7237, pp. 45–62. Springer (2012)
3. Chakraborty, D., Sarkar, P.: A General Construction of Tweakable Block Ciphers and Different Modes of Operations. In: Lipmaa, H., Yung, M., Lin, D. (eds.) Inscrypt 2006. Lecture Notes in Computer Science, vol. 4318, pp. 88–102. Springer (2006)
4. Chen, S., Lampe, R., Lee, J., Seurin, Y., Steinberger, J.P.: Minimizing the Two-Round Even-Mansour Cipher. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. Lecture Notes in Computer Science, vol. 8616, pp. 39–56. Springer (2014)
5. Chen, S., Steinberger, J.P.: Tight Security Bounds for Key-Alternating Ciphers. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. Lecture Notes in Computer Science, vol. 8441, pp. 327–350. Springer (2014)
6. Cogliati, B., Lampe, R., Seurin, Y.: Tweaking Even-Mansour Ciphers. In: Gennaro, R., Robshaw, M. (eds.) CRYPTO 2015, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 189–208. Springer (2015)

7. Cogliati, B., Seurin, Y.: EWCDM: An Efficient, Beyond-Birthday Secure, Nonce-Misuse Resistant MAC. In: Robshaw and Katz [31], pp. 121–149
8. Dunkelman, O., Keller, N.: A New Criterion for Nonlinearity of Block Ciphers. IEEE Trans. Information Theory 53(11), 3944–3957 (2007)
9. Dworkin, M.: NIST SP 800-38E: Recommendation for Block Cipher Modes of Operation: The XTS-AES Mode for Confidentiality on Storage Devices (2010)
10. Granger, R., Jovanovic, P., Mennink, B., Neves, S.: Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In: Fischlin, M., Coron, J. (eds.) EUROCRYPT 2016, Part I. Lecture Notes in Computer Science, vol. 9665, pp. 263–293. Springer (2016)
11. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust Authenticated-Encryption AEZ and the Problem That It Solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. Lecture Notes in Computer Science, vol. 9056, pp. 15–44. Springer (2015)
12. Iwata, T., Minematsu, K., Peyrin, T., Seurin, Y.: ZMAC: A Fast Tweakable Block Cipher Mode for Highly Secure Message Authentication. In: Katz and Shacham [14], pp. 34–65
13. Jean, J., Nikolić, I., Peyrin, T., Seurin, Y.: Deoxys v1.41 (2016), submission to CAESAR competition
14. Katz, J., Shacham, H. (eds.): CRYPTO 2017, Part III, Lecture Notes in Computer Science, vol. 10403. Springer (2017)
15. Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) FSE 2011. Lecture Notes in Computer Science, vol. 6733, pp. 306–327. Springer (2011)
16. Lampe, R., Seurin, Y.: Tweakable Blockciphers with Asymptotically Optimal Security. In: Moriai, S. (ed.) FSE 2013. Lecture Notes in Computer Science, vol. 8424, pp. 133–151. Springer (2013)
17. Landecker, W., Shrimpton, T., Terashima, R.S.: Tweakable Blockciphers with Beyond Birthday-Bound Security. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. Lecture Notes in Computer Science, vol. 7417, pp. 14–30. Springer (2012)
18. Lee, J., Luykx, A., Mennink, B., Minematsu, K.: Connecting Tweakable and Multi-Key Blockcipher Security. Des. Codes Cryptography 86(3), 623–640 (2018)
19. Liskov, M., Rivest, R.L., Wagner, D.A.: Tweakable Block Ciphers. In: Yung, M. (ed.) CRYPTO 2002. Lecture Notes in Computer Science, vol. 2442, pp. 31–46. Springer (2002)
20. Mennink, B., Neves, S.: Encrypted Davies-Meyer and Its Dual: Towards Optimal Security Using Mirror Theory. In: Katz and Shacham [14], pp. 556–583
21. Minematsu, K.: Improved Security Analysis of XEX and LRW Modes. In: Biham, E., Youssef, A.M. (eds.) SAC 2006. Lecture Notes in Computer Science, vol. 4356, pp. 96–113. Springer (2006)
22. Nachef, V., Patarin, J., Volte, E.: Feistel Ciphers - Security Proofs and Cryptanalysis. Springer (2017)
23. Nyberg, K.: Perfect Nonlinear S-Boxes. In: Davies, D.W. (ed.) EUROCRYPT '91. Lecture Notes in Computer Science, vol. 547, pp. 378–386. Springer (1991)
24. Nyberg, K., Knudsen, L.R.: Provable Security Against Differential Cryptanalysis. In: Brickell, E.F. (ed.) CRYPTO '92. Lecture Notes in Computer Science, vol. 740, pp. 566–574. Springer (1992)
25. Patarin, J.: Étude des Générateurs de Permutations Basés sur le Schéma du D.E.S. Ph.D. thesis, Université Paris 6, Paris, France (Nov 1991)

26. Patarin, J.: On Linear Systems of Equations with Distinct Variables and Small Block Size. In: Won, D., Kim, S. (eds.) ICISC 2005. Lecture Notes in Computer Science, vol. 3935, pp. 299–321. Springer (2005)
27. Patarin, J.: The "Coefficients H" Technique. In: Avanzi, R.M., Keliher, L., Sica, F. (eds.) SAC 2008. Lecture Notes in Computer Science, vol. 5381, pp. 328–345. Springer (2008)
28. Patarin, J.: Introduction to Mirror Theory: Analysis of Systems of Linear Equalities and Linear Non Equalities for Cryptography. Cryptology ePrint Archive, Report 2010/287 (2010)
29. Peyrin, T., Seurin, Y.: Counter-in-Tweak: Authenticated Encryption Modes for Tweakable Block Ciphers. In: Robshaw and Katz [31], pp. 33–63
30. Procter, G.: A Note on the CLRW2 Tweakable Block Cipher Construction. Cryptology ePrint Archive, Report 2014/111 (2014)
31. Robshaw, M., Katz, J. (eds.): CRYPTO 2016, Part I, Lecture Notes in Computer Science, vol. 9814. Springer (2016)
32. Rogaway, P.: Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer (2004)
33. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: Reiter, M.K., Samarati, P. (eds.) ACM CCS 2001. pp. 196–205. ACM (2001)
34. Wegman, M.N., Carter, L.: New Hash Functions and Their Use in Authentication and Set Equality. J. Comput. Syst. Sci. 22(3), 265–279 (1981)

A Proof of Theorem 1

Consider the distinguisher of Section 3.2 for any $\epsilon \geq 0$. Its success advantage satisfies

$$\begin{aligned}\mathbf{Adv}^{\mathrm{stprp}}_{\mathrm{GCL}^{f_1,f_2,f_3}}(\mathcal{D}) &= \mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 1\right) - \mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}} = 1\right) \\ &= 1 - \mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 0\right) - \mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}} = 1\right) . \qquad (29)\end{aligned}$$

The derivation relies on the following two lemmas, the proofs of which are in Sections A.1 and A.2.

Lemma 3. *Provided* $n \geq 6$, $\mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 0\right) \leq \frac{32}{2^{4\epsilon}} + \frac{80}{2^{n/2+2\epsilon}}$.

Lemma 4. *For any integral* $1 \leq \alpha \leq \sqrt{\beta} - 1$, *provided* $n \geq 16$, $\mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}} = 1\right) \leq \alpha 2^n \left(\frac{2\alpha}{2^{2\epsilon}}\right)^{3/(4\alpha^2)\cdot 2^{4\epsilon}} + \frac{2^{(\alpha+2)2\epsilon}}{2^{(\alpha-2)n/2}}$.

Putting $\epsilon = \log_2(n)/2$, we derive from (29) and Lemmas 3 and 4 that

$$\mathbf{Adv}^{\mathrm{stprp}}_{\mathrm{GCL}^{f_1,f_2,f_3}}(\mathcal{D}) \geq 1 - \frac{32}{n^2} - \frac{80}{n2^{n/2}} - \alpha 2^n \left(\frac{2\alpha}{n}\right)^{3/(4\alpha^2)\cdot n^2} - \frac{n^{(\alpha+2)}}{2^{(\alpha-2)n/2}} ,$$

provided $n \geq 16$, and for any integral $1 \leq \alpha \leq \sqrt{3/8 n} - 1$. Clearly, the bound is meaningless for $\alpha = 1, 2$. Computer verification yields optimal choice $\alpha = 5$.

A.1 Proof of Lemma 3

Putting $d^* = f_1(t) \oplus f_1(t')$, we have

$$\mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 0\right) = \mathbf{Pr}\left(\forall_{d\in\{0,1\}^n} N_d < \beta\right) \leq \mathbf{Pr}\left(N_{d^*} < \beta\right) . \tag{30}$$

Clearly, if $f_2(t) \oplus f_2(t') = 0$, then $c_i \oplus c'_j = f_3(t) \oplus f_3(t')$ for all $(i,j) \in I_{d^*}$ and thus $N_{d^*} = \binom{q'}{2} > \beta$, implying $\mathbf{Pr}\left(N_{d^*} < \beta\right) = 0$. Henceforth, assume that $d^{**} := f_2(t) \oplus f_2(t') \neq 0$.

By Chebychev's inequality:

$$\begin{aligned}\mathbf{Pr}\left(N_{d^*} < \beta\right) &= \mathbf{Pr}\left(N_{d^*} - \mathbf{Ex}\left(N_{d^*}\right) < \beta - \mathbf{Ex}\left(N_{d^*}\right)\right) \\ &\leq \mathbf{Pr}\left(\left|N_{d^*} - \mathbf{Ex}\left(N_{d^*}\right)\right| \geq \mathbf{Ex}\left(N_{d^*}\right) - \beta\right) \\ &\leq \frac{\mathbf{Var}\left(N_{d^*}\right)}{\left(\mathbf{Ex}\left(N_{d^*}\right) - \beta\right)^2} \\ &= \frac{\mathbf{Ex}\left(\left(N_{d^*}\right)^2\right) - \left(\mathbf{Ex}\left(N_{d^*}\right)\right)^2}{\left(\mathbf{Ex}\left(N_{d^*}\right) - \beta\right)^2} .\end{aligned} \tag{31}$$

For distinct $(i,j),(k,l) \in I_{d^*}$, define

$$N_{d^*}^{(i,j),(k,l)} = \begin{cases} 1, \text{ if } c_i \oplus c'_j = c_k \oplus c'_l\,, \\ 0, \text{ otherwise}\,, \end{cases} \tag{32}$$

such that

$$N_{d^*} = \sum_{\substack{(i,j),(k,l)\in I_{d^*} \\ (i,j)\neq(k,l)}} N_{d^*}^{(i,j),(k,l)} . \tag{33}$$

We have

$$\mathbf{Ex}\left(N_{d^*}\right) = \sum_{\substack{(i,j),(k,l)\in I_{d^*} \\ (i,j)\neq(k,l)}} \mathbf{Pr}\left(c_i \oplus c'_j = c_k \oplus c'_l\right) , \tag{34}$$

and

$$\begin{aligned}\mathbf{Ex}\left(\left(N_{d^*}\right)^2\right) &= \mathbf{Ex}\left(\sum_{\substack{(i,j),(k,l)\in I_{d^*} \\ (i,j)\neq(k,l)}} \sum_{\substack{(i',j'),(k',l')\in I_{d^*} \\ (i',j')\neq(k',l')}} N_{d^*}^{(i,j),(k,l)} N_{d^*}^{(i',j'),(k',l')}\right) \\ &= \sum_{\substack{(i,j),(k,l)\in I_{d^*} \\ (i,j)\neq(k,l)}} \sum_{\substack{(i',j'),(k',l')\in I_{d^*} \\ (i',j')\neq(k',l')}} \mathbf{Pr}\left(c_i \oplus c'_j = c_k \oplus c'_l\ ,\ c_{i'} \oplus c'_{j'} = c_{k'} \oplus c'_{l'}\right) .\end{aligned} \tag{35}$$

Above summation consists of $\binom{q'}{2}^2$ terms of independent probabilities, but their values differ depending on overlaps in the two sets $\{(i,j),(k,l)\}$, $\{(i',j'),(k',l')\}$. For any *distinct* $(i_1,j_1),(i_2,j_2),(i_3,j_3),(i_4,j_4) \in I_{d^*}$, define

$$\begin{aligned}
\mathbf{P}_2 &:= \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}\right) \,, \\
\mathbf{P}_3 &:= \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3}\right) \,, \\
\mathbf{P}_4 &:= \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \,,\ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4}\right) \,.
\end{aligned}$$

We can observe that the sum in (35) consists of exactly $\binom{q'}{2}$ terms satisfying $\left|\{(i,j),(k,l)\} \cup \{(i',j'),(k',l')\}\right| = 2$, in which case the corresponding probability is of the form $\mathbf{P}_2$, exactly $\binom{q'}{2}2\binom{q'-2}{1}$ terms satisfying $\left|\{(i,j),(k,l)\} \cup \{(i',j'),(k',l')\}\right| = 3$, in which case the corresponding probability is of the form $\mathbf{P}_3$, and exactly $\binom{q'}{2}\binom{q'-2}{2}$ terms satisfying $\left|\{(i,j),(k,l)\} \cup \{(i',j'),(k',l')\}\right| = 4$, in which case the corresponding probability is of the form $\mathbf{P}_4$. We obtain (using independence of the probabilities)

$$\mathbf{Ex}\left(\left(N_{d^*}\right)^2\right) = \binom{q'}{2} \cdot \mathbf{P}_2 + \binom{q'}{2}2\binom{q'-2}{1} \cdot \mathbf{P}_3 + \binom{q'}{2}\binom{q'-2}{2} \cdot \mathbf{P}_4 \,.$$

We likewise have $\mathbf{Ex}\left(N_{d^*}\right) = \binom{q'}{2} \cdot \mathbf{P}_2$, and using that $\beta = \frac{3}{2}\binom{q'}{2}/2^n$, we obtain for (30-31):

$$\begin{aligned}
\mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 0\right) &\leq \frac{\binom{q'}{2} \cdot \mathbf{P}_2 + \binom{q'}{2}2\binom{q'-2}{1} \cdot \mathbf{P}_3 + \binom{q'}{2}\binom{q'-2}{2} \cdot \mathbf{P}_4 - \left(\binom{q'}{2} \cdot \mathbf{P}_2\right)^2}{(\binom{q'}{2} \cdot \mathbf{P}_2 - \frac{3}{2}\binom{q'}{2}/2^n)^2} \\
&= \frac{\mathbf{P}_2 + 2\binom{q'-2}{1} \cdot \mathbf{P}_3 + \binom{q'-2}{2} \cdot \mathbf{P}_4 - \binom{q'}{2} \cdot \mathbf{P}_2^2}{\binom{q'}{2}(\mathbf{P}_2 - \frac{3}{2}/2^n)^2} \,. \qquad (36)
\end{aligned}$$

We can derive the following bounds on $\mathbf{P}_2, \mathbf{P}_3, \mathbf{P}_4$.

Claim. Provided $n \geq 6$, $\mathbf{P}_2 \geq 2/2^n$, $\mathbf{P}_3 \leq 5/2^{2n}$, and $\mathbf{P}_4 \leq \frac{4}{(2^n-6)(2^n-7)}$.

Proof (proof of claim). Before bounding the probabilities separately, note that in general for any distinct $(i,j),(k,l) \in I_{d^*}$, we have $i \neq k$ and $j \neq l$. Write

$$\begin{aligned}
x_{i_1} &= p_1(m_{i_1} \oplus f_1(t)) = p_1(m'_{j_1} \oplus f_1(t')) \,, \\
x_{i_2} &= p_1(m_{i_2} \oplus f_1(t)) = p_1(m'_{j_2} \oplus f_1(t')) \,, \\
x_{i_3} &= p_1(m_{i_3} \oplus f_1(t)) = p_1(m'_{j_3} \oplus f_1(t')) \,, \\
x_{i_4} &= p_1(m_{i_4} \oplus f_1(t)) = p_1(m'_{j_4} \oplus f_1(t')) \,,
\end{aligned}$$

where we recall that $d^* = f_1(t) \oplus f_1(t') = m_{i_1} \oplus m'_{j_1} = \cdots = m_{i_4} \oplus m'_{j_4}$. Above values $x_{i_1}, x_{i_2}, x_{i_3}, x_{i_4}$ are pairwise distinct as $m_{i_1}, m_{i_2}, m_{i_3}, m_{i_4}$ are pairwise distinct as i_1, i_2, i_3, i_4 are. Furthermore, write

$$\begin{aligned}
y_{i_1} &= p_2^{-1}(c_{i_1} \oplus f_3(t)) = x_{i_1} \oplus f_2(t)\,, \\
y'_{j_1} &= p_2^{-1}(c'_{j_1} \oplus f_3(t')) = x_{i_1} \oplus f_2(t')\,, \\
y_{i_2} &= p_2^{-1}(c_{i_2} \oplus f_3(t)) = x_{i_2} \oplus f_2(t)\,, \\
y'_{j_2} &= p_2^{-1}(c'_{j_2} \oplus f_3(t')) = x_{i_2} \oplus f_2(t')\,, \\
y_{i_3} &= p_2^{-1}(c_{i_3} \oplus f_3(t)) = x_{i_3} \oplus f_2(t)\,, \\
y'_{j_3} &= p_2^{-1}(c'_{j_3} \oplus f_3(t')) = x_{i_3} \oplus f_2(t')\,, \\
y_{i_4} &= p_2^{-1}(c_{i_4} \oplus f_3(t)) = x_{i_4} \oplus f_2(t)\,, \\
y'_{j_4} &= p_2^{-1}(c'_{j_4} \oplus f_3(t')) = x_{i_4} \oplus f_2(t')\,.
\end{aligned}$$

Recall that $d^{**} := f_2(t) \oplus f_2(t') \neq 0$.

We start with bounding $\mathbf{P}_2$:

$$\begin{aligned}
\mathbf{P}_2 = \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}\right) \\
&= \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \mid x_{i_1} \oplus x_{i_2} = d^{**}\right) \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**}\right) \\
&\quad + \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \mid x_{i_1} \oplus x_{i_2} \neq d^{**}\right) \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} \neq d^{**}\right)\,.
\end{aligned}$$

Given that $x_{i_1} \neq x_{i_2}$, we have

$$\mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**}\right) = \frac{1}{2^n - 1}\,.$$

Conditioned on $x_{i_1} \oplus x_{i_2} = d^{**}$, we have $y_{i_1} = y'_{j_2}$ and $y'_{j_1} = y_{i_2}$, and $c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}$ holds with probability 1. Conditioned on $x_{i_1} \oplus x_{i_2} \neq d^{**}$ and using that $d^{**} \neq 0$, the values $y_{i_1}, y'_{j_1}, y_{i_2}, y'_{j_2}$ are pairwise distinct and

$$\mathbf{Pr}\left(p_2(y_{i_1}) \oplus p_2(y'_{j_1}) = p_2(y_{i_2}) \oplus p_2(y'_{j_2}) \mid x_{i_1} \oplus x_{i_2} \neq d^{**}\right) \leq \frac{1}{2^n - 3}\,.$$

We therefore obtain

$$\mathbf{P}_2 = \frac{1}{2^n - 1} + \frac{1}{2^n - 3}\left(1 - \frac{1}{2^n - 1}\right) = \frac{2 \cdot 2^n - 5}{(2^n - 1)(2^n - 3)} \geq \frac{2}{2^n}\,.$$

We next bound $\mathbf{P}_3$:

$$\begin{aligned}
\mathbf{P}_3 = \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3}\right) \\
&= \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_1} \oplus x_{i_2} = d^{**}\right) \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**}\right) \\
&\quad + \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_1} \oplus x_{i_3} = d^{**}\right) \mathbf{Pr}\left(x_{i_1} \oplus x_{i_3} = d^{**}\right) \\
&\quad + \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_2} \oplus x_{i_3} = d^{**}\right) \mathbf{Pr}\left(x_{i_2} \oplus x_{i_3} = d^{**}\right) \\
&\quad + \mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_1} \oplus x_{i_2}, x_{i_1} \oplus x_{i_3}, x_{i_2} \oplus x_{i_3} \neq d^{**}\right) \\
&\qquad\qquad\qquad\qquad \cdot \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2}, x_{i_1} \oplus x_{i_3}, x_{i_2} \oplus x_{i_3} \neq d^{**}\right)\,,
\end{aligned}$$

using that no two or more of the events "$x_{i_1} \oplus x_{i_2} = d^{**}$," "$x_{i_1} \oplus x_{i_3} = d^{**}$," and "$x_{i_2} \oplus x_{i_3} = d^{**}$" can hold simultaneously. Starting with the first line, as before we have

$$\mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**}\right) = \frac{1}{2^n - 1} .$$

Conditioned on $x_{i_1} \oplus x_{i_2} = d^{**}$, we have $y_{i_1} = y'_{j_2}$ and $y'_{j_1} = y_{i_2}$, and $c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}$ holds with probability 1. On the other hand, $x_{i_1} \oplus x_{i_3} \neq d^{**}$, and thus, the values $y_{i_1}, y'_{j_1}, y_{i_3}, y'_{j_3}$ are pairwise distinct and

$$\mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_1} \oplus x_{i_2} = d^{**}\right) \leq \frac{1}{2^n - 3}$$

(we now need to consider an upper bound, as the probability may be 0 if the targeted value is already sampled).

The second and third line go identically. For the fourth line, conditioned on the fact that $x_{i_1} \oplus x_{i_2}, x_{i_1} \oplus x_{i_3}, x_{i_2} \oplus x_{i_3} \neq d^{**}$ and using that $d^{**} \neq 0$, the values $y_{i_1}, y'_{j_1}, y_{i_2}, y'_{j_2}, y_{i_3}, y'_{j_3}$ are pairwise distinct and

$$\mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} = c_{i_3} \oplus c'_{j_3} \mid x_{i_1} \oplus x_{i_2}, x_{i_1} \oplus x_{i_3}, x_{i_2} \oplus x_{i_3} \neq d^{**}\right) \leq \frac{1}{(2^n - 4)(2^n - 5)} .$$

We therefore obtain

$$\mathbf{P}_3 \leq \frac{3}{(2^n - 1)(2^n - 3)} + \frac{1}{(2^n - 4)(2^n - 5)} \leq \frac{4}{(2^n - 4)(2^n - 5)} \leq \frac{5}{2^{2n}} ,$$

provided $2^n \geq 45$.

We finally bound $\mathbf{P}_4$:

$$\begin{aligned}
\mathbf{P}_4 = \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \ , \ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4}\right) \\
= \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \ , \ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}\right) \\
&\qquad\qquad\qquad\qquad \cdot \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}\right) \\
+ \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \ , \ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \\
&\qquad\qquad\qquad\qquad \cdot \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \\
+ \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \ , \ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2} \neq d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}\right) \\
&\qquad\qquad\qquad\qquad \cdot \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} \neq d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}\right) \\
+ \mathbf{Pr}&\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2} \ , \ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2} \neq d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \\
&\qquad\qquad\qquad\qquad \cdot \mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} \neq d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) ,
\end{aligned}$$

For the first line, the event $x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}$ holds with probability $1/(2^n - 2)(2^n - 3)$, and conditioned on $x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} = d^{**}$, the equations $c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}$ and $c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4}$ hold with probability 1 (see the analysis of $\mathbf{P}_2$). The second and third line go as in the analysis of $\mathbf{P}_3$, giving

$$\mathbf{Pr}\left(x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \leq \frac{1}{2^n - 1} ,$$

and

$$\mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}\ ,\ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2} = d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \leq \frac{1}{2^n - 3}\,.$$

For the fourth line, conditioned on the fact that $x_{i_1} \oplus x_{i_2} \neq d^{**} \wedge x_{i_3} \oplus x_{i_4} \neq d^{**}$ and using that $d^{**} \neq 0$, the values $y_{i_1}, y'_{j_1}, y_{i_2}, y'_{j_2}$ are pairwise distinct and so are $y_{i_3}, y'_{j_3}, y_{i_4}, y'_{j_4}$, and in addition, $y_{i_1}, y_{i_2}, y_{i_3}, y_{i_4}$ are pairwise distinct and $y'_{j_1}, y'_{j_2}, y'_{j_3}, y'_{j_4}$ are. We obtain

$$\mathbf{Pr}\left(c_{i_1} \oplus c'_{j_1} = c_{i_2} \oplus c'_{j_2}\ ,\ c_{i_3} \oplus c'_{j_3} = c_{i_4} \oplus c'_{j_4} \mid x_{i_1} \oplus x_{i_2}, x_{i_3} \oplus x_{i_4} \neq d^{**}\right) \leq \frac{1}{(2^n - 6)(2^n - 7)}\,.$$

We therefore obtain

$$\mathbf{P}_4 \leq \frac{1}{(2^n - 2)(2^n - 3)} + \frac{2}{(2^n - 1)(2^n - 3)} + \frac{1}{(2^n - 6)(2^n - 7)} \leq \frac{4}{(2^n - 6)(2^n - 7)}\,.$$

□

To suit further analysis of (36), we claim that the $\mathbf{P}_4$-term cancels out to the $\mathbf{P}_2^2$-term.

Claim. Provided $6q' \leq 2^n$, $\binom{q'-2}{2} \cdot \mathbf{P}_4 \leq \binom{q'}{2} \cdot \mathbf{P}_2^2$.

Proof (proof of claim). By above claim, $\mathbf{P}_4 \leq \frac{4}{(2^n-6)(2^n-7)}$ and $\mathbf{P}_2 \geq 2/2^n$, and it remains to prove that

$$\frac{(q'-2)(q'-3)}{(2^n-6)(2^n-7)} \leq \frac{q'(q'-1)}{2^{2n}}\,.$$

This in turn follows from the fact that

$$\frac{q'-3}{2^n-7} \leq \frac{q'-2}{2^n-6} \leq \frac{q'-1}{2^n}\,,$$

as $6q' \leq 2^n$. □

From (36) and the bounds of above two claims, we directly obtain

$$\begin{aligned}
\mathbf{Pr}\left(\mathcal{D}^{\mathrm{GCL}^{f_1,f_2,f_3}} = 0\right) &\overset{a}{\leq} \frac{\mathbf{P}_2 + 2\binom{q'-2}{1} \cdot \mathbf{P}_3}{\binom{q'}{2}(\mathbf{P}_2 - \frac{3}{2}/2^n)^2} \\
&\overset{b}{\leq} \frac{2/2^n + 2\binom{q'-2}{1} \cdot 5/2^{2n}}{\binom{q'}{2}(2/2^n - \frac{3}{2}/2^n)^2} \\
&= \frac{8 \cdot 2^n + 40(q'-2)}{\binom{q'}{2}} \\
&\overset{c}{\leq} \frac{32}{2^{4\epsilon}} + \frac{80}{2^{n/2+2\epsilon}}\,,
\end{aligned}$$

where $\overset{a}{\leq}$ holds due to the second claim, $\overset{b}{\leq}$ holds as $\mathbf{P}_2 \geq 2/2^n$ and $\mathbf{P}_3 \leq 5/2^{2n}$ (note that a lower bound on $\mathbf{P}_2$ suffices for both the numerator and denominator as $A/(A-C) \leq B/(B-C)$ for $A \geq B > C > 0$), and $\overset{c}{\leq}$ holds as $\binom{q'}{2} \geq (q')^2/4$ and $q' = 2^{n/2+2\epsilon}$.

A.2 Proof of Lemma 4

For any $d \in \{0,1\}^n$, recall that N_d counts the number of collisions $c_i \oplus c'_j = c_k \oplus c'_l$ for distinct $(i,j),(k,l)$. There could be multi-collisions; for $\lambda \geq 2$ we say that $(i_1,j_1),\dots,(i_\lambda,j_\lambda) \in I_d$ form a λ-collision if $c_{i_1} \oplus c'_{j_1} = \cdots = c_{i_5} \oplus c'_{j_5}$. Denote by N_d^λ the number of λ-collisions that are *not part of* a $(\lambda+1)$-collision. Denote by $N_d^{\geq\lambda}$ the number of λ-collisions (that may be part of a $(\lambda+1)$-collision).

Fix any $1 \leq \alpha \leq \sqrt{\beta} - 1$. By basic probability theory,[1]

$$\begin{aligned}
\mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}} = 1\right) &\leq \sum_{d\in\{0,1\}^n} \mathbf{Pr}\left(N_d \geq \beta\right) \\
&\leq \sum_{d\in\{0,1\}^n} \mathbf{Pr}\left(N_d \geq \beta \mid N_d^{\geq\alpha+2} = 0\right) + \mathbf{Pr}\left(N_d^{\geq\alpha+2} \geq 1\right) \\
&\leq \sum_{d\in\{0,1\}^n} \mathbf{Pr}\left(N_d \geq \beta \mid N_d^{\geq\alpha+2} = 0\right) + \binom{q'}{\alpha+2}\frac{1}{(2^n)_{\alpha+1}}\,. \qquad (37)
\end{aligned}$$

Conditioned on the fact that there is no $(\alpha+2)$-collision, by the pigeonhole principle, $N_d \geq \beta$ only if the number of collisions arising from either 2-collisions, 3-collisions, $\dots$, or $(\alpha+1)$-collisions is at least β/α. Clearly, a 2-collision contributes 1 to N_d, a 3-collision contributes 3 to N_d, and generally, an i-collision contributes $\binom{i}{2}$ to N_d. Therefore, denoting $\mathbf{Pr}^\star(X) = \mathbf{Pr}\left(X \;\middle|\; N_d^{\geq\alpha+2} = 0\right)$ for brevity,

$$\begin{aligned}
\mathbf{Pr}^\star(N_d \geq \beta) &\leq \sum_{i=2}^{\alpha+1} \mathbf{Pr}^\star\left(N_d^i \geq \beta/\alpha\right) \\
&\leq \sum_{i=2}^{\alpha+1} \binom{q'}{i\cdot\beta/(\alpha\binom{i}{2})}\frac{1}{(2^n)_{(i-1)\cdot\beta/(\alpha\binom{i}{2})}}\,. \qquad (38)
\end{aligned}$$

As $\alpha \leq \sqrt{\beta} - 1$, we particularly have $(i-1)\cdot\beta/(\alpha\binom{i}{2}) \geq 2$ for all i, and we obtain

$$\begin{aligned}
\binom{q'}{i\cdot\beta/(\alpha\binom{i}{2})}\frac{1}{(2^n)_{(i-1)\cdot\beta/(\alpha\binom{i}{2})}} &\overset{a}{\leq} \frac{(q')_{(i-1)\cdot\beta/(\alpha\binom{i}{2})}\cdot(q'-2)^{\beta/(\alpha\binom{i}{2})}}{(2^n)_{(i-1)\cdot\beta/(\alpha\binom{i}{2})}}\cdot\left(\frac{e}{i\cdot\beta/(\alpha\binom{i}{2})}\right)^{i\cdot\beta/(\alpha\binom{i}{2})} \\
&\overset{b}{\leq} \left(\left(\frac{e\alpha(i-1)}{2}\right)^i\cdot\frac{(q')^{i-1}(q'-2)}{2^{(i-1)n}}\cdot\frac{1}{\beta^i}\right)^{\beta/(\alpha\binom{i}{2})} \\
&\overset{c}{\leq} \left(\left(\frac{2e\alpha}{3}\right)^i\cdot\frac{(i-1)^i}{2^{4\epsilon}(2^{n/2+2\epsilon}-1)^{i-2}}\right)^{(3\cdot 2^{4\epsilon})/(8\alpha\binom{i}{2})} \\
&\overset{d}{\leq} \left(\frac{2\alpha}{2^{2\epsilon}}\right)^{3/(4\alpha(i-1))\cdot 2^{4\epsilon}}\,,
\end{aligned}$$

[1] Note that a plain Markov bound or Chebychev's inequality do not help, as we have to sum over all possible $d \in \{0,1\}^n$.

where $\overset{a}{\leq}$ holds as $\binom{A}{B} \leq (A)_B \cdot (e/B)^B$ by Stirling's approximation, $\overset{b}{\leq}$ holds as $(A)_m/(B)_m \leq (A/B)^m$ if $A \leq B$, $\overset{c}{\leq}$ uses $\beta = \frac{3}{2}\binom{q'}{2}/2^n$, $q'(q'-2) \leq (q'-1)^2$, and $q' = 2^{n/2+2\epsilon}$, and, finally $\overset{d}{\leq}$ holds as $(i-1)^i \leq (2^{n/2-1})^{i-2}$ is satisfied for all i, provided that $n \geq 16$.

We obtain for (38):

$$\mathbf{Pr}^\star(N_d \geq \beta) \leq \sum_{i=2}^{\alpha+1} \left(\frac{2\alpha}{2^{2\epsilon}}\right)^{3/(4\alpha(i-1))\cdot 2^{4\epsilon}} \leq \alpha \left(\frac{2\alpha}{2^{2\epsilon}}\right)^{3/(4\alpha^2)\cdot 2^{4\epsilon}} ,$$

and for (37):

$$\mathbf{Pr}\left(\mathcal{D}^{\widetilde{p}} = 1\right) \leq \alpha 2^n \left(\frac{2\alpha}{2^{2\epsilon}}\right)^{3/(4\alpha^2)\cdot 2^{4\epsilon}} + \frac{2^{(\alpha+2)2\epsilon}}{2^{(\alpha-2)n/2}} ,$$

again using that $(A)_m/(B)_m \leq (A/B)^m$ if $A \leq B$. This bound holds for all $1 \leq \alpha \leq \sqrt{\beta} - 1$.

B Mirror Theory

We will follow the description of Patarin's mirror theory [22,26,28] by Mennink and Neves [20]. We will restrict ourselves to the simplified setting where the equations are of the form $P_a \oplus P_b = \lambda$, where the P_a's and P_b's come from independent permutations, and we will use the theory for $3n/4$-bit security at most.

B.1 System of Equations

Consider a system of $q \geq 1$ equations

$$\mathcal{E} = \{P_{\varphi(a_1)} \oplus P_{\varphi(b_1)} = \lambda_1, \cdots, P_{\varphi(a_q)} \oplus P_{\varphi(b_q)} = \lambda_q\} \tag{39}$$

over $r \geq 1$ unknowns $\mathcal{P} = \{P_1, \ldots, P_r\}$, where φ is some surjective index mapping

$$\varphi : \{a_1, b_1, \ldots, a_q, b_q\} \to \{1, \ldots, r\} .$$

In our work we consider the case that the P_a's and P_b's come from independent permutations, hence $\varphi(a_i) \neq \varphi(b_j)$ for any i, j. We write $\mathcal{I}_1 = \{\varphi(a_i) \mid i \in \{1, \ldots, q\}\}$ and $\mathcal{I}_2 = \{\varphi(b_i) \mid i \in \{1, \ldots, q\}\}$, such that $\{1, \ldots, r\} = \mathcal{I}_1 \cup \mathcal{I}_2$ is a partition. For a subset $I \subseteq \{1, \ldots, q\}$, define the multiset $\mathcal{M}_I$ as

$$\mathcal{M}_I = \bigcup_{i \in I} \{\varphi(a_i), \varphi(b_i)\} .$$

We give three definitions with respect to the system of equations $\mathcal{E}$.

Definition 1 (circle-freeness). *For any $I \subseteq \{1,\dots,q\}$, $\mathcal{M}_I$ has an element of odd multiplicity.*

Definition 2 (ξ-block-maximality). *There is a partition $\{1,\dots,r\} = \mathcal{R}_1 \cup \cdots \cup \mathcal{R}_s$ of the r indices, all of size at most ξ, such that for any $i \in \{1,\dots,q\}$ there is an $\ell \in \{1,\dots,s\}$ such that $\{\varphi(a_i),\varphi(b_i)\} \subseteq \mathcal{R}_\ell$.*

Definition 3 (non-degeneracy). *For any $I \subseteq \{1,\dots,q\}$ such that $\mathcal{M}_I$ has exactly two odd multiplicity element from either $\mathcal{I}_1$ or $\mathcal{I}_2$, it satisfies $\bigoplus_{i\in I} \lambda_i \neq 0$.*

Circle-freeness implies that there is no linear combination of the equations $\mathcal{E}$ that is independent of the unknowns, ξ-block maximality implies that there are not too many unknowns that are jointly related, and non-degeneracy implies that there is no linear combination of the equations $\mathcal{E}$ that implies equality of two distinct unknowns from either $\mathcal{I}_1$ or $\mathcal{I}_2$.

B.2 Main Result

The main theorem of Patarin's mirror theory, tailored to the case where we have a partition of the unknowns into two disjoint sets, is given below. We follow [20], with the side condition on $2^n/64$ from [22].

Theorem 3 (mirror theorem). *Let $\{1,\dots,r\} = \mathcal{I}_1 \cup \mathcal{I}_2$ be a partition of the indices. Let $\mathcal{E}$ be a system of equations over the unknowns $\mathcal{P}$ that is (i) circle-free, (ii) ξ-block-maximal, and (iii) non-degenerate. Then, as long as $\xi^2 \cdot \max\{|\mathcal{I}_1|,|\mathcal{I}_2|\} \leq 2^n/64$, the number of solutions for $\mathcal{P}$ such that $P_i \neq P_j$ for all $i,j \in \mathcal{I}_\ell$ ($\ell = 1,2$) is at least*

$$\frac{\mathrm{NonEq}(\mathcal{I}_1,\mathcal{I}_2;\mathcal{E})}{2^{nq}},$$

where $\mathrm{NonEq}(\mathcal{I}_1,\mathcal{I}_2;\mathcal{E})$ denotes the number of solutions to $\mathcal{P}$ that satisfy $P_i \neq P_j$ for all $i,j \in \mathcal{I}_\ell$ ($\ell = 1,2$) as well as the inequalities imposed by $\mathcal{E}$ (but the equalities themselves released).

A lower bound on the technical quantity $\mathrm{NonEq}(\mathcal{I}_1,\mathcal{I}_2;\mathcal{E})$ can be derived as follows. Every equation $P_{\varphi(a)} \oplus P_{\varphi(b)} = \lambda \neq 0$ in $\mathcal{E}$ imposes $P_{\varphi(a)} \neq P_{\varphi(b)}$. As $\varphi(a) \in \mathcal{I}_1$ and $\varphi(b) \in \mathcal{I}_2$ are in distinct index sets, this inequality $P_{\varphi(a)} \neq P_{\varphi(b)}$ imposes an *extra* inequality over the ones suggested by $\mathcal{I}_1, \mathcal{I}_2$. An obvious lower bound thus is

$$\mathrm{NonEq}(\mathcal{I}_1,\mathcal{I}_2;\mathcal{E}) \geq (2^n)_{|\mathcal{I}_1|}(2^n - (\xi-1))_{|\mathcal{I}_2|},$$

because every unknown of $\mathcal{I}_2$ is in exactly one block, and connects with at most $\xi - 1$ unknowns of $\mathcal{I}_1$ (as the system is ξ-block-maximal).