

Secret-key Cryptography from Ideal Primitives: A Systematic Overview

Peter Gazi

Institute of Science and Technology, Austria
peter.gazi@ist.ac.at

Stefano Tessaro

University of California, Santa Barbara
tessaro@cs.ucsb.edu

Abstract—Secret-key constructions are often proved secure in a model where one or more underlying components are replaced by an idealized oracle accessible to the attacker. This model gives rise to information-theoretic security analyses, and several advances have been made in this area over the last few years. This paper provides a systematic overview of what is achievable in this model, and how existing works fit into this view.

Index Terms—Cryptography, provable security, ideal-primitive model.

I. INTRODUCTION

Cryptographic systems are often proved secure under *computational* assumptions, such as the hardness of a certain computational problem at hand. This approach is by now standard, yet many practical schemes cannot be proved secure under such assumptions despite being seemingly sound.

An alternative is to carry out the security analysis in the so-called *ideal primitive model* (IPM): One starts by identifying a building block \mathbf{B} of a cryptographic construction $C = C[\mathbf{B}]$, and modeling an “idealized” version \mathbf{I} of \mathbf{B} . Then, one proves that the construction C is secure when \mathbf{B} is replaced by \mathbf{I} , and the attacker can evaluate \mathbf{I} . Typically, these proofs are completely information-theoretic, i.e., security holds against attackers only bounded in the number of calls to \mathbf{I} but otherwise computationally unrestricted.

An IPM security proof is a *heuristic* argument towards the security of the real construction $C[\mathbf{B}]$: It excludes any attack that would use \mathbf{B} in a black-box way, and hence, any successful attack would have to exploit some structural weakness of the real primitive \mathbf{B} . Generally, IPM proofs do not imply that \mathbf{I} can be replaced by some concrete algorithm \mathbf{B} and retain security. Still, they are often the only way to validate the security of many practical constructions.

This paper considers constructions of secret-key primitives in the IPM (and more concretely, pseudorandom functions and permutations, defined below), a problem that has gained widespread attention over the last few years. First, we are going to provide a unified treatment of the problem, highlighting the inherent barriers in terms of achievability. Second, we will survey some existing results and see how they fit within this general treatment. We also provide two results (Theorem 1 and Theorem 3) that are new to the best of our knowledge, and help better define the landscape.

The motivation for these results has been mostly practical, as these results often validate block-cipher designs. Here,

however, we focus on the development of a general theory of cryptographic constructions in ideal models, rather than on the practical implications of these results.

II. IDEAL PRIMITIVES AND SECRET-KEY CONSTRUCTIONS

A. Ideal primitives

We start by formalizing the concept of an ideal primitive using notation inspired from [27]. An *ideal primitive* \mathbf{I} is a system that initially samples as its state a function I according to some probability distribution, and then gives access to it in forms of queries, i.e., on input x it returns $I(x)$.¹ The following are typical examples:

- A *random function* $\mathbf{R}^{n,m}$ gives access to a function $R : \{0, 1\}^n \rightarrow \{0, 1\}^m$ chosen uniformly at random. It is also useful to consider a *random oracle* $\mathbf{R}^{*,m}$, which takes *arbitrary-length* inputs, instead of n -bit long ones.
- A *random permutation* \mathbf{P}^n replies to both *forward queries* $(x, +)$, and *backward queries* $(y, -)$, returning $\pi(x)$ and $\pi^{-1}(y)$, respectively, for a permutation π on n -bit strings chosen uniformly at random.
- An extension is an *ideal cipher* $\mathbf{E}^{\kappa,n}$ which replies to both *forward queries* $(k, x, +)$, and *backward queries* $(k, y, -)$, where $k \in \{0, 1\}^\kappa$, returning $\pi_k(x)$ and $\pi_k^{-1}(y)$, respectively, where the 2^κ permutations $\{\pi_k\}_{k \in \{0, 1\}^\kappa}$ on the n -bit strings are chosen independently uniformly at random. We also write $\mathbf{E}_k^{\kappa,n}$ to identify the random permutation obtained by restricting queries to π_k .

We often omit the parameters given as superscripts when clear from the context. Also, we are going to write $\mathbf{I}(x)$ to denote invoking the random primitive \mathbf{I} on input x . Often, we write $\mathbf{P}(x)$ or $\mathbf{P}^{-1}(y)$ instead of $\mathbf{P}(x, +)$ and $\mathbf{P}(y, -)$.

B. Pseudorandom functions and permutations

In this paper, we focus on the problem building pseudorandom functions (PRFs) and permutations (PRPs) from ideal primitives. PRFs and PRPs are *universal* for symmetric cryptography, meaning that all conventional secret-key primitives can be built in a black-box way from them. Concretely, we consider *deterministic* constructions $C = C[\mathbf{I}]$ that invoke

¹Note that efficient implementations of \mathbf{I} usually use *lazy sampling*, i.e., they generate I on the fly query after query, keeping this partial state as a table.

an ideal primitive \mathbf{I} to implement a *keyed* function $C[\mathbf{I}] : \{0, 1\}^w \times \{0, 1\}^d \rightarrow \{0, 1\}^r$, where $d \in \mathbb{N} \cup \{*\}$. The first argument is referred to as the *key*, and the second as the *input*.

PRF security demands that $C_K[\mathbf{I}] = C[\mathbf{I}](K, \cdot)$ (for a random secret K) behaves as $\mathbf{R}^{d,r}$ for bounded adversaries. Concretely, let A be an adversary, i.e., a computationally unbounded machine that can query one or more oracles, and finally outputs a decision bit. We define the *PRF advantage of A against C* as the quantity

$$\text{Adv}_C^{\text{prf}}(A) = \mathbb{P} \left[A^{C_K[\mathbf{I}], \mathbf{I}} \Rightarrow 1 \right] - \mathbb{P} \left[A^{\mathbf{R}^{d,r}, \mathbf{I}} \Rightarrow 1 \right],$$

where K is a uniform w -bit string. Note that here A gets access to two oracles, namely $C_K[\mathbf{I}]$ and \mathbf{I} in the “real world”, as well as $\mathbf{R}^{d,r}$ and \mathbf{I} in the “ideal world.” In particular, in the real world, both oracles are correlated, whereas in the ideal world they are independent. Queries to the first oracle are called *construction queries*, whereas queries to the second oracle are called *primitive queries*.

We also say that $E : \{0, 1\}^w \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is (κ, n) -*block cipher* if $E_k = E(k, \cdot)$ is a permutation for all k . If $C[\mathbf{I}]$ is a (κ, n) -block cipher, then we define the PRP advantage of an adversary A against C as the quantity

$$\text{Adv}_C^{\text{prp}}(A) = \mathbb{P} \left[A^{C_K[\mathbf{I}], \mathbf{I}} \Rightarrow 1 \right] - \mathbb{P} \left[A^{\mathbf{P}^n, \mathbf{I}} \Rightarrow 1 \right],$$

where we abuse notation by denoting as $C_K[\mathbf{I}]$ to oracle that answers forward queries $(x, +)$ via $C_K[\mathbf{I}](x)$ and backward queries $(y, -)$ as $C_K[\mathbf{I}]^{-1}(y)$.

Finally, we say that C is a (q_C, q_P, ε) -PRF if $\text{Adv}_C^{\text{prf}}(A) \leq \varepsilon$ for all adversaries A making q_C construction and q_P primitive queries. We informally say that C is a (q_C, q_P) -PRF if it is a (q_C, q_P, ε) -PRF for some small ε (e.g., $\varepsilon = o(1)$ or $\varepsilon = 2^{-w}$). Similarly, one can define the notions of a (q_C, q_P, ε) -PRP and a (q_C, q_P) -PRP.

III. GENERIC ATTACKS

Here, we want to understand what are the largest q_C and q_P such that a construction can be a (q_C, q_P) -secure PRF or PRP. We give two generic attacks establishing such bounds.

RANDOMNESS EXHAUSTION ATTACK: For a primitive $\mathbf{I} \in \{\mathbf{R}^{n,m}, \mathbf{P}^n, \mathbf{E}^{\kappa,n}\}$, let $q(\mathbf{I})$ be the number of queries necessary to recover its internal randomness I . (I.e., $q(\mathbf{R}^{n,m}) = q(\mathbf{P}^n) = 2^n$, $q(\mathbf{E}^{\kappa,n}) = 2^{\kappa+n}$.) Moreover, for $q < q(\mathbf{I})$, let $I_q(\mathbf{I})$ be the minimum (over all sequences of q queries to \mathbf{I}) of the number of possible states of the internal randomness of \mathbf{I} consistent with these q queries. Also, let $R_q(\mathbf{I}) = \log I_q(\mathbf{I})$. For example, $R_q(\mathbf{R}^{n,m}) = (2^n - q) \cdot m$.

Theorem 1: Let $C = C[\mathbf{I}] : \{0, 1\}^w \times \{0, 1\}^d \rightarrow \{0, 1\}^r$. For all $q_P < q(\mathbf{I})$, there exists an adversary A asking q_P primitive queries and $q_C = \left\lceil \frac{R_{q_P}(\mathbf{I}) + w + \Delta}{r} \right\rceil$ construction queries, such that

$$\text{Adv}_C^{\text{prf}}(A) \geq 1 - 2^{-\Delta}.$$

Proof: The adversary A asks $q = q_P$ primitive queries minimizing the number of consistent states I to be $I_q(\mathbf{I}) = 2^{R_q(\mathbf{I})}$, and q_C distinct construction queries. Then, given the




Fig. 1. The construction M_3 .

query-answer pairs x_i, y_i for $i \in [q_C]$, the adversary A checks whether there exists some I consistent with the q_P primitive queries, and a key $k \in \{0, 1\}^w$ such that $C_k[I](x_i) = y_i$ for all $i \in [q_C]$. If so, it outputs 1, if not it outputs 0. Clearly, A always outputs 1 in the real world. In the ideal world, the adversary obtains at least $R_q(\mathbf{I}) + w + \Delta$ bits of randomness, yet there are at most $2^{R_q(\mathbf{I}) + w}$ choices of I and k , and thus at most so many sequences of q_C pairs (x_i, y_i) that can appear in the real world. Thus, the distinguisher outputs 1 with probability at most $2^{-\Delta}$, i.e., the probability that the random outputs hit one of these sequences. ■

The above attack implies that any PRF construction based on \mathbf{I} can only be secure against adversaries with $q_P \ll q(\mathbf{I})$ and $q_C \ll q(\mathbf{I})$. The same attack can also be used against PRP security, with the term $2^{-\Delta}$ replaced by a slightly larger one.

KEY-SEARCH ATTACK: Another simple attack tries out all keys and checks them for consistency with construction queries. This is summarized by the following theorem, whose proof is omitted. (The same attack applies to PRP security, with only a slightly larger term instead of 2^{w-rq_C} .)

Theorem 2: Let $C = C[\mathbf{I}] : \{0, 1\}^w \times \{0, 1\}^d \rightarrow \{0, 1\}^r$, and let t be an upper bound on the number of queries C makes to \mathbf{I} upon each invocation. Then there exists an adversary A asking q_C construction queries and $q_P \leq tq_C 2^w$ primitive queries such that $\text{Adv}_C^{\text{prf}}(A) \geq 1 - 2^{w-rq_C}$.

IV. CONSTRUCTIONS FROM RANDOM PERMUTATIONS AND FUNCTIONS

In this section, we discuss the question of building a PRF or a PRP from a random function $\mathbf{R}^{n,m}$ or a random permutation \mathbf{P}^n . We note that the best we can hope for here is to build a $(2^n, 2^n)$ -PRF or a $(2^n, 2^n)$ -PRP, and we will see how close we get to this. We note that such constructions have been considered in the context of justifying block-cipher designs from some (unkeyed) function or permutation.

A. Constructing PRFs

As far as we know, all PRF constructions (without invertibility properties) from a random function $\mathbf{R} = \mathbf{R}^{n,m}$ in the literature fall short of achieving $(2^n, 2^n)$ -PRF security. Here, we show that the following construction $M_\ell[\mathbf{R}]$ (cf. also




Fig. 2. The key-alternating cipher $\text{KAC}[\mathbf{P}_1, \dots, \mathbf{P}_\ell]$ of length ℓ .




Fig. 3. The cascade construction $\text{CE}[\mathbf{E}]$ of length ℓ .

Fig. 1) is sufficient to achieve this: Given key $k = (k_1, \dots, k_\ell)$, on input $x \in \{0, 1\}^n$, it outputs

$$\mathbf{M}_\ell[\mathbf{R}](k, x) = \mathbf{R}(x \oplus k_1) \oplus \dots \oplus \mathbf{R}(x \oplus k_\ell).$$

This construction is similar to one used by Myers [29] for the different problem of amplifying security of weak pseudorandom functions. We are not aware that it was ever analyzed for $\ell \geq 2$ (the analysis for $\ell = 1$ is folklore). The following theorem establishes its security, and implies for example that for $\ell = \Omega(1/\varepsilon)$ the construction is a (q_C, q_P) -PRF for $q_C = q_P = 2^{n(1-\varepsilon)}$ for any $\varepsilon > 0$.

Theorem 3: The construction \mathbf{M}_ℓ is a (q_C, q_P, ε) -PRF for $\varepsilon = q_C \cdot (q_C + q_P)^\ell / 2^{n\ell}$.

Here, we only give the main idea behind the proof. Using a standard argument, one can show that to break PRP security an attacker needs to issue a construction query x with the property that there exists no $i \in [\ell]$ such that $x \oplus k_i$ is fresh, i.e., for every i , \mathbf{R} is evaluated on input $x \oplus k_i$ by either another construction query or directly as a primitive query, and that the probability that this event occurs bounds $\text{Adv}_{\mathbf{M}_\ell}^{\text{prf}}(\mathbf{A})$ for every \mathbf{A} . Again by a standard argument (cf. e.g. [27]), it is sufficient to consider the probability of this event in the *ideal* world, where the keys k_1, \dots, k_ℓ are independent of the interaction. Using the independence of the keys, it is not hard to see that the probability of this event is at most $q_C(q_C + q_P)^\ell / 2^{n\ell}$.

B. PRPs from random permutations

The question of building PRPs from a random permutation \mathbf{P}^n was first considered by Even and Mansour [13]. Their construction, given keys (k_1, k_2) and input x , outputs

$$\text{EM}[\mathbf{P}](k_1, k_2, x) = k_2 \oplus \mathbf{P}(x \oplus k_1).$$

The construction is a (q_C, q_P) -PRP as long as $q_C \cdot q_P \leq 2^n$, even when $k_1 = k_2$ [12]. This construction can be generalized to multiple rounds, and is usually called a *key-alternating cipher (KAC)* (see Fig. 2). This structure is in fact used also within the AES block cipher, and this fact has motivated a recent line of works [5], [30], [22], [7] showing that for ℓ rounds, KACs are a (q_C, q_P) -PRP as long as $q_C q_P^\ell \ll 2^{n\ell}$. This for example tolerates $q_C = 2^n$ and $q_P = 2^{n \frac{\ell-1}{\ell}}$, which is nearly optimal for growing ℓ . Several variants of length-two KACs with repeated keys or permutations were considered in [6]. The results on KACs also have implications for the security of related constructions called XOR-cascades, as discussed below in Section V-C.

C. PRPs from Random Functions

The question of building PRPs from a random function was considered in the works of Gentry and Ramzan [19] and of Lampe and Seurin [24], who gave constructions based on Feistel networks of n -bit input PRPs from a random function $\mathbf{R}^{n/2, n/2}$. In particular, the construction of [24] is (nearly) an $(2^{n/2}, 2^{n/2})$ -PRP for sufficiently many rounds, which is the best one can hope for by Theorem 1.

V. KEY-LENGTH EXTENSION FOR BLOCK CIPHERS

A. Motivation

The key length inherently bounds the PRP security of a block cipher. For some existing block ciphers (such as DES) a short key represents *the* main security shortcoming, even if they otherwise do not exhibit any weaknesses. This motivated the problem of *key-length extension (KLE)*: The goal is to find constructions \mathbf{C} transforming any (κ, n) -block cipher E into a (w, n) -block cipher $\mathbf{C}[E]$ with $w > \kappa$ and with the property that any attack should require significantly more than 2^κ efforts, assuming that E itself contains no non-generic weaknesses, and hence can be modeled as $\mathbf{E} = \mathbf{E}^{\kappa, n}$. Therefore, KLE formally considers constructions $\mathbf{C}[\mathbf{E}]$ of a (w, n) -block cipher for $w > \kappa$ which we want to prove to be secure PRPs, noting that the best we can hope for is a (q_C, q_P) -PRP for q_C and q_P approaching 2^n and $2^{\kappa+n}$, respectively.

B. Plain Cascades

The most natural KLE approach is to apply the block cipher repeatedly using an independent key at each step – this is usually referred to as *cascading*. Formally, the cascade of length ℓ for $\mathbf{E} = \mathbf{E}^{\kappa, n}$ is the $(\ell \cdot \kappa, n)$ -block cipher which, on input key $k = (k_1, \dots, k_\ell)$ and plaintext x , returns

$$\text{CE}_\ell[\mathbf{E}](k, x) = (\mathbf{E}_{k_\ell} \circ \dots \circ \mathbf{E}_{k_1})(x). \quad (1)$$

A cascade of length ℓ is depicted in Figure 3. In practice, the cascade of length $\ell = 3$ underlies the widely deployed *Triple-DES (3DES)* standard [1].

We note that plain cascades have been the object of security analyses in the standard-model, showing that they can be used to amplify the security of weak pseudorandom permutations. Describing these works is beyond the scope of this paper, but we refer the reader to [31] for an overview.

SECURITY BOUNDS: A length-two cascade does not increase security in terms of the number of tolerable primitive queries due to the meet-in-the-middle attack [11]. Nonetheless, a slight security increase in terms of smaller distinguishing advantage ε was shown in [2] when $q_P < 2^\kappa$.

Bellare and Rogaway [4] were the first to prove that CE_3 is a (q_C, q_P) -PRP whenever $q_C \leq 2^n$ and

$$\log(q_P) \ll \kappa + \min \left\{ \frac{\kappa}{2}, \frac{n}{2} \right\}.$$

Gaži and Maurer [16] subsequently showed an improvement for odd $\ell \geq 4$ and $\kappa \leq n$, showing CE_ℓ is a (q_C, q_P) -PRP when $q_C \leq 2^n$ and

$$\log(q_P) \ll \min \left\{ \frac{2\ell}{\ell+1} \cdot \kappa, \kappa + \frac{n}{2} \right\}.$$

With increasing ℓ the right-hand side approaches $\min \{2\kappa, \kappa + \frac{n}{2}\}$. Lee [25] improved the right-hand side, showing a better bound approaching the optimal value of $\kappa + \min \{\kappa, n\}$ with increasing $\ell \rightarrow \infty$, however his result only gives useful bounds for large ℓ (say $\ell \geq 16$). A tight bound (matching the attacks mentioned below) was finally given by Dai, Lee, Mennink, and Steinberger [10], establishing the security of a cascade of length ℓ as long as $q_C \leq 2^n$ and

$$\log(q_P) \ll \kappa + \min \left\{ \frac{\ell'-2}{2} \cdot \kappa, \frac{\ell'-2}{\ell'} \cdot n \right\},$$

where $\ell' = 2 \cdot \lceil \ell/2 \rceil$ denotes the smallest even integer not smaller than ℓ .

A recent work [15] initiated the study of plain cascades in a more fine-grained (and cryptographically more appropriate) setting where q_C can be smaller than 2^n . In this setting, [15] prove that plain cascades of length $\ell = 2r + 1$ are secure whenever

$$\begin{aligned} \log(q_C) + r \log(q_P) &\ll r(\kappa + n) \quad \wedge \quad \log(q_C) \ll \kappa \\ &\quad \wedge \quad \log(q_P) \ll 2\kappa. \end{aligned}$$

They also show a very similar bound for the two-key variant in the $\ell = 3$ case, where the first and the third encryption keys are identical (as proposed in the 3DES standard).

GENERIC ATTACKS: A parallel (and overlapping) line of works investigates so-called *generic* attacks against cascades. These are attacks that compromise its PRP-security (as described above) in the ideal-cipher model, i.e., by accessing the block-cipher as a black box and hence *not* exploiting any potential weaknesses of its inner workings.

Lucks [26] presented the to-date best known attack on cascades of length $\ell = 3$, which requires roughly $2^{\kappa+n/2}$ queries. A generalization of this attack [18] is applicable against cascades of length $\ell \in \{2r + 1, 2r + 2\}$, requiring q_C construction queries and q_P primitive queries, where q_C, q_P satisfy

$$\log(q_C) + r \log(q_P) \approx r(\kappa + n) \quad \wedge \quad q_C \leq q_P/2^\kappa.$$

C. XOR-Cascades

An alternative KLE approach uses variant of the key-whitening technique, generalizing the DESX block-cipher construction due to Rivest. In its simplest form, one obtains the following construction $\text{FX}[\mathbf{E}]$ abstracting DESX which given three keys k_i, k_o, k , on input x outputs

$$\text{FX}[\mathbf{E}]((k_i, k_o, k), x) = k_o \oplus \mathbf{E}_k(k_i \oplus x).$$




Fig. 4. The XOR-cascade construction XCE of length ℓ .

Combining cascading and key whitening, Gaži and Tesaro [17] proposed the so-called 2-XOR-cascade (or randomized cascade) construction. It maps each n -bit message x under a key $(k, z) \in \{0, 1\}^\kappa \times \{0, 1\}^n$ to

$$\text{2XOR}[\mathbf{E}]((k, z), x) = \mathbf{E}_{\tilde{k}}(\mathbf{E}_k(x \oplus z) \oplus z)$$

where \tilde{k} is derived from k in a deterministic way (e.g. by flipping a single bit). The natural generalization to length ℓ in [25], [18] is referred to as the ℓ -XOR-cascade XCE_ℓ . Given key (k, z) (where $z = (z_0, z_1, \dots, z_\ell)$) and input x , it returns

$$\begin{aligned} \text{XCE}_\ell[\mathbf{E}]((k, z), x) &= (\oplus_{z_\ell} \circ \mathbf{E}_{\phi_\ell(k)} \circ \oplus_{z_{\ell-1}} \circ \mathbf{E}_{\phi_{\ell-1}(k)} \\ &\quad \circ \dots \circ \oplus_{z_1} \circ \mathbf{E}_{\phi_1(k)} \circ \oplus_{z_0})(x), \end{aligned}$$

where \oplus_z maps x' to $x' \oplus z$, and ϕ_1, \dots, ϕ_ℓ are permutations on the κ -bit strings such that $\phi_i(k) \neq \phi_j(k)$ for all k and $i \neq j$. The general XOR-cascade is depicted in Figure 4.

SECURITY BOUNDS: Kilian and Rogaway [21] showed that FX is a (q_C, q_P) -PRP whenever $q_C \cdot q_P \ll 2^{\kappa+n}$ and also provided an attack matching their bound. In contrast, the 2-XOR-cascade construction [17] was proved secure as long as $q_C \leq 2^n$ and $q_P \leq 2^{\kappa+n/2}$ queries and this bound is also shown tight by a matching attack.

The work by Lee [25] first considered the general case of XOR-cascade of length ℓ (with independent keys) and proved that its security approaches the optimal bound $2^{\kappa+n}$, while again giving useful statements only for large ℓ .

The security of a variant of XOR-cascades (not containing the last whitening step) was considered in [18], where it was reduced to the security of key-alternating ciphers discussed in Section IV-B. This reduction, together with a tight analysis of KACs given in [7], results in tight bounds for XOR-cascades in the setting with 2^n construction queries. A recent, more fine-grained reduction given in [15] achieves the same for the general case of arbitrary $q_C \leq 2^n$, establishing that an XOR-cascade of length ℓ is secure, roughly speaking, as long as $\log(q_C) + \ell \log(q_P) \ll \ell(\kappa + n)$.

GENERIC ATTACKS: A generic attack against XOR-cascades was given in [18], requiring roughly q_C construction queries and q_P block cipher queries for any values q_C, q_P such that $\log(q_C) + \ell \log(q_P) \approx \ell(\kappa + n)$ is satisfied. This attack hence proves the above-mentioned bound to be tight.

VI. ON STRONGER SECURITY NOTIONS

There have been works targeting strictly stronger security notions than PRF and PRP security. In particular, a series of works considered constructions of block ciphers from random functions [9], [20] and from random permutations [3], [23] in

the sense of indifferentiability [28]. Very recently, the notion of PRF and PRP security against related-key attacks has also been shown to be attainable in [14], [8]. However, the concrete security of the constructions is far lower than that for all aforementioned results.

Acknowledgements. We would like to thank all the colleagues we worked with on these questions: Yooyoung Lee, Ueli Maurer, Yannick Seurin and John Steinberger.

Peter Gaži was partly funded by the European Research Council under an ERC Starting Grant (259668-PSPC). Stefano Tessaro was partially supported by NSF grant CNS-1423566.

REFERENCES

- [1] “Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher.” National Institute of Standards and Technology, Special Publication 800-67, 2004.
- [2] W. Aiello, M. Bellare, G. D. Crescenzo, and R. Venkatesan, “Security amplification by composition: The case of doubly-iterated, ideal ciphers,” in *Advances in Cryptology — CRYPTO ’98*, vol. 1462 of *Lecture Notes in Computer Science*, pp. 390–407, Springer Berlin Heidelberg, 1998.
- [3] E. Andreeva, A. Bogdanov, Y. Dodis, B. Mennink, and J. P. Steinberger, “On the indifferentiability of key-alternating ciphers,” in *CRYPTO 2013, Part I* (R. Canetti and J. A. Garay, eds.), vol. 8042 of *LNCS*, (Santa Barbara, CA, USA), pp. 531–550, Springer, Berlin, Germany, Aug. 18–22, 2013.
- [4] M. Bellare and P. Rogaway, “Code-based game-playing proofs and the security of triple encryption,” in *Advances in Cryptology — EUROCRYPT 2006*, vol. 4004 of *Lecture Notes in Computer Science*, pp. 409–426, Springer Berlin Heidelberg, 2006. Full version at <http://eprint.iacr.org/2004/331>.
- [5] A. Bogdanov, L. R. Knudsen, G. Leander, F.-X. Standaert, J. Steinberger, and E. Tischhauser, “Key-alternating ciphers in a provable setting: encryption using a small number of public permutations,” in *Advances in Cryptology — EUROCRYPT 2012* (D. Pointcheval and T. Johansson, eds.), vol. 7237 of *Lecture Notes in Computer Science*, pp. 45–62, Springer Berlin Heidelberg, 2012.
- [6] S. Chen, R. Lampe, J. Lee, Y. Seurin, and J. P. Steinberger, “Minimizing the two-round Even-Mansour cipher,” in *CRYPTO 2014, Part I* (J. A. Garay and R. Gennaro, eds.), vol. 8616 of *LNCS*, (Santa Barbara, CA, USA), pp. 39–56, Springer, Berlin, Germany, Aug. 17–21, 2014.
- [7] S. Chen and J. P. Steinberger, “Tight security bounds for key-alternating ciphers,” in *EUROCRYPT 2014* (P. Q. Nguyen and E. Oswald, eds.), vol. 8441 of *LNCS*, (Copenhagen, Denmark), pp. 327–350, Springer, Berlin, Germany, May 11–15, 2014.
- [8] B. Cogliati and Y. Seurin, “On the provable security of the iterated even-mansour cipher against related-key and chosen-key attacks,” in *EUROCRYPT 2015*, LNCS, 2015. To appear.
- [9] J.-S. Coron, J. Patarin, and Y. Seurin, “The random oracle model and the ideal cipher model are equivalent,” in *CRYPTO 2008* (D. Wagner, ed.), vol. 5157 of *LNCS*, (Santa Barbara, CA, USA), pp. 1–20, Springer, Berlin, Germany, Aug. 17–21, 2008.
- [10] 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* (J. A. Garay and R. Gennaro, eds.), vol. 8616 of *LNCS*, (Santa Barbara, CA, USA), pp. 20–38, Springer, Berlin, Germany, Aug. 17–21, 2014.
- [11] W. Diffie and M. E. Hellman, “Exhaustive Cryptanalysis of the NBS Data Encryption Standard,” *Computer*, vol. 10, no. 6, pp. 74–84, 1977.
- [12] O. Dunkelman, N. Keller, and A. Shamir, “Minimalism in cryptography: The Even-Mansour scheme revisited,” in *EUROCRYPT 2012* (D. Pointcheval and T. Johansson, eds.), vol. 7237 of *LNCS*, (Cambridge, UK), pp. 336–354, Springer, Berlin, Germany, Apr. 15–19, 2012.
- [13] S. Even and Y. Mansour, “A construction of a cipher from a single pseudorandom permutation,” *Journal of Cryptology*, vol. 10, no. 3, pp. 151–162, 1997.
- [14] P. Farshim and G. Procter, “The related-key security of iterated even-mansour ciphers,” in *FSE 2015*, LNCS, 2015. To appear.
- [15] P. Gaži, J. Lee, Y. Seurin, J. Steinberger, and S. Tessaro, “Relaxing Full-Codebook Security: A Refined Analysis of Key-Length Extension Schemes.” to appear in *Fast Software Encryption*, 2015.
- [16] P. Gaži and U. Maurer, “Cascade encryption revisited,” in *Advances in Cryptology — ASIACRYPT 2009* (M. Matsui, ed.), vol. 5912 of *Lecture Notes in Computer Science*, pp. 37–51, Springer Berlin Heidelberg, 2009.
- [17] P. Gaži and S. Tessaro, “Efficient and optimally secure key-length extension for block ciphers via randomized cascading,” in *Advances in Cryptology — EUROCRYPT 2012* (D. Pointcheval and T. Johansson, eds.), vol. 7237 of *Lecture Notes in Computer Science*, pp. 63–80, Springer Berlin Heidelberg, 2012.
- [18] P. Gaži, “Plain versus randomized cascading-based key-length extension for block ciphers,” in *CRYPTO 2013, Part I* (R. Canetti and J. A. Garay, eds.), vol. 8042 of *LNCS*, (Santa Barbara, CA, USA), pp. 551–570, Springer, Berlin, Germany, Aug. 18–22, 2013.
- [19] C. Gentry and Z. Ramzan, “Eliminating random permutation oracles in the Even-Mansour cipher,” in *ASIACRYPT 2004* (P. J. Lee, ed.), vol. 3329 of *LNCS*, (Jeju Island, Korea), pp. 32–47, Springer, Berlin, Germany, Dec. 5–9, 2004.
- [20] T. Holenstein, R. Künzler, and S. Tessaro, “The equivalence of the random oracle model and the ideal cipher model, revisited,” in *43rd ACM STOC* (L. Fortnow and S. P. Vadhan, eds.), (San Jose, California, USA), pp. 89–98, ACM Press, June 6–8, 2011.
- [21] J. Kilian and P. Rogaway, “How to Protect DES Against Exhaustive Key Search (an Analysis of DESX),” *Journal of Cryptology*, vol. 14, pp. 17–35, 2001.
- [22] R. Lampe, J. Patarin, and Y. Seurin, “An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher,” in *Advances in Cryptology — ASIACRYPT 2012* (X. Wang and K. Sako, eds.), vol. 7658 of *Lecture Notes in Computer Science*, pp. 278–295, Springer Berlin Heidelberg, 2012.
- [23] R. Lampe and Y. Seurin, “How to construct an ideal cipher from a small set of public permutations,” in *ASIACRYPT 2013, Part I* (K. Sako and P. Sarkar, eds.), vol. 8269 of *LNCS*, (Bengalore, India), pp. 444–463, Springer, Berlin, Germany, Dec. 1–5, 2013.
- [24] R. Lampe and Y. Seurin, “Security analysis of key-alternating feistel ciphers,” in *FSE 2014*, *Lecture Notes in Computer Science*, 2015.
- [25] J. Lee, “Towards Key-Length Extension with Optimal Security: Cascade Encryption and Xor-cascade Encryption,” in *Advances in Cryptology — EUROCRYPT 2013* (T. Johansson and P. Nguyen, eds.), vol. 7881 of *Lecture Notes in Computer Science*, pp. 405–425, Springer Berlin Heidelberg, 2013.
- [26] S. Lucks, “Attacking triple encryption,” in *Fast Software Encryption* (S. Vaudenay, ed.), vol. 1372 of *Lecture Notes in Computer Science*, pp. 239–253, Springer Berlin Heidelberg, 1998.
- [27] U. M. Maurer, “Indistinguishability of random systems,” in *EUROCRYPT 2002* (L. R. Knudsen, ed.), vol. 2332 of *LNCS*, (Amsterdam, The Netherlands), pp. 110–132, Springer, Berlin, Germany, Apr. 28 – May 2, 2002.
- [28] U. M. Maurer, R. Renner, and C. Holenstein, “Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology,” in *TCC 2004* (M. Naor, ed.), vol. 2951 of *LNCS*, (Cambridge, MA, USA), pp. 21–39, Springer, Berlin, Germany, Feb. 19–21, 2004.
- [29] S. Myers, “Efficient amplification of the security of weak pseudo-random function generators,” in *EUROCRYPT 2001* (B. Pfitzmann, ed.), vol. 2045 of *LNCS*, (Innsbruck, Austria), pp. 358–372, Springer, Berlin, Germany, May 6–10, 2001.
- [30] J. Steinberger, “Improved Security Bounds for Key-Alternating Ciphers via Hellinger Distance.” *Cryptology ePrint Archive*, Report 2012/481, 2012. <http://eprint.iacr.org/>.
- [31] S. Tessaro, “Security amplification for the cascade of arbitrarily weak PRPs: Tight bounds via the interactive hardcore lemma,” in *TCC 2011* (Y. Ishai, ed.), vol. 6597 of *LNCS*, (Providence, RI, USA), pp. 37–54, Springer, Berlin, Germany, Mar. 28–30, 2011.