

Relaxing Full-Codebook Security: A Refined Analysis of Key-Length Extension Schemes

Peter Gaži*, Jooyoung Lee**, Yannick Seurin***, John Steinberger†, and
Stefano Tessaro‡

Abstract. We revisit the security (as a pseudorandom permutation) of cascading-based constructions for block-cipher key-length extension. Previous works typically considered the extreme case where the adversary is given the *entire codebook* of the construction, the only complexity measure being the number q_e of queries to the underlying ideal block cipher, representing adversary’s secret-key-independent computation. Here, we initiate a systematic study of the more natural case of an adversary restricted to adaptively learning a number q_c of plaintext/ciphertext pairs that is *less than the entire codebook*. For any such q_c , we aim to determine the highest number of block-cipher queries q_e the adversary can issue without being able to successfully distinguish the construction (under a secret key) from a random permutation.

More concretely, we show the following results for key-length extension schemes using a block cipher with n -bit blocks and κ -bit keys:

- Plain cascades of length $\ell = 2r + 1$ are secure whenever $q_c q_e^r \ll 2^{r(\kappa+n)}$, $q_c \ll 2^\kappa$ and $q_e \ll 2^{2\kappa}$. The bound for $r = 1$ also applies to two-key triple encryption (as used within Triple DES).
- The r -round XOR-cascade is secure as long as $q_c q_e^r \ll 2^{r(\kappa+n)}$, matching an attack by Gaži (CRYPTO 2013).
- We fully characterize the security of Gaži and Tessaro’s two-call 2XOR construction (EUROCRYPT 2012) for all values of q_c , and note that the addition of a third whitening step strictly increases security for $2^{n/4} \leq q_c \leq 2^{3/4n}$. We also propose a variant of this construction *without* re-keying and achieving comparable security levels.

Keywords: block ciphers, key-length extension, provable security, ideal-cipher model

* IST Austria, E-mail: peter.gazi@ist.ac.at

** Sejong University, Seoul, Korea. E-mail: jlee05@sejong.ac.kr

*** ANSSI, Paris, France. E-mail: yannick.seurin@m4x.org

† Tsinghua University, P.R. China. E-mail: jpsteinb@gmail.com

‡ UC Santa Barbara. E-mail: tessaro@cs.ucsb.edu

1 Introduction

1.1 Block Ciphers and Key-Length Extension

Block ciphers (like DES [10] and AES [2]) are the workhorses of cryptography. Most importantly, they constitute the basic building block within several modes of operation for secret-key message encryption and authentication.

Formally, a block cipher with *key length* κ and *block length* n (often referred to as a (κ, n) -*block cipher*) is a family of efficiently computable (and invertible) permutations E_k on the set of n -bit strings indexed by a κ -bit key k . For example, $n = 64$ and $\kappa = 56$ for DES, and $n = 128$ and $\kappa \in \{128, 192, 256\}$ for AES.

BLOCK-CIPHER SECURITY. Most applications assume and require that the underlying block cipher behaves as a *pseudorandom permutation* (PRP), i.e., under a random secret key, it cannot be efficiently distinguished from a uniformly random permutation. To capture this notion, the PRP-security level of a block cipher is defined as the complexity required to distinguish it from a random permutation with non-negligible advantage.

The security level of a block cipher E is inherently limited by its key length κ : Given very few plaintext-ciphertext pairs $(x_i, E_k(x_i))$, a generic brute-force attack can easily recover the secret key k with roughly 2^κ evaluations of E . This easily yields a PRP distinguishing attack with the same complexity. Clearly, this attack directly affects legacy designs with short keys, such as DES, for which 2^{56} is well within the boundaries of feasible computation.

KEY-LENGTH EXTENSION AND THE ICM. Nonetheless, legacy designs often remain attractive in niche applications, like e.g. in the financial sector, where DES-based constructions are used to encipher PIN numbers due to their short block length (as in the EMV standard [1]). In order to mitigate the effects of the above generic attacks, the well-known *key-length extension* (KLE) problem addresses the following question:

“Does there exist a construction C transforming any (κ, n) -block cipher E into a (κ', n) -block cipher $C[E]$ (for $\kappa' > \kappa$), such that $C[E]$ is secure against generic attackers (using E as a black-box) investing more than 2^κ effort?”

Starting with the work of Killian and Rogaway on DESX [16], and followed by a series of subsequent works [5, 14, 15, 19, 12, 9], KLE has been formalized and studied in the *ideal cipher model* (ICM), where the underlying block cipher is modeled as an *ideal cipher*, i.e., E_k is an independent random permutation for every individual key k . Then, ICM PRP security of a KLE construction $C[E]$ is captured by considering a random experiment where the attacker (also known as a *distinguisher*) issues two types of queries:

- **Block-cipher queries** to evaluate the block cipher $E_k(x)$ and $E_k^{-1}(y)$ for any k, x , and y chosen by the distinguisher.

- **Construction queries** to evaluate on a chosen n -bit input x either the KLE construction $C[E]_{K'}$ with a uniformly random secret κ' -bit key K' , or a uniform random permutation P independent of E . The respective inverses can also be evaluated.

The distinguisher’s goal is to decide whether construction queries are answered by the construction or by P , and its power is measured in terms of the number $q_e \leq 2^{n+\kappa}$ of queries of the former type, and the number $q_c \leq 2^n$ of queries of the latter type. Security of C is measured in terms of which values of (q_c, q_e) do not allow distinguishing with non-negligible advantage.

RELAXING FULL-CODEBOOK SECURITY. So far, the security of KLE constructions (with the notable exception of the work of Killian and Rogaway [16]) has been analyzed in the *full-codebook regime*, i.e., where we allow $q_c = 2^n$, and then see how large q_e can be while still retaining pseudorandomness. However, there is often no rational reason to assume that $q_c = 2^n$. Not only this value is usually unreasonably large, but also, we can either easily restrict the number of block cipher evaluations on a certain secret key at the application level (by enforcing re-keying) *or* when using the KLE construction within a certain mode of operation, the security analysis of the latter may simply force security to only hold for smaller q_c anyway (e.g., $q_c \leq 2^{n/2}$ for CBC modes).

In this paper, we relax the unreasonably strong requirement of full-codebook security, and undertake the first in-depth investigation of the security of KLE constructions in the realistic scenario where $q_c \ll 2^n$.

1.2 Plain and Randomized Cascades

Before we turn to our contributions, let us first review previous works on KLE in the full-codebook regime $q_c = 2^n$. A summary of the attainable PRP security levels is given in Appendix A.

CASCADING-BASED KLE. The most natural KLE approach is perhaps *cascading*, generalizing the idea behind triple DES. Formally, the cascade of length ℓ for a (κ, n) -block cipher E is the $(\ell \cdot \kappa, n)$ -block cipher which takes an $\ell\kappa$ -bit key $\text{mk} = (k_1, \dots, k_\ell) \in (\{0, 1\}^\kappa)^\ell$ and encrypts a plaintext $x \in \{0, 1\}^n$ by computing

$$y = \text{CE}_{\text{mk}}[E](x) = E_{k_\ell} \circ E_{k_{\ell-1}} \circ \dots \circ E_{k_2} \circ E_{k_1}(x).$$

It is well known that the case $\ell = 2$ still allows for a 2^κ -query meet-in-the-middle attack (even though only a smaller distinguishing advantage is achievable for $q_e < 2^\kappa$ as shown by Aiello *et al.* [3]). For the case $r = 3$ (which generalizes 3DES), Bellare and Rogaway [5] first proved PRP security for $q_e \leq 2^{\kappa+n/2}$. This result was later generalized to arbitrary length ℓ by Gaži and Maurer [14]. Their bound on q_e was however far from tight, and was first improved by Lee [19], and a tight bound (matching an attack by Gaži [12]) was only recently given by Dai, Lee, Mennink, and Steinberger [9].

RANDOMIZED CASCADES. Another approach to key-length extension generalizes the DESX construction, using additional key material to randomize inputs and

Table 1. Overview of our results. Parameters q_c, q_e for which we prove security.

Construction	Security
XCE with r rounds	$q_c q_e^r \ll 2^{r(n+\kappa)}$ (Tight)
2XOR	$q_e \ll \max\{2^{\kappa+n/2}, 2^{\kappa+n-\log(q_c)}\}$ (Tight)
3XOR	$q_c \leq 2^{2/3n} : q_e \ll 2^{\kappa + \frac{4-\log(q_e)}{5}n}$ $q_c \in [2^{\frac{2}{3}n}, 2^{\frac{3}{4}n}] : q_e \ll 2^{\kappa + (2-2\log(q_c))n}$
3XSK	$q_c \leq 2^{2/3n} : q_e \ll 2^{\kappa + \frac{4-\log(q_c)}{5}n}$
CE with $\ell = 2r + 1$ rounds	$q_c q_e^r \ll 2^{r(\kappa+n)}, q_c \ll 2^\kappa, q_e \ll 2^{2\kappa}$
Two-key triple encryption	$q_c q_e \ll 2^{\kappa+n}, q_c \ll 2^\kappa, q_e \ll 2^{2\kappa}$

outputs of block-cipher calls. (This technique is often called *whitening*.) For example, the r -round XOR-cascade of a (κ, n) -block cipher E is the $(\kappa + (r + 1)n, n)$ -block cipher which, on input key (k, z) (where $z = (z_0, z_1, \dots, z_r)$ consists of $(r + 1)$ n -bit strings) and message x , returns

$$\text{XCE}[E]((k, z), x) = \oplus_{z_r} \circ E_{\phi_r(k)} \circ \oplus_{z_{r-1}} \circ E_{\phi_{r-1}(k)} \circ \dots \circ \oplus_{z_1} \circ E_{\phi_1(k)} \circ \oplus_{z_0}(x),$$

where \oplus_z maps x' to $x' \oplus z$, and ϕ_1, \dots, ϕ_r are permutations on the κ -bit strings such that $\phi_i(k) \neq \phi_j(k)$ for all k and $i \neq j$. Security bounds for XOR-cascades in the full-codebook regime were proved by Lee [19] and by Gaži [12]. The latter work considered a variant without the last whitening step, and in combination with the result on key-alternating ciphers [8] led to tight bounds.

A simple variant of two-round XOR-cascades, called 2XOR, was studied by Gaži and Tessaro [15], where the third key z_2 is omitted, and $z_0 = z_1$. They prove PRP security for $q_e \leq 2^{\kappa+n/2}$, and that this security level is optimal (for $q_c = 2^n$) with respect to a large class of two-call constructions. We finally emphasize that the work by Killian and Rogaway [16] analyzing DESX (which is the case $r = 1$) is a notable exception to the above restriction to the full-domain regime, and exhibits a smooth security trade-off for any q_c and q_e as long as $q_c \cdot q_e \leq 2^{n+\kappa}$.

1.3 Our Contributions

While tight bounds are known in the full-codebook regime, the landscape is still mostly uncharted when moving to the case $q_c \ll 2^n$. This paper proves lower and upper bounds on the PRP security level of existing and new KLE constructions in the setting where $q_c \ll 2^n$. While a summary of our bounds is given in Table 1, we now discuss our contributions a bit more in detail.

We start with the randomized case:

- **Tight bounds for XOR cascades.** We provide *tight* bounds for XOR-cascades, matching an attack previously given by Gaži [12].
- **Characterizing 2XOR.** We complete the picture of the security of 2XOR for all $q_c \leq 2^n$, showing that $q_e \leq 2^{\kappa+n/2}$ is tight when $q_c \in [2^{n/2}, 2^n]$,

and observing that otherwise $q_c \cdot q_e \leq 2^{\kappa+n}$ is necessary and sufficient for $q_c \leq 2^{n/2}$.

- **The 3XOR construction.** We show that adding the whitening key to the output of 2XOR yields a construction—that we name 3XOR—which is always at least as secure as 2XOR and strictly more secure for $2^{n/4} < q_c < 2^{3/4n}$.
- **A two-call construction with no re-keying.** We finally propose a variant of 3XOR (called 3XSK) where both block-cipher calls are with the *same* key, whereas the middle whitening key is a permutation of the original one. The security is comparable to that of 3XOR for $q_c \leq 2^{2/3n}$.

Our results also improve our picture with respect to plain cascading.

- **Odd-length plain cascades.** We prove that cascades of odd length $\ell = 2r + 1$ are secure whenever $q_c q_e^r \ll 2^{r(\kappa+n)}$, $q_c \ll 2^\kappa$, and $q_e \ll 2^{2\kappa}$. For κ and n satisfying $\kappa \geq \frac{rn}{r+1}$, this improves on the security bound of Dai *et al.* [9] when $q_c \leq 2^{\frac{rn}{r+1}}$. Moreover, when $\kappa \geq n$, this yields a tight bound (matching Gaži’s attack [12]) for *all* parameters (for $\kappa \leq n$, the situation is more involved, see Section 5 for a complete discussion).
- **Two-key triple encryption.** We prove a similar bound for two-key triple encryption, where the first and third keys are identical, as in Triple DES.

OVERVIEW OF OUR TECHNIQUES. It turns out that the techniques behind our results are fairly general. We start by defining a general class of KLE constructions called *randomized KLE schemes*, that capture both plain cascades, XOR-cascades and others. Our core technical tool is then a lemma relating the security of a construction from this class to a particular cipher that can be derived from it, called a *sequential cipher*, also introduced here. Such ciphers constitute a generalization of key-alternating ciphers (or KACs, for short), studied in [6, 24, 17, 4, 18], and implement a block cipher by invoking a number of permutations in a sequential manner. Our lemma generalizes a previous result by Gaži [12], which only considered the case of KACs but neither our relaxation to randomized KLEs, nor the case without a full codebook.

To instantiate some of our bounds, we provide a generalized analysis of sequential ciphers, extending recent bounds by Chen and Steinberger [8].

1.4 Further Related Works

We note in passing that an orthogonal line of works devoted to cascade-like construction was initiated by Luby and Rackoff [20]. These works study *standard model* security amplification achieved by plain and randomized cascades, and in particular show how, when instantiated with a block cipher which is a weak PRP in the sense of the attacker achieving a large distinguishing advantage, these constructions reduce the best possible advantage with an increasing number of rounds. Increasingly tighter bounds have been given by Maurer and Tessaro [22], and by Tessaro [25]. An information-theoretic version of this question was also studied [26, 21].

2 Preliminaries

2.1 Basic Notation and Block Ciphers

In all the following, we fix integers $n, \kappa > 0$, and denote $N = 2^n$ and $K = 2^\kappa$. The set of all permutations on $\{0, 1\}^n$ will be denoted \mathcal{P}_n . For a set T and an integer $\ell \geq 1$, $(T)_\ell$ denotes the set of all sequences that consist of ℓ distinct elements of T . For integers $1 \leq \ell \leq t$, we will write $(t)_\ell = t(t-1) \cdots (t-\ell+1)$. If $|T| = t$, then $(t)_\ell$ becomes the size of $(T)_\ell$.

A block cipher is a function family $E : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that for all $k \in \mathcal{K}$ the mapping $E(k, \cdot)$ is a permutation on $\{0, 1\}^n$. We denote by $\text{BC}(\mathcal{K}, n)$ the set of all such block ciphers, shortening to $\text{BC}(\kappa, n)$ when $\mathcal{K} = \{0, 1\}^\kappa$. In the ideal-cipher model, a block cipher E is chosen from $\text{BC}(\kappa, n)$ uniformly at random and made available to the participants through oracle queries. It allows for two types of oracle queries $E(k, x)$ and $E^{-1}(k, y)$ for $x, y \in \{0, 1\}^n$ and $k \in \{0, 1\}^\kappa$.¹ The answer to an inverse query $E^{-1}(k, y)$ is $x \in \{0, 1\}^n$ such that $E(k, x) = y$.

For a set $\mathcal{Q} = ((x_1, y_1), \dots, (x_q, y_q)) \in (\{0, 1\}^n \times \{0, 1\}^n)^q$ and a permutation $P \in \mathcal{P}_n$, we say that P extends \mathcal{Q} , denoted $P \vdash \mathcal{Q}$, if $P(x_i) = y_i$ for $i = 1, \dots, q$. The *domain* and the *range* of \mathcal{Q} are defined as

$$\text{Dom}(\mathcal{Q}) = \{x \in \{0, 1\}^n : (x, y) \in \mathcal{Q}\}, \text{Rng}(\mathcal{Q}) = \{y \in \{0, 1\}^n : (x, y) \in \mathcal{Q}\}$$

respectively. By an abuse of notation, we will sometimes denote \mathcal{Q} the bijection from $\text{Dom}(\mathcal{Q})$ to $\text{Rng}(\mathcal{Q})$ such that $\mathcal{Q}(x_i) = y_i$ for $i = 1, \dots, q$. Thus, for another set (bijection) $\mathcal{Q}' \in (\{0, 1\}^n \times \{0, 1\}^n)^{q'}$, we have

$$\begin{aligned} \text{Dom}(\mathcal{Q}' \circ \mathcal{Q}) &= \{x \in \{0, 1\}^n : (x, y) \in \mathcal{Q} \wedge y \in \text{Dom}(\mathcal{Q}')\} \\ \text{Rng}(\mathcal{Q}' \circ \mathcal{Q}) &= \{y \in \{0, 1\}^n : (x, y) \in \mathcal{Q}' \wedge x \in \text{Rng}(\mathcal{Q})\} \end{aligned}$$

and similar definitions for the composition of more than two bijections. For a set $\mathcal{Q} = ((x_1, k_1, y_1), \dots, (x_q, k_q, y_q)) \in (\{0, 1\}^n \times \{0, 1\}^\kappa \times \{0, 1\}^n)^q$ and a block cipher $E \in \text{BC}(\kappa, n)$, we say that E extends \mathcal{Q} , denoted $E \vdash \mathcal{Q}$, if $E(k_i, x_i) = y_i$ for $i = 1, \dots, q$.

2.2 Indistinguishability in Idealized Models

In this paper, we consider block ciphers that are built (in a black-box way) on top of an existing primitive F . The primitive F is modeled as an ideal oracle (publicly accessible to the adversary), whose answers follow some probability distribution. Namely, we will consider two slightly different settings: the so-called *KLE-setting* where F will be the ideal cipher E ; and the so-called *SC-setting*² where F will be a tuple of random permutations $\mathbf{P} = (P_1, \dots, P_m)$. In both settings, we consider

¹ We interchangeably use both notations $E(k, x)$ and $E_k(x)$, and similarly $E^{-1}(k, y)$ and $E_k^{-1}(y)$.

² This refers to the notion of a *sequential cipher* defined in Section 3.1.

a construction C which encrypts a message $x \in \{0, 1\}^n$ with some (master) key³ $\text{mk} \in \{0, 1\}^{\kappa'}$ by making calls to F , and denote $C[F]$ the resulting block cipher (hence $C[F] \in \text{BC}(\kappa', n)$, and $C_{\text{mk}}[F]$ is the permutation associated to key mk).

To define security, we consider an adversary (*a.k.a.* distinguisher) D which interacts with a pair of oracles that we denote generically (P, F) . The goal of D is to distinguish whether it is interacting with $(C_{\text{mk}}[F], F)$ for some uniformly random key mk (a case we will informally refer to as the “real” world) or with (P, F) where P is a random n -bit permutation independent from F (the “ideal” world). Note that in both worlds the first oracle P is a permutation that can be queried in both directions. The distinguisher’s advantage is defined as

$$\text{Adv}_C^{\text{cca}}(D) = \left| \Pr \left[D^{C_{\text{mk}}[F], F} = 1 \right] - \Pr \left[D^{P, F} = 1 \right] \right|$$

where the first probability is taken over the random choice of mk and the random answers of F , and the second probability is taken over the random choice of P and F . We refer to D ’s queries to its first and second oracle as *construction* and *primitive* queries, respectively. In the KLE-setting (SC-setting), the primitive queries are sometimes referred to more concretely as block-cipher queries (permutation queries), respectively.

In the KLE-setting, for $q_c, q_e \geq 0$ we define

$$\text{Adv}_C^{\text{cca}}(q_c, q_e) = \max_D \text{Adv}_C^{\text{cca}}(D)$$

where the maximum is taken over all distinguishers making exactly q_c construction and q_e ideal-cipher queries. Similarly, in the SC-setting, for $q_c, q_p \geq 0$,

$$\text{Adv}_C^{\text{cca}}(q_c, q_p) = \max_D \text{Adv}_C^{\text{cca}}(D)$$

where the maximum is taken over all distinguishers making exactly q_c construction queries and q_p permutation queries to *each* permutation oracle P_i .

In all the paper, we assume that the distinguisher is computationally unbounded, deterministic, and that it never makes redundant queries (these last two assumptions being *wlog*). In accordance with several recent works on the topic, we are using Patarin’s H-coefficients technique [23] in some of our proofs. Our use of the H-coefficients technique will be self-contained, for a more detailed introduction to this method see for example [8].

3 From Randomized KLE Schemes to Sequential Ciphers

In this section we study the relationship of two general classes of constructions that we first define. On one hand, we consider *randomized KLE schemes* that

³ We use the wording *master key* to emphasize that it will usually be used to derive sub-keys for calling the underlying block cipher. We also write $\text{mk} = (k, z)$ for key-length extension schemes and $\text{mk} = z$ for sequential ciphers (see Section 3.1).

generalize both cascades and XOR-cascades,⁴ on the other hand we introduce *sequential ciphers* that are in turn a generalization of key-alternating ciphers (whose definition is also provided below). We show that every randomized KLE scheme induces a sequential cipher and the security properties of these two constructions are tightly connected.

3.1 Definitions

RANDOMIZED KLE. Let $n, \kappa > 0$ be some fixed parameters denoting the block- and key-length of the underlying block cipher, respectively. Fix additional parameters $\lambda, r, m > 0$. Let (ϕ_1, \dots, ϕ_m) be m permutations of $\{0, 1\}^\kappa$ with the property that for any $k \in \{0, 1\}^\kappa$, the values $(\phi_1(k), \dots, \phi_m(k))$ are distinct. (Note that this imposes $m \leq 2^\kappa$.) Let $\sigma : \{1, \dots, r\} \rightarrow \{1, \dots, m\}$ be a surjective function.⁵ For $i = 0, \dots, r$, let

$$\rho^i : \{0, 1\}^\lambda \times \{0, 1\}^n \rightarrow \{0, 1\}^n$$

be a function such that for each $z \in \{0, 1\}^\lambda$, $\rho^i(z, \cdot)$ (also denoted $\rho_z^i(\cdot)$) is a permutation on $\{0, 1\}^n$.⁶

A *randomized key-length extension scheme* \mathbf{R} transforms a block cipher $E \in \text{BC}(\kappa, n)$ into a new block cipher $\mathbf{R}[E] \in \text{BC}(\kappa + \lambda, n)$ specified as follows: for a plaintext $x \in \{0, 1\}^n$ and a key $(k, z) \in \{0, 1\}^\kappa \times \{0, 1\}^\lambda$, the ciphertext is defined as (see Figure 1)

$$\mathbf{R}[E]((k, z), x) = \rho_z^r \circ E_{k_{\sigma(r)}} \circ \rho_z^{r-1} \circ E_{k_{\sigma(r-1)}} \circ \dots \circ E_{k_{\sigma(2)}} \circ \rho_z^1 \circ E_{k_{\sigma(1)}} \circ \rho_z^0(x).$$

where we simply write $(k_1, \dots, k_m) = (\phi_1(k), \dots, \phi_m(k))$. For a fixed key (k, z) , we also denote $\mathbf{R}_{k,z}[E]$ the permutation $x \mapsto \mathbf{R}[E]((k, z), x)$.

SEQUENTIAL CIPHER. With the same primitives σ and (ρ^0, \dots, ρ^r) , a *sequential cipher* \mathbf{S} transforms a set of permutations $\mathbf{P} = (P_1, \dots, P_m)$ into a block cipher $\mathbf{S}[\mathbf{P}] \in \text{BC}(\lambda, n)$ specified as follows: for a plaintext $x \in \{0, 1\}^n$ and a key $z \in \{0, 1\}^\lambda$, the ciphertext is defined as (again see Figure 1)

$$\mathbf{S}[\mathbf{P}](z, x) = \rho_z^r \circ P_{\sigma(r)} \circ \rho_z^{r-1} \circ P_{\sigma(r-1)} \circ \dots \circ P_{\sigma(2)} \circ \rho_z^1 \circ P_{\sigma(1)} \circ \rho_z^0(x).$$

For a fixed key z , we denote $\mathbf{S}_z[\mathbf{P}]$ the permutation $x \mapsto \mathbf{S}[\mathbf{P}](z, x)$.

⁴ Our randomized KLE schemes also cover the notion of *sequential constructions* introduced in [12]. Since the latter are KLE schemes, they syntactically differ from the notion of a sequential cipher considered here.

⁵ In case this causes confusion, r is the number of rounds, m is the number of distinct keys that are used to call the underlying block cipher, and σ specifies which key is used at each round.

⁶ Though each ρ^i is syntactically a block cipher, we prefer to avoid this wording since in most of the paper the ρ^i 's will be much simpler than E , the block cipher underlying the key-length extension scheme.




Fig. 1. The randomized key-length extension construction $R[E]$ (top), and its induced sequential cipher $\bar{R}[P]$ (bottom).

3.2 Induced Sequential Ciphers

When the key k is fixed in some key-length extension scheme R , the resulting scheme can be regarded as a sequential cipher with key space $\{0, 1\}^\lambda$ using independent random permutations P_1, \dots, P_m in place of $E_{\phi_1(k)}, \dots, E_{\phi_m(k)}$ in the ideal cipher model. We formalize this remark as follows.

Definition 1. Let R be a randomized key-length extension scheme defined as above. The induced sequential cipher of R , denoted \bar{R} , is a sequential cipher which specifies a block cipher $\bar{R}[P] \in \text{BC}(\lambda, n)$ from an m -tuple of permutations $P = (P_1, \dots, P_m)$ of $\{0, 1\}^n$ by replacing each call to $E(\phi_i(k), \cdot)$, resp. $E^{-1}(\phi_i(k), \cdot)$ when computing $R_{k,z}[E](x)$ by a call to $P_i(\cdot)$, resp. $P_i^{-1}(\cdot)$ in the computation of $\bar{R}_z[P](x)$.

Example 1. If we let σ be the identity, $\lambda = (r + 1)n$ and $\rho^i(z, u) = u \oplus z_i$ for $i = 0, \dots, r$ where z is split as $z = (z_0, \dots, z_r) \in (\{0, 1\}^n)^{r+1}$, then the resulting randomized KLE and sequential cipher constructions are called an *XOR-cascade scheme* and a *key-alternating cipher (KAC)*, respectively.

More formally, the r -round XOR-cascade construction XCE turns a block cipher $E \in \text{BC}(\kappa, n)$ into a new block cipher $\text{XCE}[E] \in \text{BC}(\kappa + (r + 1)n, n)$ as follows. Let (ϕ_1, \dots, ϕ_r) be r permutations of $\{0, 1\}^\kappa$ with the property that for any $k \in \{0, 1\}^\kappa$, the values $(\phi_1(k), \dots, \phi_r(k))$ are distinct. Then for a plaintext $x \in \{0, 1\}^n$ and a key $(k, z) \in \{0, 1\}^\kappa \times (\{0, 1\}^n)^{r+1}$ with $z = (z_0, \dots, z_r)$, the ciphertext is defined as (see also Figure 2):

$$\text{XCE}[E]((k, z), x) = \oplus_{z_r} \circ E_{\phi_r(k)} \circ \oplus_{z_{r-1}} \circ E_{\phi_{r-1}(k)} \circ \dots \circ \oplus_{z_1} \circ E_{\phi_1(k)} \circ \oplus_{z_0}(x),$$

where \oplus_{z_i} denotes the mapping $x \mapsto x \oplus z_i$. Its induced sequential cipher $\bar{\text{XCE}}$ is the key-alternating cipher (hence denoted KAC)

$$\bar{\text{XCE}}[P](z, x) = \text{KAC}[P](z, x) = \oplus_{z_r} \circ P_r \circ \oplus_{z_{r-1}} \circ P_{r-1} \circ \dots \circ \oplus_{z_1} \circ P_1 \circ \oplus_{z_0}(x).$$

A tight bound for the security of key-alternating ciphers was given in [8]. We show how to extend their approach to the more general case of sequential ciphers, as this will turn out to be useful in our later proofs. Due to space constraints, we present this extension in the full version of this paper [13].

Example 2. A more specialized case of the XOR-cascade scheme (i.e., taking $r = 1$) is the FX construction of [16] (the generic variant of DESX, attributed to Rivest) which turns a block cipher $E \in \text{BC}(\kappa, n)$ into a block cipher $\text{FX}[E] \in \text{BC}(\kappa + 2n, n)$ defined as

$$\text{FX}_{k,(z_0,z_1)}[E](x) = E_k(x \oplus z_0) \oplus z_1.$$

The resulting construction $\overline{\text{FX}}$, for a permutation $P \in \mathcal{P}_n$, is

$$\overline{\text{FX}}_{(z_0,z_1)}[P](x) = P(x \oplus z_0) \oplus z_1,$$

which is exactly the Even-Mansour cipher [11].

3.3 The Reduction

In this section we prove our main lemma that reduces the security of a randomized key-length extension scheme to the security of the corresponding induced sequential cipher. It can be seen as a generalization of [12, Theorem 2] to more general classes of constructions and as well to the setting where the number of construction queries q_c is arbitrary (rather than $q_c = 2^n$).

Lemma 1. *Let \mathbb{R} be a randomized key-length extension scheme and let $\overline{\mathbb{R}}$ be its induced sequential cipher. Then for $q_c, q_e, M > 0$, one has*

$$\text{Adv}_{\mathbb{R}}^{\text{cca}}(q_c, q_e) \leq \frac{mq_e}{KM} + \text{Adv}_{\overline{\mathbb{R}}}^{\text{cca}}(q_c, M).$$

Proof. Consider a distinguisher \mathbb{D} interacting with (P, E) , where E is an ideal cipher and P is either the construction $\mathbb{R}_{k,z}[E]$ for a uniformly random key $(k, z) \in \{0, 1\}^\kappa \times \{0, 1\}^\lambda$, or a random permutation independent from E . Following the H-coefficients technique [23, 8], we summarize all the information gathered by the distinguisher when interacting with the system (P, E) in the *raw query transcript* which is simply the ordered list of queries of \mathbb{D} to its oracles together with their answers. From this raw query transcript we can build the *construction query transcript*

$$\mathcal{Q}_C = ((x_1, y_1), \dots, (x_{q_c}, y_{q_c})),$$

where the i -th pair (x_i, y_i) indicates that the i -th query to the construction/random permutation oracle was either $P(x_i)$ with answer y_i or $P^{-1}(y_i)$ with answer x_i . Similarly, we can build the *ideal cipher query transcript*

$$\mathcal{Q}_E = ((u_1, k_1, v_1), \dots, (u_{q_e}, k_{q_e}, v_{q_e})),$$

where the i -th triple (u_i, k_i, v_i) indicates that the i -th query to the ideal cipher was either $E(k_i, u_i)$ with answer v_i or $E^{-1}(k_i, v_i)$ with answer u_i . (Since the distinguisher is deterministic, the raw query transcript can unambiguously be reconstructed from the pair $(\mathcal{Q}_C, \mathcal{Q}_E)$.)

Moreover, in the real world, the key k (but not z) is given for free to D at the end of its queries, while in the ideal world (where no such key exists), a dummy key k is drawn uniformly at random and given to D . (This can only *increase* the distinguishing advantage since D can disregard this additional information.) This results in what we simply call the *transcript* $\tau = (\mathcal{Q}_C, \mathcal{Q}_E, k)$ of the attack. We will say that a transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_E, k)$ is *attainable* if there exists a permutation P and a block cipher E such that the interaction of D with (P, E) yields queries transcripts $(\mathcal{Q}_C, \mathcal{Q}_E)$ (said otherwise, the probability to obtain this transcript in the “ideal” world is non-zero). Finally, we let T_{re} , resp. T_{id} denote the probability distribution of the transcript τ induced by the real world, resp. the ideal world (note that these two probability distributions depend on the distinguisher). By extension, we use the same notation to denote a random variable distributed according to each distribution.

Let D be an optimal distinguisher making q_c construction queries and q_e ideal-cipher queries such that⁷

$$\mathbf{Adv}_{\mathbf{R}}^{\text{cca}}(q_c, q_e) = \mathbf{Adv}_{\mathbf{R}}^{\text{cca}}(D) = \sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{id}} = \tau] - \sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{re}} = \tau]$$

and let \mathcal{T}_1 be the set of attainable transcripts $\tau = (\mathcal{Q}_C, \mathcal{Q}_E, k)$ such that the distinguisher outputs 1 when obtaining τ . Given an ideal-cipher queries transcript \mathcal{Q}_E , we also define the set of bad keys as

$$\text{Bad}(\mathcal{Q}_E) = \{k \in \{0, 1\}^k : |\{(x, y) : (x, k, y) \in \mathcal{Q}_E\}| > M\}.$$

(Hence, a key k is bad if it appears strictly more than M times in \mathcal{Q}_E .) We say that an attainable transcript $\tau = (\mathcal{Q}_C, \mathcal{Q}_E, k)$ is *bad* if $\phi_i(k) \in \text{Bad}(\mathcal{Q}_E)$ for some $i = 1, \dots, m$, and *good* otherwise. We denote resp. \mathcal{T}_{bad} and $\mathcal{T}_{\text{good}}$ the sets of bad and good transcripts (which form a partition of the set of attainable transcripts \mathcal{T}). Then we have

$$\begin{aligned} \mathbf{Adv}_{\mathbf{R}}^{\text{cca}}(D) &= \sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{id}} = \tau] - \sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{re}} = \tau] \\ &\leq \sum_{\tau \in \mathcal{T}_{\text{bad}}} \Pr[T_{\text{id}} = \tau] + \sum_{\tau \in \mathcal{T}_1 \cap \mathcal{T}_{\text{good}}} \Pr[T_{\text{id}} = \tau] - \sum_{\tau \in \mathcal{T}_1 \cap \mathcal{T}_{\text{good}}} \Pr[T_{\text{re}} = \tau] \quad (1) \end{aligned}$$

where the inequality follows from the fact that $\mathcal{T}_{\text{good}}$ and \mathcal{T}_{bad} form a partition of the set of attainable transcripts \mathcal{T} . We upper bound each summand in turn.

Since in the ideal world the key k is drawn uniformly at random at the end of the interaction of the distinguisher with its oracles, we clearly can bound

⁷ Without loss of generality, we can assume $\sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{id}} = \tau] \geq \sum_{\tau \in \mathcal{T}_1} \Pr[T_{\text{re}} = \tau]$ by slightly modifying D if necessary.

$\sum_{\tau \in \mathcal{T}_{\text{bad}}} \Pr[T_{\text{id}} = \tau] = \Pr[T_{\text{id}} \in \mathcal{T}_{\text{bad}}]$ as

$$\Pr[T_{\text{id}} \in \mathcal{T}_{\text{bad}}] \leq \sum_{i=1}^m \Pr[k \leftarrow_{\S} \{0, 1\}^{\kappa} : \phi_i(k) \in \text{Bad}(\mathcal{Q}_E)] \leq \frac{mq_e}{KM}, \quad (2)$$

where the last inequality follows from the fact that each ϕ_i is a permutation (hence $\phi_i(k)$ is uniformly random) and that the size of $\text{Bad}(\mathcal{Q}_E)$ is at most q_e/M by definition.

To upper bound the second term, we consider the following (probabilistic) distinguisher \bar{D} against construction \bar{R} (in the random permutation model), which uses D as a subroutine. \bar{D} has access to $m+1$ permutation oracles (P_0, P_1, \dots, P_m) , where P_0 is either the construction $\bar{R}_z[P_1, \dots, P_m]$ for some random key $z \leftarrow_{\S} \{0, 1\}^{\lambda}$, or a random permutation independent from (P_1, \dots, P_m) . At the beginning of the experiment, \bar{D} draws a key $k \leftarrow_{\S} \{0, 1\}^{\kappa}$ uniformly at random. Then, \bar{D} runs D and answers its queries as follows. First, it relays any construction query from D to its own construction oracle and relays back the corresponding answer to D . When D makes any ideal cipher query for some key $k' \notin \{\phi_1(k), \dots, \phi_m(k)\}$, \bar{D} simulates a perfectly random permutation associated with k' . If D makes an ideal cipher query for some key $\phi_i(k)$, $i = 1, \dots, m$, \bar{D} relays this query to permutation oracle P_i and forwards the corresponding answer to D . However, if D attempts to make more than M queries corresponding to some key $\phi_i(k)$, $i = 1, \dots, m$, then \bar{D} aborts and outputs 0. (Hence \bar{D} always makes at most M queries to each permutation oracle P_i , $i = 1, \dots, m$.) Otherwise, once D has finished its queries, \bar{D} forwards k to D (recall that we include k in the transcript) and outputs the same value as D . Clearly, when \bar{D} is interacting with (P_0, P_1, \dots, P_m) , where P_0 is the construction $\bar{R}_z[P_1, \dots, P_m]$ then it is perfectly simulating the real world $(R_{k,z}[E], E)$ to D , while when \bar{D} is interacting with (P_0, P_1, \dots, P_m) where P_0 is independent from (P_1, \dots, P_m) , then it is perfectly simulating the ideal world (P, E) to D . Hence, the distinguishing advantage of \bar{D} is

$$\text{Adv}_{\bar{R}}^{\text{cca}}(\bar{D}) = \left| \sum_{\tau \in \mathcal{T}_1 \cap \mathcal{T}_{\text{good}}} \Pr[T_{\text{re}} = \tau] - \sum_{\tau \in \mathcal{T}_1 \cap \mathcal{T}_{\text{good}}} \Pr[T_{\text{id}} = \tau] \right|. \quad (3)$$

Since \bar{D} makes at most q_c queries to its construction and at most M queries to each permutation oracle P_i , $i = 1, \dots, m$, and since in the information-theoretic setting the advantage of a probabilistic adversary cannot be larger than the one of the best deterministic adversary, one has

$$\text{Adv}_{\bar{R}}^{\text{cca}}(\bar{D}) \leq \text{Adv}_{\bar{R}}^{\text{cca}}(q_c, M). \quad (4)$$

Combining (1), (2), (3), and (4), we obtain

$$\text{Adv}_{\bar{R}}^{\text{cca}}(D) = \text{Adv}_{\bar{R}}^{\text{cca}}(q_c, q_e) \leq \frac{mq_e}{KM} + \text{Adv}_{\bar{R}}^{\text{cca}}(q_c, M). \quad \square$$

The following corollary can be easily obtained after optimization of M in Lemma 1 when one has a simple enough upper bound on $\text{Adv}_{\bar{R}}^{\text{cca}}(q_c, q_p)$.




Fig. 2. The XOR-cascade key-length extension scheme $\text{XCE}[E]$.

Corollary 1. *Let \mathbf{R} be a randomized key-length extension scheme and let $\bar{\mathbf{R}}$ be its induced sequential cipher. Assume that*

$$\text{Adv}_{\mathbf{R}}^{\text{cca}}(q_c, q_p) \leq A + B \frac{q_c^\alpha q_p^\beta}{N^\gamma},$$

where $A, B, \alpha, \beta, \gamma$ do not depend on q_p , and $B \geq 1, \beta \geq 1$. Then

$$\text{Adv}_{\mathbf{R}}^{\text{cca}}(q_c, q_e) \leq A + mB(\beta + 1) \left(\frac{q_c^\alpha q_e^\beta}{K^\beta N^\gamma} \right)^{\frac{1}{\beta+1}}.$$

4 Randomized Key-Length Extension Schemes

In this section, we derive security bounds for various randomized KLE schemes.

4.1 XOR-Cascades: Tight Bounds

As a first application, we complete the picture of the security of the XOR-cascade key-length extension scheme with independent whitening keys introduced in [12]. We derive a tight security bound for the setting with less than 2^n construction queries. Recall the definition of the r -round XOR-cascade construction XCE given in Section 3.2, Example 1.

Lemma 1 shows that the security of the r -round XOR-cascade construction is directly related to the security of the corresponding r -round key-alternating cipher KAC. It was observed in [12] to be related to the security of the $(r - 1)$ -round key-alternating cipher, but in hindsight this rather appears as an artifact of the setting $q_c = 2^n$.

Combining our improved result on the security of KAC (following from [8] and given in the full version of this paper [13]) with Lemma 1, we obtain that for any integer M such that $q_c + M \leq N/2$,

$$\text{Adv}_{\text{XCE}}^{\text{cca}}(q_c, q_e) \leq \frac{rq_e}{KM} + 4(r + 2) \left(\frac{rq_c M^r}{(r + 2)N^r} \right)^{\frac{1}{r+1}}.$$

After the optimization of M (by equating the two summands), we arrive at the following theorem.

Theorem 1. Consider the r -round XOR-cascade construction XCE. Then

$$\text{Adv}_{\text{XCE}}^{\text{cca}}(q_c, q_e) \leq C_r \left(\frac{q_c q_e^r}{K^r N^r} \right)^{\frac{1}{2r+1}},$$

where C_r is a constant that depends only on r , namely

$$C_r = (2^{4r+3} \cdot r^{r+1} (r+2)^r)^{\frac{1}{2r+1}} \in \mathcal{O}(r).$$

In short, XOR-cascade encryption is secure as long as $q_c q_e^r$ is small compared to $2^{r(\kappa+n)}$. We note that this security bound is matched by a generic attack on sequential constructions given in [12, Theorem 3], since this attack can be easily generalized for arbitrary q_c as observed there.

4.2 2XOR: Tight Bounds

The construction 2XOR was proposed by Gaži and Tessaro [15] to turn a block cipher $E \in \text{BC}(\kappa, n)$ into a new block cipher $2\text{XOR}[E] \in \text{BC}(\kappa+n, n)$ defined as

$$2\text{XOR}_{k,z}[E](x) = E_{\phi(k)}(E_k(x \oplus z) \oplus z),$$

where ϕ is any (fixed) permutation of $\{0, 1\}^\kappa$ without fixed points. They showed the following result.

Theorem 2 ([15, Theorem 3]). For any integer q_e ,

$$\text{Adv}_{2\text{XOR}}^{\text{cca}}(q_c = 2^n, q_e) \leq 4 \cdot \left(\frac{q_e}{2^{\kappa+n/2}} \right)^{\frac{2}{3}}.$$

We describe how to attack 2XOR for any $1 \leq c \leq n/2$ using roughly 2^c construction queries and $2^{\kappa+n-c}$ block-cipher queries.

Theorem 3. Let $1 \leq c \leq n/2$ and $1 \leq t \leq c$ be integers such that t is even. There exists a distinguisher D which makes at most $q_c = 2^{c+t/2}$ (forward) construction queries and $q_e = 2^{\kappa+n-c+t/2+1}$ ideal cipher queries, and which achieves

$$\text{Adv}_{2\text{XOR}}^{\text{cca}}(\text{D}) \geq 1 - 2^{\kappa+n-2^t(n-1)}.$$

In particular, for $c = n/2$ and $t = \lfloor \log_2(\kappa/n+1) \rfloor + 1$, its advantage is negligibly close to 1 (asymptotically in n), and its complexity is $q_c = \mathcal{O}(2^{n/2})$ and $q_e = \mathcal{O}(2^{\kappa+n/2})$ (for κ/n constant).

Proof. Consider the distinguisher D depicted in Figure 3 (we assume n to be even for simplicity). For its analysis, first note that for any $z \in \{0, 1\}^n$, the size of the set V_z determined on line 10 is exactly 2^t , due to the choice of the sets X and U . When D interacts with $(2\text{XOR}_{k,z}[E], E)$, it always outputs 1 since the check on line 11 succeeds for the real key (k, z) . In the ideal world (P, E) , we can upper-bound the probability that the distinguisher outputs 1 as follows: for each key (k, z) , the values $\tilde{v}(k, u)$ and $\tilde{u}(k, x)$ for the 2^t pairs $(x, u) \in V_z$ are independent, so that the probability that the check on line 11 succeeds is exactly $\frac{1}{(2^n)_{2^t}} \leq 2^{-2^t(n-1)}$. By the union bound over the $2^{\kappa+n}$ pairs (k, z) , the probability that D returns 1 is at most $2^{\kappa+n-2^t(n-1)}$. \square

Distinguisher $D^{S,E}$: 1: let $X := \{x'0^{n-c-t/2} : x' \in \{0,1\}^{c+t/2}\} \subseteq \{0,1\}^n$ 2: let $U := \{0^{c-t/2}u' : u' \in \{0,1\}^{n-c+t/2}\} \subseteq \{0,1\}^n$ 3: for all $x \in X$ do 4: query $y(x) := S(x)$ 5: for all $(k, x) \in \{0,1\}^\kappa \times X$ do 6: query $\tilde{u}(k, x) := E_{\phi(k)}^{-1}(y(x))$ 7: for all $(k, u) \in \{0,1\}^\kappa \times U$ do 8: query $\tilde{v}(k, u) := E_k(u)$ 9: for all $(k, z) \in \{0,1\}^\kappa \times \{0,1\}^n$ do 10: let $V_z := \{(x, u) \in X \times U \mid x \oplus u = z\}$ 11: if $\forall (x, u) \in V_z : \tilde{v}(k, u) \oplus \tilde{u}(k, x) = z$ then 12: return 1 13: return 0	where $S \in \{2\text{XOR}_{k,z}[E], P\}$
---	---

Fig. 3. Distinguisher D for the proof of Theorem 3, attacking the construction 2XOR and parametrized by c, t .

To illustrate that the tradeoff $q_c q_e \leq 2^{\kappa+n}$ imposed by the attack above is tight for $0 \leq \log_2(q_c) \leq n/2$, consider the sequential cipher $\overline{2\text{XOR}}$ induced by 2XOR. By a trivial reduction (simulating its last, independent random permutation), one can show that 2XOR is at least as secure as the Even-Mansour cipher $\overline{\text{FX}}$ described in Example 2 in Section 3.2. However, it follows from [11] that $\overline{\text{FX}}$ is secure as long as $q_c q_p \leq 2^n$, which, via $\overline{2\text{XOR}}$ and the application of Lemma 1, implies that 2XOR is secure roughly as long as $q_c q_e \leq 2^{\kappa+n}$. This completes the picture for 2XOR on the interval $0 \leq \log_2(q_c) \leq n/2$. For any $q_c \geq 2^{n/2}$, a tight bound for 2XOR is $q_e = 2^{\kappa+n/2}$ as follows from Theorems 2 and 3, hence the security of 2XOR is now understood for the full spectrum of parameters (q_c, q_e) (see Figure 5).

4.3 3XOR: Final Whitening Step Helps

It was also argued in [15] that the 2XOR construction has optimal security within a large class of (so-called sequential) two-query constructions in the following sense: They give a generic attack on any construction from this class requiring roughly 2^n construction queries and $2^{\kappa+n/2}$ block-cipher queries, hence matching the security bound from Theorem 2. However, this only shows the optimality of the 2XOR construction (and in particular, no need to add a final XOR step at its end) in the setting where $q_c = 2^n$ is assumed. As we show below, the situation changes as soon as we also consider lower values of q_c . In this general case, adding a third randomization step actually *does improve* security for some range of the parameters (q_c, q_e) .

We define the 3XOR construction similarly to the 2XOR construction, but with a final whitening step (see Figure 4), i.e.,




Fig. 4. The $3\text{XOR}[E]$ key-length extension scheme.

$$3\text{XOR}_{k,z}[E](x) = E_{\phi(k)}(E_k(x \oplus z) \oplus z) \oplus z .$$

Note that 3XOR is simply the 2-round XCE construction with identical whitening keys in-between the block-cipher calls. The induced sequential cipher $\overline{3\text{XOR}}$ is hence the 2-round Even-Mansour cipher with independent permutations and identical round keys:

$$\overline{3\text{XOR}}_z[P_1, P_2](x) = P_2(P_1(x \oplus z) \oplus z) \oplus z .$$

The security of this construction was analyzed by Chen *et al.* [7]. We recall their result in the full version of this paper [13]. Combining it with Corollary 1, we obtain the following theorem for the security of the 3XOR construction.

Theorem 4. *Assume that $n \geq 11$, $q_c \geq 9n$, $q_e \geq 9n$, and $2q_c + 2q_e \leq N$. Then the following upper bounds hold:*

(i) *When $q_c \leq 2^{\frac{n}{4}}$, one has*

$$\mathbf{Adv}_{3\text{XOR}}^{\text{cca}}(q_c, q_e) \leq 24 \left(\frac{q_c q_e}{KN} \right)^{\frac{1}{2}} .$$

(ii) *When $2^{\frac{n}{4}} \leq q_c \leq 2^{\frac{2n}{3}}$, one has*

$$\mathbf{Adv}_{3\text{XOR}}^{\text{cca}}(q_c, q_e) \leq \frac{6}{N} + 4 \times (13 + 9\sqrt{n}) \left(\frac{q_c^{\frac{1}{5}} q_e}{KN^{\frac{4}{5}}} \right)^{\frac{1}{2}} .$$

(iii) *When $2^{\frac{2n}{3}} \leq q_c \leq 2^{\frac{3n}{4}}$, one has*

$$\mathbf{Adv}_{3\text{XOR}}^{\text{cca}}(q_c, q_e) \leq \frac{6}{N} + 4 \times (13 + 9\sqrt{n}) \left(\frac{q_c^2 q_e}{KN^2} \right)^{\frac{1}{2}} .$$

(iv) *When $q_c \geq 2^{\frac{3n}{4}}$, one has,*

$$\mathbf{Adv}_{3\text{XOR}}^{\text{cca}}(q_c, q_e) \leq \frac{1}{eN} + 6n \left(\frac{q_e}{KN^{\frac{1}{2}}} \right)^{\frac{2}{3}} .$$

This security bound is qualitatively similar to the one of 2XOR for $q_c \leq 2^{\frac{n}{4}}$ and $q_c \geq 2^{\frac{3n}{4}}$, but strictly better for $2^{\frac{n}{4}} \leq q_c \leq 2^{\frac{3n}{4}}$ (see Figure 5). Regarding the tightness of the bound, we note that the general attack against sequential




Fig. 5. The security of the 3XOR key-length extension scheme. All parameters below the (red) solid line are secure due to Theorem 4, while all parameters above the (black) dashed line are insecure due to the attack [12]. The status for parameters between these two lines remains unknown. The (blue) dotted line (which merges with the red solid line for $q_c \leq 2^{\frac{n}{4}}$ and $q_c \geq 2^{\frac{3n}{4}}$) also indicates the (tight) security bound for 2XOR.

constructions given in [12, Theorem 3] applies to 3XOR, so that for any q_c , the construction is insecure for $q_e \approx 2^{\kappa+n-\frac{1}{2}\log_2 q_c}$. This matches the security bound for the special cases $q_c \approx 1$, $q_c \approx 2^{\frac{2n}{3}}$, and $q_c \approx 2^n$ (see Figure 5).

In conclusion, our results in Sections 4.2 and 4.3 show that 3XOR is *always* at least as secure as 2XOR for all possible values of q_c , and strictly more secure for $2^{n/4} < q_c < 2^{3n/4}$.

4.4 3XSK: A 2-Call Construction without Rekeying

A drawback of the 3XOR construction is that the underlying block cipher E is called under two distinct keys. Since rekeying is typically a costly operation for a block cipher, it would be appealing to have a key-length extension construction providing the same level of security as 3XOR, but calling the underlying block cipher E with a single key. We describe such a construction in this section.

Let π be a linear orthomorphism of \mathbb{F}_2^n (a permutation π of $\{0, 1\}^n$ is an orthomorphism if $z \mapsto z \oplus \pi(z)$ is also a permutation).⁸ We define the 3XSK (*3 XOR, single key*) construction which turns a block cipher $E \in \text{BC}(\kappa, n)$ into a new block cipher $3\text{XSK}[E] \in \text{BC}(\kappa + n, n)$ as follows (see Figure 6):

$$3\text{XSK}_{k,z}[E](x) = E_k(E_k(x \oplus z) \oplus \pi(z)) \oplus z.$$

The induced sequential cipher $\overline{3\text{XSK}}$ is exactly the two-round Even-Mansour cipher with a single permutation and the sequence of round keys $(z, \pi(z), z)$,

$$\overline{3\text{XSK}}_z[P](x) = P(P(x \oplus z) \oplus \pi(z)) \oplus z.$$

⁸ For example, assuming n even, $\pi : (z_L, z_R) \mapsto (z_R, z_L \oplus z_R)$, where z_L and z_R are respectively the left and right halves of z , is an \mathbb{F}_2 -linear orthomorphism.




Fig. 6. The 3XSK[E] key-length extension scheme.

Again, the security of this construction was studied by Chen *et al.* [7] and we restate their findings in the full version of this paper [13]. Combining it with Corollary 1, we obtain the following theorem for the security of the 3XSK construction.

Theorem 5. *Assume that $n \geq 9$, $q_c \geq 9n$, $q_e \geq 9n$, and $4q_c + 2q_e \leq N$. Then the following upper bounds hold:*

(i) *When $q_c \leq 2^{\frac{n}{3}}$, one has*

$$\mathbf{Adv}_{3\text{XSK}}^{\text{cca}}(q_c, q_p) \leq \frac{23}{N^{\frac{1}{3}}} + 32 \left(\frac{q_c q_e}{KN} \right)^{\frac{1}{2}}.$$

(ii) *When $q_c \geq 2^{\frac{n}{3}}$, one has*

$$\mathbf{Adv}_{3\text{XSK}}^{\text{cca}}(q_c, q_p) \leq \frac{10}{N} + (23 + 6\sqrt{n}) \frac{q_c}{N^{\frac{2}{3}}} + 2 \times (39 + 9\sqrt{n}) \left(\frac{q_e}{KN^{\frac{2}{3}}} \right)^{\frac{1}{2}}.$$

(Note that this bound becomes vacuous for $q_c \geq 2^{\frac{2n}{3}}$.)

This matches the security bound for 2XOR for $q_c < 2^{\frac{n}{3}}$, (and hence the lower bound proven for 3XOR in Section 4.3 for $q_c < 2^{\frac{n}{4}}$) while for $2^{\frac{n}{3}} \leq q_c \leq 2^{\frac{2n}{3}}$ it caps at $q_p \approx 2^{\kappa + \frac{2n}{3}}$ (hence it is slightly worse than the security lower bound of 3XOR in that case). The security for q_c larger than $2^{\frac{2n}{3}}$ remains unknown. Note that the attack given in [12] also applies to 3XSK exactly in the same way as to 3XOR, providing an upper bound on its security.

5 Plain Cascade Encryption

In this section, we give another application of Lemma 1, this time to analyze the security of plain cascade encryption in the setting where the number of construction queries is smaller than 2^n . Recall that the ℓ -round cascade encryption using a (κ, n) -block cipher E , denoted $\text{CE}[E]$, takes an $\ell\kappa$ -bit key $\text{mk} = (k_1, \dots, k_\ell) \in (\{0, 1\}^\kappa)^\ell$ and encrypts a plaintext $x \in \{0, 1\}^n$ by computing

$$y = \text{CE}_{\text{mk}}[E](x) = E_{k_\ell} \circ E_{k_{\ell-1}} \circ \dots \circ E_{k_2} \circ E_{k_1}(x).$$

We focus on the security of CE for odd length $\ell = 2r + 1$ and our result is summarized in the following theorem.

Theorem 6. Consider the ℓ -round cascade encryption CE where $\ell = 2r + 1$ for some $r \geq 1$. Then, assuming $q_c \leq N/4$ and $q_e \leq KN/8$, one has

$$\text{Adv}_{\text{CE}}^{\text{cca}}(q_c, q_e) \leq \frac{\ell^2}{K} + e^{-n} + r^2(r+1) \left(\frac{3nq_c}{K} \right)^{\frac{1}{2}} + \max \left\{ A_r \left(\frac{q_c q_e^r}{K^r N^r} \right)^{\frac{1}{2r+1}}, B_r \left(\frac{nq_e}{K^2} \right)^{\frac{1}{3}} \right\},$$

where $A_r \in \mathcal{O}(r^2)$ and $B_r \in \mathcal{O}(r^{\frac{7}{3}})$ only depend on r , namely

$$A_r = 6r^2(r+1) \left(\frac{2^r}{r^{r+1}(r+1)^{r+1}} \right)^{\frac{1}{2r+1}}, \quad B_r = 6r^2(r+1) \left(\frac{3}{4r(r+1)} \right)^{\frac{1}{3}}.$$

Proof (sketch). From a high-level perspective, the proof consists of the following steps. First, we modify the cascade to use two *independent* ideal ciphers E and E' in an interleaving manner and show that this does not introduce a large distinguishing gap. Second, we need to assume that the block cipher E' used in the odd steps of the cascade is *good* in some well-defined sense and hence we show that the opposite is unlikely (over the randomness of E'). Third, we publish the complete function table of E' (but not the keys being used with it), thus arriving at a randomized KLE scheme of length r . Then we can apply Lemma 1 to reduce its security to the security of the induced sequential cipher. Finally, we analyze the latter directly, using an H-coefficient analysis inspired by [8] that employs the assumption that E' is good. The full proof discussing each of the individual steps in greater detail can be found in the full version of this paper [13]. \square

DISCUSSION. In terms of the number of threshold queries, cascade encryption of length $\ell = 2r + 1$ is hence secure when $q_c q_e^r \ll 2^{r(\kappa+n)}$, $q_c \ll 2^\kappa$, and $q_e \ll 2^{2\kappa}$ (asymptotically, ignoring constants). Our bound must be compared with the security result of Dai *et al.* [9], who considered the full-codebook regime $q_c = 2^n$. They showed that, for $\kappa \geq n/(r+1)$ (which is satisfied for virtually any real block cipher we know of), cascade encryption of length $\ell = 2r + 1$ is secure when $q_e \ll 2^{\kappa + \frac{rn}{r+1}}$ (and, obviously, this also holds for any $q_c < 2^n$). Hence, our new bound improves on [9] when $q_c \leq 2^{\frac{rn}{r+1}}$, but only assuming $\kappa \geq \frac{rn}{r+1}$ since otherwise the condition $q_e \ll 2^{2\kappa}$ in our bound becomes more restrictive than Dai *et al.*'s one. This is depicted on Figure 7. We remark that our bound also applies to cascade encryption of length $2r + 2$, since adding a round cannot decrease security.

TIGHTNESS. As observed in [12], the attack against cascades given there can be adjusted to provide a trade-off between block-cipher and construction queries. This results in an attack against plain cascade of length $\ell = 2r + 1$ that achieves a constant distinguishing advantage as long as $q_c q_e^r \approx 2^{r(\kappa+n)}$ and $q_e \geq 2^\kappa q_c$ (again, ignoring constants). Note that the second condition only comes into play




Fig. 7. The security of plain cascade encryption with $2r+1$ or $2r+2$ rounds, depending on κ and n . All parameters below the solid line are secure due either to Theorem 6 or the results of [9]. All parameters above the dashed line are insecure due to the attack of [12]. The status for parameters between these lines remains unknown.

when $q_c \geq 2^{\frac{rn}{r+1}}$, in which case the attack requires $q_e \approx 2^{\kappa + \frac{rn}{r+1}}$ (instead of $q_e \approx 2^{\kappa + n - \frac{1}{r} \log_2 q_c}$). Hence, this matches the bound of [9] for $q_c \geq 2^{\frac{rn}{r+1}}$. When $\kappa \geq n$, this also matches our own new bound for $q_c \leq 2^{\frac{rn}{r+1}}$, yielding a tight bound for all parameters. When $\frac{rn}{r+1} \leq \kappa \leq n$, the attack matches our new bound only for $q_c \geq 2^{r(n-\kappa)}$ since otherwise the security bound caps at $q_e \ll 2^{2\kappa} < 2^{\kappa + n - \frac{1}{r} \log_2 q_c}$. When $\kappa \leq \frac{rn}{r+1}$, there is a provable security gap between this attack and the bound of [9] for any $q_c \leq 2^{\frac{rn}{r+1}}$. This is also summarized on Figure 7. Again, all this applies to the case of cascade encryption of length $2r+2$ since Gaži's attack [12] was given for cascades of even length. Note that the case of 3DES ($\kappa = 56$, $n = 64$, and $r = 1$) corresponds to the middle graph.

TWO-KEY TRIPLE ENCRYPTION. Let TTE denote a variant of triple encryption where the first and the third keys are identical. So TTE accepts a 2κ -bit key $\mathbf{mk} = (k_1, k_2) \in (\{0, 1\}^\kappa)^2$ and encrypts a plaintext $x \in \{0, 1\}^n$ by computing $y = \text{TTE}_{\mathbf{mk}}[E](x) = E_{k_1} \circ E_{k_2} \circ E_{k_1}(x)$. We prove the following result.

Theorem 7. *For the two-key triple encryption TTE, we have, assuming $q_c \leq N/4$ and $q_e \leq KN/8$,*

$$\text{Adv}_{\text{TTE}}^{\text{cca}}(q_c, q_e) \leq e^{-n} + 2 \left(\frac{3nq_c}{K} \right)^{\frac{1}{2}} + 12 \max \left\{ \left(\frac{q_c q_e}{2KN} \right)^{\frac{1}{3}}, \left(\frac{3nq_e}{8K^2} \right)^{\frac{1}{3}} \right\}.$$

Proof (sketch). Similar to the analysis of cascade encryption in the proof of Theorem 6, we slightly modify the key-sampling process from **A** to **B**:

- A:** Choose $\mathbf{mk} \in (\{0, 1\}^\kappa)^2$ uniformly at random.
- B:** Randomly partition $T_1 \cup T_2 = \{0, 1\}^\kappa$ so that $|T_1| = |T_2|$, choose $z_1 \in T_1$ and $k_2 \in T_2$ uniformly at random, and then define $\mathbf{mk} = (z_1, k_2)$.

It is easy to show that these two processes have the same probability distribution. The rest of the proof follows exactly the same line of arguments as the proof of Theorem 6 for cascade encryption of length 3. \square

ACKNOWLEDGMENT. Peter Gaži was partly funded by the European Research Council under an ERC Starting Grant (259668-PSPC).

Jooyoung Lee was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2013R1A1A2007488).

Yannick Seurin was partially supported by the French National Agency of Research through the BLOC project (contract ANR-11-INS-011).

John Steinberger was funded by National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003, and by the China Ministry of Education grant number 20121088050.

Stefano Tessaro was partially supported by NSF grant CNS-1423566.

References

- [1] *EMV Integrated Circuit Card Specification for Payment Systems, Book 2: Security and Key Management, v.4.2*. June 2008.
- [2] Advanced encryption standard (aes). National Institute of Standards and Technology (NIST), FIPS PUB 197, U.S. Department of Commerce, Nov. 2001.
- [3] W. Aiello, M. Bellare, G. Di Crescenzo, and R. Venkatesan. Security amplification by composition: The case of doubly-iterated, ideal ciphers. In *CRYPTO'98*, volume 1462 of *LNCS*, pages 390–407. Springer, Aug. 1998.
- [4] E. Andreeva, A. Bogdanov, Y. Dodis, B. Mennink, and J. P. Steinberger. On the indistinguishability of key-alternating ciphers. In *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 531–550. Springer, Aug. 2013.
- [5] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, May / June 2006.
- [6] 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 *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62. Springer, Apr. 2012.
- [7] S. Chen, R. Lampe, J. Lee, Y. Seurin, and J. P. Steinberger. Minimizing the two-round Even-Mansour cipher. In *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 39–56. Springer, Aug. 2014.
- [8] S. Chen and J. P. Steinberger. Tight security bounds for key-alternating ciphers. In *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, May 2014.
- [9] Y. Dai, J. Lee, B. Mennink, and J. P. Steinberger. The security of multiple encryption in the ideal cipher model. In *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 20–38. Springer, Aug. 2014.
- [10] Data encryption standard. National Bureau of Standards, NBS FIPS PUB 46, U.S. Department of Commerce, Jan. 1977.
- [11] S. Even and Y. Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–162, 1997.
- [12] P. Gaži. Plain versus randomized cascading-based key-length extension for block ciphers. In *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 551–570. Springer, Aug. 2013.
- [13] P. Gaži, J. Lee, Y. Seurin, J. Steinberger, and S. Tessaro. Relaxing Full-Codebook Security: A Refined Analysis of Key-Length Extension Schemes. Full version of this paper. Available at <http://eprint.iacr.org/>.
- [14] P. Gaži and U. M. Maurer. Cascade encryption revisited. In *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 37–51. Springer, Dec. 2009.
- [15] P. Gaži and S. Tessaro. Efficient and optimally secure key-length extension for block ciphers via randomized cascading. In *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 63–80. Springer, Apr. 2012.
- [16] J. Kilian and P. Rogaway. How to protect DES against exhaustive key search (an analysis of DESX). *Journal of Cryptology*, 14(1):17–35, 2001.
- [17] R. Lampe, J. Patarin, and Y. Seurin. An asymptotically tight security analysis of the iterated even-mansour cipher. In *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 278–295. Springer, Dec. 2012.
- [18] R. Lampe and Y. Seurin. How to construct an ideal cipher from a small set of public permutations. In *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 444–463. Springer, Dec. 2013.

- [19] J. Lee. Towards key-length extension with optimal security: Cascade encryption and xor-cascade encryption. In *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 405–425. Springer, May 2013.
- [20] M. Luby and C. Rackoff. Pseudo-random Permutation Generators and Cryptographic Composition. In *Symposium on Theory of Computing - STOC '86*, pages 356–363. ACM, 1986.
- [21] U. M. Maurer, K. Pietrzak, and R. Renner. Indistinguishability amplification. In *CRYPTO 2007*, volume 4622 of *LNCS*, pages 130–149. Springer, Aug. 2007.
- [22] U. M. Maurer and S. Tessaro. Computational indistinguishability amplification: Tight product theorems for system composition. In *CRYPTO 2009*, volume 5677 of *LNCS*, pages 355–373. Springer, Aug. 2009.
- [23] J. Patarin. The “coefficients H” technique (invited talk). In *SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, Aug. 2008.
- [24] J. Steinberger. Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Report 2012/481, 2012. <http://eprint.iacr.org/2012/481>.
- [25] S. Tessaro. Security amplification for the cascade of arbitrarily weak PRPs: Tight bounds via the interactive hardcore lemma. In *TCC 2011*, volume 6597 of *LNCS*, pages 37–54. Springer, Mar. 2011.
- [26] S. Vaudenay. Adaptive-attack norm for decorrelation and super-pseudorandomness. In *SAC 1999*, volume 1758 of *LNCS*, pages 49–61. Springer, Aug. 1999.

A Overview of Previous Results

The following tables summarize the currently known best bounds for the full-codebook regime where $q_c = 2^n$. For a given ℓ (resp. r), ℓ' (resp. r') is the smallest even integer greater or equal to ℓ (resp. r).

ℓ -Round Plain Cascades

Ref.	Max. value of $\log(q_e)$	Note
[5]	$\kappa + \min\{n/2, \kappa/2\}$	$\ell = 3$
[14]	$\kappa + \min\{\kappa(\ell' - 2)/\ell', n/2\}$	
[19]	$\kappa + \min\{\kappa, n\} - 8n/\ell$	
[9]	$\kappa + \min\{\kappa(\ell' - 2)/2, n(\ell' - 2)/\ell'\}$	

r -Round Randomized Cascades

Ref.	Max. value of $\log(q_e)$	Note
[16]	$n + \kappa - \log(q_c)$	$r = 1$, FX $r = 2$, 2XOR
[15]	$\kappa + n/2$	
[19]	$\kappa + \min\{\kappa, n\} - 4n/r$	
[12]	$\kappa + \frac{r'-2}{r'}n$	