On the Provable Security of the Iterated Even-Mansour Cipher against Related-Key and Chosen-Key Attacks*

Benoît Cogliati** and Yannick Seurin***

April 20, 2015

Abstract. The iterated Even-Mansour cipher is a construction of a block cipher from r public permutations $P_1, \ldots, P_r$ which abstracts in a generic way the structure of key-alternating ciphers. The indistinguishability of this construction from a truly random permutation by an adversary with oracle access to the inner permutations $P_1, \ldots, P_r$ has been investigated in a series of recent papers. This construction has also been shown to be (fully) indifferentiable from an ideal cipher for a sufficient number of rounds (five or twelve depending on the assumptions on the key-schedule). In this paper, we extend this line of work by considering the resistance of the iterated Even-Mansour cipher to xor-induced related-key attacks (i.e., related-key attacks where the adversary is allowed to xor any constant of its choice to the secret key) and to chosen-key attacks. For xor-induced related-key attacks, we first provide a distinguishing attack for two rounds, assuming the key-schedule is linear. We then prove that for a linear key-schedule, three rounds yield a cipher which is secure against xor-induced related-key attacks up to $\mathcal{O}(2^{\frac{n}{2}})$ queries of the adversary, whereas for a nonlinear key-schedule, one round is sufficient to obtain a similar security bound. We also show that the iterated Even-Mansour cipher with four rounds offers some form of provable resistance to chosen-key attacks, which is the minimal number of rounds to achieve this property. The main technical tool that we use to prove this result is *sequential indifferentiability*, a weakened variant of (full) indifferentiability introduced by Mandal *et al.* (TCC 2010).

Keywords: block cipher, ideal cipher, related-key attacks, chosen-key attacks, iterated Even-Mansour cipher, key-alternating cipher, indifferentiability, correlation intractability

** University of Versailles, France. E-mail: benoitcogliati@hotmail.fr

*** ANSSI, Paris, France. E-mail: yannick.seurin@m4x.org. This author was partially supported by the French National Agency of Research through the BLOC project (contract ANR-11-INS-011).

1 Introduction

Background. The Even-Mansour construction, and its generalization, the iterated Even-Mansour (*IEM* for short) construction, is a very simple way to define a block cipher from a set of r public permutations $P_1, \ldots, P_r$ of $\{0,1\}^n$. Given a plaintext $x \in \{0,1\}^n$, the ciphertext y is computed as

$$y = k_r \oplus P_r(k_{r-1} \oplus P_{r-1}(\cdots P_2(k_1 \oplus P_1(k_0 \oplus x))\cdots)),$$

where the n-bit round keys $k_0, \ldots, k_r$ are either independent or derived from a master key k through key derivation functions $(\gamma_0, \ldots, \gamma_r)$. It abstracts in a generic way the high-level structure of most key-alternating ciphers such as AES. The nonexistence of *generic* attacks (i.e., attacks that are possible independently of a particular instantiation of the permutations $P_1, \ldots, P_r$) against this construction can be studied in the Random Permutation Model, where the P_i's are modeled as public random permutations to which the adversary is only given black-box (oracle) access.

The security of this construction in the traditional (single-key) indistinguishability framework (in other words, its pseudorandomness) has been extensively studied, starting with the seminal work of Even and Mansour for $r = 1$ round [EM97]. For an arbitrary number r of rounds, the case where all round keys are independent is by now well understood [BKL+12, Ste12, LPS12, CS14], and a tight security bound of $\mathcal{O}(2^{\frac{rn}{r+1}})$ queries has been established [CS14]. Chen *et al.* [CLL+14] also considered, for $r = 2$, the more complex case where the round keys are derived from an n-bit master key (as well as the case where the two inner permutations P_1 and P_2 are identical), and showed that a $\mathcal{O}(2^{\frac{2n}{3}})$-security bound still holds in that case.

On the other hand, two recent papers [ABD+13, LS13] explored a very strong security property of this construction, namely *(full) indifferentiability from an ideal cipher* (where "full" indifferentiability refers to the notion of Maurer *et al.* [MRH04]), which roughly ensures that the construction "behaves" in some well-defined sense as an ideal cipher, i.e., a block cipher drawn at random from the set of all block ciphers of some given block- and key-length. Andreeva *et al.* [ABD+13] showed that this property is achieved by the 5-round IEM cipher, assuming the key derivation function is modeled as a random oracle, while Lampe and Seurin [LS13] showed this for the 12-round IEM cipher, lifting the cryptographic requirement on the key derivation (namely, their result holds for the *trivial* key-schedule, i.e., when all round keys are equal to the n-bit master key).

In this paper, we complete the picture of the security of the IEM construction by considering security notions that lie between mere pseudorandomness and full indifferentiability from an ideal cipher, namely security against xor-induced related-key attacks (*XRKA* for short), i.e., related-key attacks where the adversary is allowed to xor any constant of its choice to the secret key, and against chosen-key attacks (*CKA* for short).

Related-Key Attacks. We start by considering XRKAs, which are important for at least two reasons. First, they arise naturally in a number of contexts, such as the f8 and f9 protocols of the 3GPP standard [IK04]. Second, from a theoretical point of view, they are the simplest kind of attacks to have the *completeness* property [GL10], namely, for any keys $k, k' \in \{0,1\}^n$, there exists $\Delta \in \{0,1\}^n$ such that $k \oplus \Delta = k'$. In order to study the resistance of the r-round IEM cipher to XRKAs, we use the traditional indistinguishability-based model of Bellare

and Kohno [BK03], albeit adapted to the Random Permutation Model. This means that the adversary has access to $r+1$ oracles: a *related key oracle* which takes as input an offset $\Delta \in \{0,1\}^n$ and a plaintext $x \in \{0,1\}^n$ (or a ciphertext $y \in \{0,1\}^n$), and r permutation oracles that we denote $P = (P_1, \ldots, P_r)$. The goal of the adversary is to distinguish two worlds: the "real" world, where on input (Δ, x), the related key oracle returns $\mathsf{EM}^P_{k\oplus\Delta}(x)$, where EM^P is the iterated Even-Mansour construction instantiated with permutations P and $k \in \{0,1\}^n$ is a random key, and the "ideal" world, where the related key oracle returns $E_{k\oplus\Delta}(x)$ for a random block cipher E independent from P. We start by describing a very efficient distinguishing XRKA on the 2-round IEM construction whenever the key derivation functions γ_i are linear (with respect to xor).[1] This somehow comes as a surprise since Bogdanov *et al.* [BKL+12] had previously conjectured that two rounds should be sufficient to prevent "certain types" of related-key attacks.[2] Motivated by this finding, we then consider what is the minimal number of rounds required to achieve provable security against XRKAs.[3] We first show that for the trivial key-schedule (all round keys are equal to the n-bit master key), the 3-round IEM cipher is secure against XRKAs up to $\mathcal{O}(2^{\frac{n}{2}})$ queries of the adversary. We conjecture that this bound is tight, but we were unable to find a matching attack (we also conjecture that a matching attack must be adaptive and make two-sided queries to the related-key oracle). If one is willing to use a cryptographically strong key-schedule, we show that a similar security bound is already attained with one round, assuming the key derivation functions are nonlinear (i.e., they have a small maximal differential probability). In this latter case, we note that our security bound is matched by a standard (i.e., non related-key) attack, namely Daemen's attack [Dae91].

Application to Tweakable Block Ciphers. We note that our results about the XRKA-security of the IEM construction have a direct application to the construction of tweakable block ciphers [LRW02] provably secure in the Random Permutation Model. Indeed, as hinted by Liskov *et al.* [LRW02] and formally proved by Bellare and Kohno [BK03, Theorem 7.1], an XRKA-secure block cipher E with key space $\{0,1\}^\kappa$ immediately gives rise to a secure tweakable block cipher $\widetilde{E}$ (with tweak space $\{0,1\}^\kappa$) by letting

$$\widetilde{E}(k,t,x) \stackrel{\text{def}}{=} E(k \oplus t, x),$$

where t is the tweak input. This, combined with Theorem 2, implies that the construction depicted on Figure 1 is a tweakable block cipher secure up to the birthday bound (in the Random Permutation Model). More precisely, the advantage of the best distinguisher against this construction is upper bounded by the same quantity as in Theorem 2.[4] A similar transformation

[1] Usually, in the case of standard (single-key) attacks, a distinguishing attack immediately gives rise to a key-recovery attack with similar complexity. This does not seem to be the case here, and we do not know whether our distinguishing XRKA can be converted into a key-recovery XRKA of similar complexity.

[2] The authors of [BKL+12] did not formulate any precise conjecture, but they mention that the best related-key attack they are aware of for two rounds and identical round keys is a *key-recovery* attack requiring $\mathcal{O}(2^{\frac{n}{2}})$ queries (see Appendix C.3 of the full version of their paper). Our own attack does not really improve on Bogdanov *et al.*'s one since it is a distinguishing attack, yet it implies that two rounds cannot be deemed secure against XRKAs.

[3] We only consider the case where all round keys are derived from the same n-bit master key k. Indeed, it is not hard to see that when round keys are independent, there are trivial XRKAs [BKL+12].

[4] Strictly speaking, Theorem 7.1 of [BK03] was proven in the standard model, but it can easily be seen to apply also in idealized models, such as the Random Permutation Model.

can be applied to the (1-round) Even-Mansour construction with a nonlinear key-schedule. We note that, at the high-level, the construction of Figure 1 is very similar in spirit to the "superposition TWEAKEY" constructions recently proposed by Jean *et al.* [JNP14], and our results can be seen as a first step towards proving the soundness of the TWEAKEY approach.

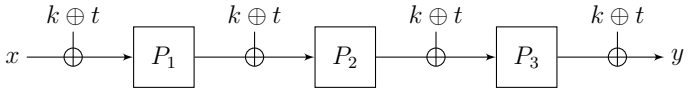

Fig. 1. A tweakable block cipher construction with tweak space $\{0,1\}^n$ secure up to birthday bound in the Random Permutation Model (the key k, the tweak t, and the plaintext x are all n-bit strings). The best distinguishing advantage against this construction is given by the upper bound of Theorem 2.

CHOSEN-KEY ATTACKS. We then turn our attention to an even stronger adversarial setting, namely chosen-key attacks [KR07, BKN09]. In this model, the adversary is given a block cipher, and its goal is, very informally, to exhibit some non-random behavior of the cipher, for keys and plaintext/ciphertext pairs of its choice. Rigorously formalizing what a non-random behavior means without ending with an unachievable definition turns out to be elusive for similar reasons that it is hard to rigorously define what collision resistance means for a single hash function [CGH98, Rog06].[5] Luckily, working in the Random Permutation Model allows us to escape those complications since it is somehow equivalent to considering a *large class of ciphers* consisting of all key-alternating ciphers of a given block-length and with a given key-schedule (rather than a single fully specified one, say, AES-128). In this setting, we are able to rigorously define resistance to CKAs thanks to the notion of *correlation intractability* first introduced by Canetti *et al.* [CGH98] in the context of hash functions.

The most convenient way we are aware of to prove that a block cipher construction is correlation intractable is to use a weakened variant of "full" indifferentiability [MRH04], named *sequential indifferentiability* (seq-indifferentiability for short), introduced by Mandal *et al.* [MPS12] to prove that the 6-round Feistel construction is correlation intractable. In a nutshell, a block cipher construction $\mathcal{C}^F$ based on an underlying ideal primitive F is (fully) indifferentiable from an ideal cipher if there exists a simulator $\mathcal{S}$ such that the two systems $(\mathcal{C}^F, F)$, where F is random, and $(E, \mathcal{S}^E)$, where E is an ideal (random) cipher, are indistinguishable by any (polynomially bounded) adversary $\mathcal{D}$. The distinguisher can query its two oracles as it wishes, and in the ideal world $(E, \mathcal{S}^E)$, the simulator is not aware of the queries made by $\mathcal{D}$ directly to E. Seq-indifferentiability is defined as full indifferentiability, except that the distinguisher is restricted to only query its right oracle in a first phase (F or $\mathcal{S}^E$), and then only its left oracle ($\mathcal{C}^F$ or E). Seq-indifferentiability is closely related to the notion of public indifferentiability [DRS09, YMO09], where in the ideal world the simulator gets to know all the queries of the distinguisher to the ideal primitive (i.e., the ideal cipher E in our context). We first give a "composition" theorem which relates seq-indifferentiability and correlation intractability (a similar one was already proved in [MPS12], but here we explicitly relate the various parameters since it is important for concrete security statements). Then, we

[5] For example, the fact that for any fixed block cipher E, $E_0(0)$ has some fixed, non-random value may be seen as a non-random behavior, yet arguably a harmless one.

prove that the 4-round IEM cipher, with the trivial key-schedule, is seq-indifferentiable from an ideal cipher (by a previous attack by Lampe and Seurin [LS13], this is also the minimal number of rounds to obtain this property). This implies by our composition theorem that the 4-round IEM cipher is correlation intractable, and hence offers some form of resistance to CKAs, but we warn that due to the quadratic query complexity of our simulator, the provable guarantee one obtains is not as tight as one might wish.

A Note on Known-Key Attacks. Known-key attacks refer, informally, to the setting where the adversary is given a block cipher E and a random key k, and must exhibit some non-random behavior of the permutation E_k [KR07]. In order to capture this security property, Andreeva *et al.* [ABM13] have introduced the notion of known-key indifferentiability (KK-indifferentiability), and they have proved that the 1-round Even-Mansour cipher is KK-indifferentiable from an ideal cipher. This might seem surprising at first sight since KKAs seem stronger than RKAs, yet the 1-round Even-Mansour cipher withstands the former but not the latter. We argue however that this is due to the fact that the KK-indifferentiability notion of [ABM13] is slightly too restrictive because it involves one single random key. We defer the details to Appendix C.

Table 1. Summary of provable security results for the iterated Even-Mansour cipher $\mathsf{EM}[n, r, \gamma]$ (with independent inner permutations). The *trivial* key-schedule means that all round keys are equal to the n-bit master key.

Sec. notion	# rounds	Key sched.	Sec. bound	Sim. complexity (query / time)	Ref.
Single-key	$r \geq 1$	independent	$2^{\frac{rn}{r+1}}$	—	[CS14]
	1	trivial	$2^{\frac{n}{2}}$	—	[EM97, DKS12]
	2	trivial	$2^{\frac{2n}{3}}$	—	[CLL+14]
XOR Related-Key	3	trivial	$2^{\frac{n}{2}}$	—	this paper
	1	nonlinear	$2^{\frac{n}{2}}$	—	this paper
Chosen-Key (Seq-indiff.)	4	trivial	$2^{\frac{n}{4}}$	q^2 / q^2	this paper
Full indiff.	5	random oracle	$2^{\frac{n}{10}}$	q^2 / q^3	[ABD+13]
	12	trivial	$2^{\frac{n}{12}}$	q^4 / q^6	[LS13]

Related Work. Provable security against RKAs was already considered in previous work. However, this was either for weak classes of RKAs (in particular, lacking the completeness property) [BK03, Luc04], or for inefficient number-theoretic constructions [BC10]. Our own results seem to be the first that hold both for a natural class of RKAs and for a practically-oriented construction. For provable security against CKAs, the only previous work we are aware of is [MPS12], which considered the 6-round Feistel construction.

In a concurrent and independent work, Farshim and Procter [FP15] also analyze the related-key security of the iterated Even-Mansour cipher. One of their main results (Corollary 3) is very similar to Theorem 2 in this paper; their bound is slightly worse than ours,

but their analysis is more general and applies to other families of related-key deriving functions than the xor-induced family. They also consider chosen-plaintext (related-key) attacks, whereas we directly consider chosen-plaintext and ciphertext attacks.

Open Problems. Regarding related-key security, it seems natural to conjecture that four rounds and the trivial key-schedule on one hand, or two rounds and a nonlinear key-schedule on the other hand, should deliver a $\mathcal{O}(2^{\frac{2n}{3}})$-security bound. If true, this should be provable by combining the techniques of [CLL+14] and the techniques of this paper. Regarding chosen-key security, an interesting open problem would be to find a construction of a block cipher from some underlying primitive (e.g., a random oracle or a small set of random permutations) which is seq-indifferentiable from an ideal cipher with a linear simulator complexity (indeed, by our composition theorem, this would imply an optimal resistance to CKAs). A first step in this direction was taken by Kiltz *et al.* [KPS13] in the context of digital signatures.

Organization. We set the notation and give some useful definitions in Section 2. We then consider the security of the IEM cipher against RKAs in Section 3 and against CKAs in Section 4.

2 Preliminaries

General Notation. In all the following, we fix an integer $n \geq 1$ and denote $N = 2^n$. The set of all permutations on $\{0,1\}^n$ will be denoted $\mathcal{P}_n$. A block cipher with key space $\{0,1\}^\kappa$ and message space $\{0,1\}^n$ is a mapping $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ such that for any key $k \in \{0,1\}^\kappa$, $x \mapsto E(k,x)$ is a permutation. We interchangeably use the notations $E(k,x)$ and $E_k(x)$. We denote $\mathsf{BC}(\kappa, n)$ the set of all block ciphers with key space $\{0,1\}^\kappa$ and message space $\{0,1\}^n$. For integers $1 \leq s \leq t$, we will write $(t)_s = t(t-1)\cdots(t-s+1)$ and $(t)_0 = 1$ by convention. For a function $f : \{0,1\}^n \to \{0,1\}^n$, let

$$\delta(f) = \max_{a,b \in \{0,1\}^n, a \neq 0} |\{x \in \{0,1\}^n : f(x \oplus a) \oplus f(x) = b\}|.$$

Note that $\delta(f)$ is a measure of the nonlinearity of f. A permutation f of $\{0,1\}^n$ is said *almost perfect nonlinear* [NK92] if $\delta(f) = 2$.

The Iterated Even-Mansour Cipher. Fix integers $n, r \geq 1$. Let $\gamma = (\gamma_0, \ldots, \gamma_r)$ be a $(r+1)$-tuple of permutations of $\{0,1\}^n$. The r-round iterated Even-Mansour construction $\mathsf{EM}[n, r, \gamma]$ specifies, from any r-tuple $P = (P_1, \ldots, P_r)$ of permutations of $\{0,1\}^n$, a block cipher with n-bit keys and n-bit messages, simply denoted EM^P in all the following (parameters $[n, r, \gamma]$ will always be clear from the context), which maps a plaintext $x \in \{0,1\}^n$ and a key $k \in \{0,1\}^n$ to the ciphertext defined by (see Figure 2):

$$\mathsf{EM}^P(k,x) = \gamma_r(k) \oplus P_r(\gamma_{r-1}(k) \oplus P_{r-1}(\cdots P_2(\gamma_1(k) \oplus P_1(\gamma_0(k) \oplus x))\cdots)).$$

The pseudorandomness of the IEM cipher was mostly studied for the case of *independent* round keys [BKL+12, LPS12, CS14], with the notable exception of [CLL+14]. In this paper, we focus on the case where the round keys are derived from an n-bit master key.

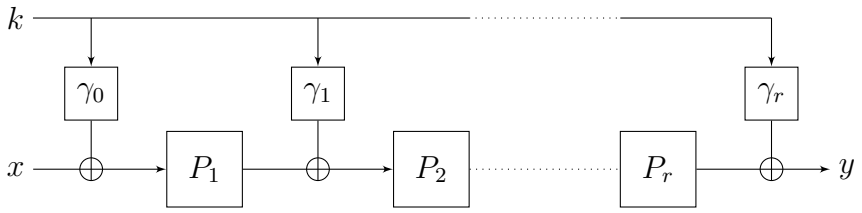

Fig. 2. The r-round iterated Even-Mansour cipher.

Related-Key Oracle. Let $E \in \mathsf{BC}(\kappa, n)$ be a block cipher, and fix a key $k \in \{0,1\}^\kappa$. We define the xor-restricted related-key oracle $\mathsf{RK}[E_k]$, which takes as input an "offset" $\Delta \in \{0,1\}^\kappa$ and a plaintext $x \in \{0,1\}^n$, and returns $\mathsf{RK}[E_k](\Delta, x) := E_{k\oplus\Delta}(x)$. The oracle can be queried backward, namely $\mathsf{RK}[E_k]^{-1}(\Delta, y) := E^{-1}_{k\oplus\Delta}(y)$.

3 Resistance to Related-Key Attacks

3.1 Security Definitions

To formalize related-key attacks against the r-round IEM cipher, we extend in a straightforward way the classical Bellare-Kohno model [BK03] to the case where the adversary has access to additional oracles. Formally, we consider a xor-restricted related-key adversary $\mathcal{D}$ which has access to $r+1$ oracles, a related-key oracle and r permutation oracles, and must distinguish between the following two worlds:

- the "real" world, where it interacts with $(\mathsf{RK}[\mathsf{EM}_k^P], P)$ where $P = (P_1, \ldots, P_r)$ is a tuple of random permutations and k is a randomly drawn key;
- the "ideal" world where it interacts with $(\mathsf{RK}[E_k], P)$ where $P = (P_1, \ldots, P_r)$ is a tuple of random permutations, E is an ideal cipher independent from P, and k a randomly drawn key.

The distinguisher is adaptive, and can make two-sided queries to each oracle. As usual, we assume that it is computationally unbounded, deterministic, and never makes pointless queries. Note that in the ideal world, the key k is meaningless, and the related-key oracle $\mathsf{RK}[E_k]$ simply implements an independent random permutation for each offset $\Delta \in \{0,1\}^n$.

The distinguishing advantage of $\mathcal{D}$ is defined as

$$\mathbf{Adv}(\mathcal{D}) = \left|\Pr\left[\mathcal{D}^{\mathsf{RK}[\mathsf{EM}_k^P],P} = 1\right] - \Pr\left[\mathcal{D}^{\mathsf{RK}[E_k],P} = 1\right]\right|,$$

where the first probability is taken over the random choice of k and P, and the second probability is taken over the random choice of E, k, and P.

For q_e, q_p non-negative integers, we define the insecurity of the iterated Even-Mansour cipher against xor-restricted related-key attacks as

$$\mathbf{Adv}^{\text{xor-rka}}_{\mathsf{EM}[n,r,\gamma]}(q_e, q_p) = \max_{\mathcal{D}} \mathbf{Adv}(\mathcal{D}),$$

where the maximum is taken over all distinguishers making exactly q_e queries to the related-key oracle and exactly q_p queries to each inner permutation oracle.

TRANSCRIPT. We summarize the information gathered by the distinguisher in what we call the *query transcript* $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_r})$, defined as follows. The tuple

$$\mathcal{Q}_E = ((\Delta_1, x_1, y_1), \ldots, (\Delta_{q_e}, x_{q_e}, y_{q_e}))$$

summarizes the queries to the related-key oracle, and means that the j-th query was either a forward query (Δ_j, x_j) and the answer y_j, or a backward query (Δ_j, y_j) and the answer x_j. Similarly, the tuple

$$\mathcal{Q}_{P_i} = ((u_{i,1}, v_{i,1}), \ldots, (u_{i,q_p}, v_{i,q_p}))$$

summarizes the queries to the i-th inner permutation P_i, and means that the j-th query was either a forward query $u_{i,j}$ and the answer $v_{i,j}$, or a backward query $v_{i,j}$ and the answer $u_{i,j}$. (Recall that the distinguisher is deterministic, so that there is a one-to-one mapping between this directionless representation and the raw transcript of the interaction of the distinguisher with the oracles). A query transcript is said *attainable* if the probability to obtain it in the ideal world is non-zero (hence, the set of attainable query transcripts depends on the distinguisher). To simplify the security proof (in particular, the definition of *bad* transcripts), we reveal to the distinguisher the key k at the end of its query phase (this is without loss of generality since $\mathcal{D}$ is free to ignore this additional information to compute its output bit). Formally, we append k to the query transcript $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_r})$, obtaining what we will simply call the *transcript* $\tau = (\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_r}, k)$ of the attack. A transcript τ is said attainable if the corresponding query transcript is attainable. We denote $\mathcal{T}$ the set of attainable transcripts. In all the following, we denote T_{re}, resp. T_{id}, 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.

ADDITIONAL NOTATION. Given a block cipher $E \in \mathsf{BC}(n, n)$, a key $k \in \{0,1\}^n$, and a related-key oracle query transcript $\mathcal{Q}_E$, we say that (E, k) *extends* $\mathcal{Q}_E$, written $(E, k) \vdash \mathcal{Q}_E$, if $E_{k \oplus \Delta}(x) = y$ for each $(\Delta, x, y) \in \mathcal{Q}_E$. Similarly, given a permutation P and a permutation query transcript $\mathcal{Q}_P$, we say that P extends $\mathcal{Q}_P$, written $P \vdash \mathcal{Q}_P$, if $P(u) = v$ for each $(u, v) \in \mathcal{Q}_P$. It is easy to see that for any attainable transcript $\tau = (\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_r}, k)$, the interaction of the distinguisher with oracles $(\mathsf{RK}[E_k], P_1, \ldots, P_r)$ produces τ *iff* $(E, k) \vdash \mathcal{Q}_E$ and $P_i \vdash \mathcal{Q}_{P_i}$ for $i = 1, \ldots, r$.

THE H-COEFFICIENTS TECHNIQUE. We will use the H-coefficients technique [Pat08], which relies on the following lemma. See e.g. [CS14, CLL+14] for a proof.

Lemma 1. *Fix a distinguisher $\mathcal{D}$. Let $\mathcal{T} = \mathcal{T}_{\text{good}} \sqcup \mathcal{T}_{\text{bad}}$ be a partition of the set of attainable transcripts. Assume that there exists ε_1 such that for any $\tau \in \mathcal{T}_{\text{good}}$, one has*[6]

$$\frac{\Pr[T_{\text{re}} = \tau]}{\Pr[T_{\text{id}} = \tau]} \geq 1 - \varepsilon_1,$$

and that there exists ε_2 such that $\Pr[T_{\text{id}} \in \mathcal{T}_{\text{bad}}] \leq \varepsilon_2$. Then $\mathbf{Adv}(\mathcal{D}) \leq \varepsilon_1 + \varepsilon_2$.

[6] Recall that for an attainable transcript, one has $\Pr[T_{\text{id}} = \tau] > 0$.

3.2 The Linear Key-Schedule Case

In this section, we consider xor-induced related-key attacks against the IEM cipher with independent permutations and a linear key-schedule. We give attacks for up to two rounds, and then prove a $\mathcal{O}(2^{\frac{n}{2}})$-security bound for three rounds.

A Simple Attack on One Round. We start with a very simple attack for one round. Given a permutation P on $\{0,1\}^n$ and two linear permutations $\gamma_0, \gamma_1 : \{0,1\}^n \to \{0,1\}^n$, consider the 1-round Even-Mansour cipher which maps a key $k \in \{0,1\}^n$ and a plaintext $x \in \{0,1\}^n$ to the ciphertext defined as

$$\mathsf{EM}^P(k,x) = \gamma_1(k) \oplus P(\gamma_0(k) \oplus x).$$

Consider the distinguisher which simply queries the related-key oracle on two inputs $(0,x)$ and $(\Delta, x \oplus \gamma_0(\Delta))$, where $\Delta \neq 0$, getting respective answers y and y', and checks whether $y' = y \oplus \gamma_1(\Delta)$. This holds with probability 1 in the real world, but only with probability $1/N$ in the ideal world, so that the distinguishing advantage of this adversary is negligibly close to one.

An attack on Two Rounds. We then show a more intricate distinguishing attack for two rounds (and, again, a linear key-schedule). This attack does not require to query the internal permutation oracles, and makes only four queries to the related-key oracle. It can be seen as a very efficient boomerang related-key attack [BDK05]. Formally, we prove the following theorem.

Theorem 1. *Let $\gamma = (\gamma_0, \gamma_1, \gamma_2)$ be a linear key-schedule. Then*

$$\mathbf{Adv}^{\text{xor-rka}}_{\mathsf{EM}[n,2,\gamma]}(4,0) \geq 1 - \frac{1}{N}.$$

Proof. We denote generically $(\mathsf{RK}, (P_1, P_2))$ the oracles to which the adversary has access. Consider the following distinguisher (see Figure 3 for a diagram of the attack):

(1) choose arbitrary values $x_1, \Delta_1 \in \{0,1\}^n$, and query $y_1 := \mathsf{RK}(\Delta_1, x_1)$;
(2) choose an arbitrary value $\Delta_2 \in \{0,1\}^n \setminus \{\Delta_1\}$, compute $x_2 := x_1 \oplus \gamma_0(\Delta_2 \oplus \Delta_1)$, and query $y_2 := \mathsf{RK}(\Delta_2, x_2)$;
(3) choose an arbitrary $\Delta_3 \in \{0,1\}^n \setminus \{\Delta_1, \Delta_2\}$, compute $y_3 := y_1 \oplus \gamma_2(\Delta_1 \oplus \Delta_3)$, and query $x_3 := \mathsf{RK}^{-1}(\Delta_3, y_3)$;
(4) compute $\Delta_4 := \Delta_3 \oplus \Delta_2 \oplus \Delta_1$ and $y_4 := y_2 \oplus \gamma_2(\Delta_2 \oplus \Delta_4)$, and query $x_4 := \mathsf{RK}^{-1}(\Delta_4, y_4)$;
(5) if $x_4 = x_3 \oplus \gamma_0(\Delta_3 \oplus \Delta_4)$, output 1, else output 0.

When the distinguisher is interacting with the ideal world $(\mathsf{RK}[E], (P_1, P_2))$, where E is an ideal cipher independent from P_1 and P_2, the value x_4 is uniformly random and independent from x_3, Δ_3, and Δ_4 (indeed the offsets Δ_i for $i = 1,2,3,4$ are pairwise distinct, so that y_4 is the first query to the random permutation corresponding to offset Δ_4). Hence, the probability that the distinguisher returns 1 in the ideal case is 2^{-n}.

Now we show that when the distinguisher is interacting with the real world, i.e., with $(\mathsf{RK}[\mathsf{EM}_k^{P_1,P_2}], (P_1, P_2))$, it always returns 1, independently of k, P_1, and P_2. Noting that, by definition, $x_2 = x_1 \oplus \gamma_0(\Delta_2 \oplus \Delta_1)$, we denote u_1 the common value

$$u_1 \stackrel{\text{def}}{=} x_1 \oplus \gamma_0(k \oplus \Delta_1) = x_2 \oplus \gamma_0(k \oplus \Delta_2),$$

and we denote $v_1 = P_1(u_1)$. We also denote

$$\begin{aligned} u_2 &= v_1 \oplus \gamma_1(k \oplus \Delta_1) && (1)\\ v_2 &= P_2(u_2) \\ u'_2 &= v_1 \oplus \gamma_1(k \oplus \Delta_2) && (2)\\ v'_2 &= P_2(u'_2). \end{aligned}$$

Hence, one has

$$\begin{aligned} y_1 &= v_2 \oplus \gamma_2(k \oplus \Delta_1) && (3)\\ y_2 &= v'_2 \oplus \gamma_2(k \oplus \Delta_2). && (4) \end{aligned}$$

Since $y_3 = y_1 \oplus \gamma_2(\Delta_1 \oplus \Delta_3)$, we can see, using (3), that

$$y_3 \oplus \gamma_2(k \oplus \Delta_3) = y_1 \oplus \gamma_2(k \oplus \Delta_1) = v_2.$$

Define

$$\begin{aligned} v'_1 &= u_2 \oplus \gamma_1(k \oplus \Delta_3) && (5)\\ u'_1 &= P_1^{-1}(v'_1). \end{aligned}$$

This implies that

$$x_3 = u'_1 \oplus \gamma_0(k \oplus \Delta_3). \quad (6)$$

Since $y_4 = y_2 \oplus \gamma_2(\Delta_2 \oplus \Delta_4)$, we see by (4) that

$$y_4 \oplus \gamma_2(k \oplus \Delta_4) = y_2 \oplus \gamma_2(k \oplus \Delta_2) = v'_2.$$

Moreover, since $\Delta_4 = \Delta_3 \oplus \Delta_2 \oplus \Delta_1$, we have

$$\begin{aligned} u'_2 \oplus \gamma_1(k \oplus \Delta_4) &= u'_2 \oplus \gamma_1(k \oplus \Delta_2) \oplus \gamma_1(\Delta_1 \oplus \Delta_3) \\ &= v_1 \oplus \gamma_1(k \oplus \Delta_1) \oplus \gamma_1(k \oplus \Delta_3) && \text{by (2)} \\ &= u_2 \oplus \gamma_1(k \oplus \Delta_3) && \text{by (1)} \\ &= v'_1 && \text{by (5)}. \end{aligned}$$

This finally implies by (6) that

$$x_4 = u'_1 \oplus \gamma_0(k \oplus \Delta_4) = x_3 \oplus \gamma_0(\Delta_3 \oplus \Delta_4),$$

which concludes the proof. □

Security Proof for Three Rounds. We consider the 3-round IEM cipher with the trivial key schedule (the result can be straightforwardly extended to the general case where the key derivation functions $(\gamma_0, \ldots, \gamma_3)$ are any permutations). Given three permutations P_1, P_2, P_3 on $\{0,1\}^n$, we denote $\mathsf{EM}^{P_1,P_2,P_3}$ the 3-round IEM cipher which maps a key $k \in \{0,1\}^n$ and a plaintext $x \in \{0,1\}^n$ to the ciphertext defined as

$$\mathsf{EM}^{P_1,P_2,P_3}(k,x) = k \oplus P_3(k \oplus P_2(k \oplus P_1(k \oplus x))).$$

We prove the following result.

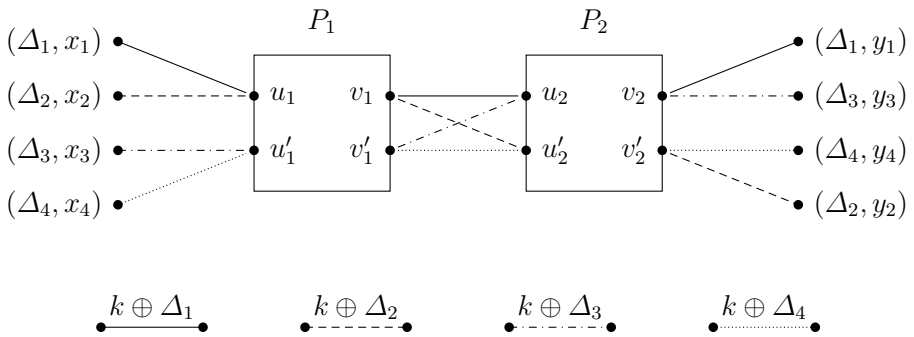

Fig. 3. A related-key attack on the iterated Even-Mansour cipher with two rounds and a linear key-schedule.

Theorem 2. *Let q_e, q_p be positive integers, $N = 2^n$, and $\mathcal{I}$ be the trivial key-schedule. Then*

$$\mathbf{Adv}^{\text{xor-rka}}_{\mathsf{EM}[n,3,\mathcal{I}]}(q_e, q_p) \leq \frac{6q_eq_p}{N} + \frac{4q_e^2}{N}.$$

Proof. The proof follows from Lemma 1, and Lemmas 2 and 3 proven below. □

The security bound of Theorem 2 is plotted in the (q_e, q_p) plane on Figure 4. Regarding tightness (with respect to information-theoretic adversaries) of this bound, the best attack we are aware of is the traditional (single-key) attack against the 3-round IEM construction [BKL+12, Gaz13], which requires $q_eq_p^3 \sim 2^{3n}$ (see Figure 4).

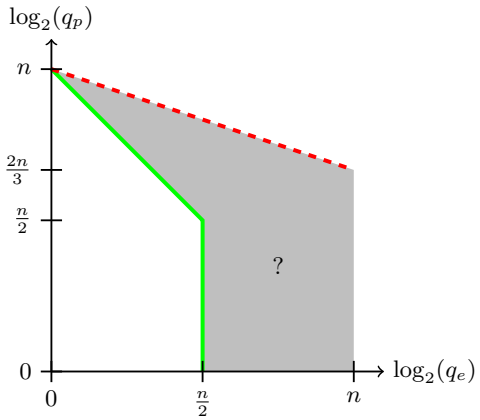

Fig. 4. The XRKA security of the 3-round IEM construction with the trivial key-schedule. All parameters below the solid line are secure by Theorem 2, while all parameters above the dashed line are insecure by the single-key attack of [BKL+12]. The security for parameters between these two lines remains unknown.

The remaining of this section is devoted to the proof of Theorem 2. Following the H-coefficient technique, we start by defining bad transcripts.

Definition 1. *Let $\tau = (\mathcal{Q}_E, \mathcal{Q}_{P_1}, \mathcal{Q}_{P_2}, \mathcal{Q}_{P_3}, k)$ be an attainable transcript. We say that τ is* bad *if*

$$k \in \mathsf{BadK} = \bigcup_{1 \le i \le 2} \mathsf{BadK}_i$$

where:

$k \in \mathsf{BadK}_1 \Leftrightarrow$ *there exists* $(\Delta, x, y) \in \mathcal{Q}_E$ *and* $(u_1, v_1) \in \mathcal{Q}_{P_1}$ *such that* $k \oplus \Delta = x \oplus u_1$
$k \in \mathsf{BadK}_2 \Leftrightarrow$ *there exists* $(\Delta, x, y) \in \mathcal{Q}_E$ *and* $(u_3, v_3) \in \mathcal{Q}_{P_3}$ *such that* $k \oplus \Delta = y \oplus v_3$.

Otherwise, τ is said good. *We denote $\mathcal{T}_{\text{bad}}$ the set of bad transcripts, and $\mathcal{T}_{\text{good}} = \mathcal{T} \setminus \mathcal{T}_{\text{bad}}$ the set of good transcripts.*

First, we upper bound the probability to get a bad transcript in the ideal world.

Lemma 2.

$$\Pr[T_{\text{id}} \in \mathcal{T}_{\text{bad}}] \le \frac{2q_e q_p}{N}.$$

Proof. Since we are in the ideal case, the key k is drawn uniformly at random at the end of the query phase. Hence, we only need to upper bound the number of possible bad values for k for every attainable query transcripts $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \mathcal{Q}_{P_2}, \mathcal{Q}_{P_3})$. Fix any query transcript $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \mathcal{Q}_{P_2}, \mathcal{Q}_{P_3})$. Then, for every $(\Delta, x, y) \in \mathcal{Q}_E$ and every $(u_1, v_1) \in \mathcal{Q}_{P_1}$, there is exactly one key k such that $k = x \oplus \Delta \oplus u_1$. Hence, $|\mathsf{BadK}_1| \le q_e q_p$. Similarly, $|\mathsf{BadK}_2| \le q_e q_p$. Hence, for $i = 1, 2$,

$$\Pr\left[k \leftarrow_{\$} \{0,1\}^n : k \in \mathsf{BadK}_i\right] \le \frac{q_e q_p}{N}.$$

The result follows. □

We then consider good transcripts in the following lemma.

Lemma 3. *For any good transcript $\tau \in \mathcal{T}_{\text{good}}$, one has*

$$\frac{\Pr\left[T_{\text{re}} = \tau\right]}{\Pr\left[T_{\text{id}} = \tau\right]} \ge 1 - \frac{4q_e q_p}{N} - \frac{4q_e^2}{N}.$$

Proof. If $\mathcal{T}_{\text{good}} = \emptyset$, there is nothing to prove. Otherwise, fix a good transcript $\tau = (\mathcal{Q}_E, \mathcal{Q}_{P_1}, \mathcal{Q}_{P_2}, \mathcal{Q}_{P_3}, k)$. Let m denote the number of different offsets Δ appearing in $\mathcal{Q}_E$ and q_i the number of queries using the i-th offset (ordering the offsets arbitrarily). Note that $q_e = \sum_{i=1}^m q_i$. In the ideal world, one simply has

$$\begin{aligned}\Pr[T_{\text{id}} = \tau] &= \Pr[k' \leftarrow_{\$} \{0,1\}^n : k' = k] \times \Pr[P_i \leftarrow_{\$} \mathcal{P}_n : P_i \vdash \mathcal{Q}_{P_i}, i = 1,2,3] \\ &\qquad \times \Pr[E \leftarrow_{\$} \mathsf{BC}(n,n) : (E,k) \vdash \mathcal{Q}_E] \\ &= \frac{1}{N} \cdot \frac{1}{((N)_{q_p})^3} \cdot \frac{1}{\prod_{i=1}^m (N)_{q_i}}. \end{aligned} \tag{7}$$

Now we have to lower bound the probability

$$\Pr\left[T_{\text{re}} = \tau\right] = \frac{1}{N} \times \Pr\left[P_1, P_2, P_3 \leftarrow_{\$} \mathcal{P}_n : (\mathsf{EM}^{P_1,P_2,P_3}, k) \vdash \mathcal{Q}_E \wedge P_i \vdash \mathcal{Q}_{P_i}, i = 1,2,3\right].$$

Let

$$U_1 = \{u_1 \in \{0,1\}^n : (u_1, v_1) \in \mathcal{Q}_{P_1}\}, \qquad V_1 = \{v_1 \in \{0,1\}^n : (u_1, v_1) \in \mathcal{Q}_{P_1}\},$$
$$U_2 = \{u_2 \in \{0,1\}^n : (u_2, v_2) \in \mathcal{Q}_{P_2}\}, \qquad V_2 = \{v_2 \in \{0,1\}^n : (u_2, v_2) \in \mathcal{Q}_{P_2}\},$$
$$U_3 = \{u_3 \in \{0,1\}^n : (u_3, v_3) \in \mathcal{Q}_{P_3}\}, \qquad V_3 = \{v_3 \in \{0,1\}^n : (u_3, v_3) \in \mathcal{Q}_{P_3}\}$$

denote the domains and ranges of $\mathcal{Q}_{P_1}$, $\mathcal{Q}_{P_2}$, and $\mathcal{Q}_{P_3}$ respectively. For $u'_1 \in \{0,1\}^n$, let $X(u'_1) = \{(\Delta, x, y) \in \mathcal{Q}_E : x \oplus k \oplus \Delta = u'_1\}$, and let $U'_1 = \{u'_1 \in \{0,1\}^n : X(u'_1) \neq \emptyset\}$. Similarly, for $v'_3 \in \{0,1\}^n$, let $Y(v'_3) = \{(\Delta, x, y) \in \mathcal{Q}_E : y \oplus k \oplus \Delta = v'_3\}$, and let $V'_3 = \{v'_3 \in \{0,1\}^n : Y(v'_3) \neq \emptyset\}$. Note that by definition of a good transcript, one has $U_1 \cap U'_1 = \emptyset$ and $V_3 \cap V'_3 = \emptyset$. Let also $\alpha = |U'_1|$ and $\beta = |V'_3|$. For clarity, we denote

$$U'_1 = \{u'_{1,1}, \ldots, u'_{1,\alpha}\}$$
$$V'_3 = \{v'_{3,1}, \ldots, v'_{3,\beta}\}$$

using an arbitrary order. Note that

$$q_e = \sum_{i=1}^{\alpha} |X(u'_{1,i})| = \sum_{i=1}^{\beta} |Y(v'_{3,i})|. \tag{8}$$

It is now sufficient for our result to lower bound the number of possible tuple of values $(v'_{1,1}, \ldots, v'_{1,\alpha})$ and $(u'_{3,1}, \ldots, u'_{3,\beta})$ such that, conditioned on $P_1(u'_{1,i}) = v'_{1,i}$ for $1 \leq i \leq \alpha$ and $P_3(u'_{3,j}) = v'_{3,j}$ for $1 \leq j \leq \beta$, the event $E_k^{P_1,P_2,P_3} \vdash \mathcal{Q}_E$ is equivalent to q_e "new" equations on P_2 (i.e., distinct from equations imposed by $P_2 \vdash \mathcal{Q}_{P_2}$). More precisely, let N_1 be the number of tuples of pairwise distinct values $(v'_{1,1}, \ldots, v'_{1,\alpha})$ such that, for every $i = 1, \ldots, \alpha$:

(i) $v'_{1,i} \neq v_1$ for every $v_1 \in V_1$,
(ii) $v'_{1,i} \neq k \oplus \Delta \oplus u_2$ for every $(\Delta, x, y) \in X(u'_{1,i})$, $u_2 \in U_2$,
(iii) $v'_{1,i} \neq \Delta \oplus v'_{1,j} \oplus \Delta'$ for every $(\Delta, x, y) \in X(u'_{1,i})$, $1 \leq j \leq i-1$, $(\Delta', x', y') \in X(u'_{1,j})$.

Then

$$\begin{aligned} N_1 &\geq \prod_{i=1}^{\alpha} \left(N - q_p - i + 1 - |X(u'_{1,i})|(q_p + \sum_{j=1}^{i-1} |X(u'_{1,j})|) \right) \\ &\geq \prod_{i=1}^{\alpha} \Big(N - q_p - q_e - |X(u'_{1,i})|(q_p + q_e) \Big) \qquad \text{by (8)}. \end{aligned}$$

Similarly, let N_3 be the number of tuples of pairwise distinct values $(u'_{3,1}, \ldots, u'_{3,\beta})$ such that, for every $i = 1, \ldots, \beta$:

(i') $u'_{3,i} \neq u_3$ for every $u_3 \in U_3$,
(ii') $u'_{3,i} \neq k \oplus \Delta \oplus v_2$ for every $(\Delta, x, y) \in Y(v'_{3,i})$, $v_2 \in V_2$,
(iii') $u'_{3,i} \neq \Delta \oplus u'_{3,j} \oplus \Delta'$ for every $(\Delta, x, y) \in Y(v'_{3,i})$, $1 \leq j \leq i-1$, $(\Delta', x', y') \in Y(v'_{3,j})$.

Then

$$\begin{aligned} N_3 &\geq \prod_{i=1}^{\beta} \left(N - q_p - i + 1 - |Y(v'_{3,i})|(q_p + \sum_{j=1}^{i-1} |Y(v'_{3,j})|) \right) \\ &\geq \prod_{i=1}^{\beta} \Big(N - q_p - q_e - |Y(v'_{3,i})|(q_p + q_e) \Big) \qquad \text{by (8)}. \end{aligned}$$

For every possible choice of $(v'_{1,1}, \ldots, v'_{1,\alpha})$ and $(u'_{3,1}, \ldots, u'_{3,\beta})$ satisfying these conditions, P_1 will be fixed on exactly $q_p + \alpha$ points, P_2 on $q_p + q_e$ points and P_3 on $q_p + \beta$ points. In more details, assume $N_1 \cdot N_3 > 0$, fix any tuples $(v'_{1,1}, \ldots, v'_{1,\alpha})$ and $(u'_{3,1}, \ldots, u'_{3,\beta})$ satisfying these conditions, and let Ev_1 be the event that $P_1(u'_{1,i}) = v'_{1,i}$ for $1 \leq i \leq \alpha$ and Ev_3 be the event that $P_3(u'_{3,j}) = v'_{3,j}$ for $1 \leq j \leq \beta$. Then by conditions (i) and (i') we have

$$\Pr\left[\mathsf{Ev}_1 \wedge (P_1 \vdash \mathcal{Q}_{P_1})\right] = \frac{1}{(N)_{q_p+\alpha}}$$
$$\Pr\left[\mathsf{Ev}_3 \wedge (P_3 \vdash \mathcal{Q}_{P_3})\right] = \frac{1}{(N)_{q_p+\beta}}.$$

Fix now P_1 and P_3 satisfying Ev_1 and Ev_3. For each $(\Delta, x, y) \in \mathcal{Q}_E$, let u'_2 and v'_2 be respectively the corresponding input and output to P_2 for this query, viz., $u'_2 = v'_{1,i} \oplus k \oplus \Delta$ for i such that $x \oplus k \oplus \Delta = u'_{1,i}$, and $v'_2 = u'_{3,j} \oplus k \oplus \Delta$ for j such that $y \oplus k \oplus \Delta = v'_{3,j}$. Then, the q_e values u'_2 are all outside U_2 by condition (ii), and pairwise distinct by condition (iii), and similarly the q_e values v'_2 are all outside V_2 by condition (ii'), and pairwise distinct by condition (iii'). It follows that

$$\Pr\left[(\mathsf{EM}^{P_1,P_2,P_3}, k) \vdash \mathcal{Q}_E \wedge (P_2 \vdash \mathcal{Q}_{P_2}) \middle| \mathsf{Ev}_1 \wedge (P_1 \vdash \mathcal{Q}_{P_1}) \wedge \mathsf{Ev}_3 \wedge (P_3 \vdash \mathcal{Q}_{P_3})\right] = \frac{1}{(N)_{q_p+q_e}}.$$

Hence, summing over the at least $N_1 \cdot N_3$ possible pairs of tuples, we obtain

$$\Pr\left[T_{\mathrm{re}} = \tau\right] \geq \frac{N_1 \cdot N_3}{N \cdot (N)_{q_p+\alpha} \cdot (N)_{q_p+q_e} \cdot (N)_{q_p+\beta}}. \tag{9}$$

This last inequality is also trivially true if $N_1 \cdot N_3 = 0$. Using (7) and (9), one has

$$\begin{aligned}
\frac{\Pr\left[T_{\mathrm{re}} = \tau\right]}{\Pr\left[T_{\mathrm{id}} = \tau\right]} &\geq \frac{N_1 \cdot N_3 \cdot N \cdot (N)_{q_p}^3 \prod_{i=1}^{m} (N)_{q_i}}{N \cdot (N)_{q_p+\alpha} \cdot (N)_{q_p+q_e} \cdot (N)_{q_p+\beta}} \\
&\geq \frac{N_1 \cdot N_3 \cdot \prod_{i=1}^{m} (N)_{q_i}}{(N-q_p)_{\alpha} \cdot (N-q_p)_{q_e} \cdot (N-q_p)_{\beta}} \\
&\geq \frac{N_1 \cdot N_3 \cdot (N)_{q_e}}{(N-q_p)_{\alpha} \cdot (N-q_p)_{q_e} \cdot (N-q_p)_{\beta}} \\
&\geq \frac{N_1 \cdot N_3}{N^{\alpha+\beta}}.
\end{aligned}$$

Finally, one has, since $\alpha \leq q_e$,

$$\begin{aligned}
\frac{N_1}{N^{\alpha}} &= \frac{\prod_{i=1}^{\alpha} \left(N - q_p - q_e - |X(u'_{1,i})|(q_p + q_e)\right)}{N^{\alpha}} \\
&\geq 1 - \sum_{i=1}^{\alpha} \frac{q_p + q_e + |X(u'_{1,i})|(q_p + q_e)}{N} \\
&\geq 1 - \frac{q_e q_p}{N} - \frac{q_e^2}{N} - (q_p + q_e) \sum_{i=1}^{\alpha} \frac{|X(u'_{1,i})|}{N} \\
&\geq 1 - \frac{2q_e q_p}{N} - \frac{2q_e^2}{N} \qquad \text{by (8)}.
\end{aligned}$$

The same lower bound holds for $\frac{N_3}{N^\beta}$. Hence

$$\frac{\Pr[T_{\mathrm{re}} = \tau]}{\Pr[T_{\mathrm{id}} = \tau]} \geq \left(1 - \frac{2q_e q_p}{N} - \frac{2q_e^2}{N}\right)^2$$
$$\geq 1 - \frac{4q_e q_p}{N} - \frac{4q_e^2}{N}.$$

□

3.3 The Nonlinear Key-Schedule Case

In this section, we show that when the key-schedule is nonlinear, one round is sufficient to achieve a $\mathcal{O}(2^{\frac{n}{2}})$-security bound against xor-induced related-key attacks.

Given a permutation P on $\{0,1\}^n$ and two permutations $\gamma_0, \gamma_1 : \{0,1\}^n \to \{0,1\}^n$, we denote EM^P the 1-round Even-Mansour cipher which maps a key $k \in \{0,1\}^n$ and a plaintext $x \in \{0,1\}^n$ to the ciphertext defined as

$$\mathsf{EM}^P(k,x) = \gamma_1(k) \oplus P(\gamma_0(k) \oplus x).$$

We prove the following result.

Theorem 3. *Let q_e, q_p be positive integers, $N = 2^n$, and $\gamma = (\gamma_0, \gamma_1)$. Then*

$$\mathbf{Adv}^{\text{xor-rka}}_{\mathsf{EM}[n,1,\gamma]}(q_e, q_p) \leq \frac{2q_e q_p}{N} + \frac{(\delta(\gamma_0) + \delta(\gamma_1))q_e^2}{2N}.$$

In particular, if γ_0 and γ_1 are almost perfect nonlinear permutations, then

$$\mathbf{Adv}^{\text{xor-rka}}_{\mathsf{EM}[n,1,\gamma]}(q_e, q_p) \leq \frac{2q_e q_p + 2q_e^2}{N}.$$

Proof. Deferred to Appendix B. □

The security bound of Theorem 3 is plotted in the (q_e, q_p) plane on Figure 5. Regarding tightness (with respect to information-theoretic adversaries) of this bound, we note that usual (single-key) attacks against the traditional Even-Mansour cipher [Dae91, DKS12] apply in this setting as well. It requires $q_e q_p \sim 2^n$, so that our bound is tight for $q_e \leq 2^{n/2}$ (Figure 5).

4 Resistance to Chosen-Key Attacks and Sequential Indifferentiability

4.1 Formalizing Chosen-Key Attacks in Idealized Models

In this section, we see a block cipher $E \in \mathsf{BC}(\kappa, n)$ as a primitive which takes as input a triple $\alpha = (\delta, k, z)$, where $\delta \in \{+,-\}$ indicates whether this is a direct (plaintext) or inverse (ciphertext) query, $k \in \{0,1\}^\kappa$ is the key, and $z \in \{0,1\}^n$ is the plaintext/ciphertext (depending on δ), and returns the corresponding ciphertext/plaintext (again, depending on δ) $z' \in \{0,1\}^n$. This allows the block cipher to be described as having a single interface rather than two interfaces E and E^{-1}. In the following, we denote $\mathsf{Dom} = \{+,-\} \times \{0,1\}^\kappa \times \{0,1\}^n$ and $\mathsf{Rng} = \{0,1\}^n$ respectively the domain and the range of E. For an integer $m \geq 1$, an m-ary *relation* $\mathcal{R}$ is simply a subset $\mathcal{R} \subset \mathsf{Dom}^m \times \mathsf{Rng}^m$.

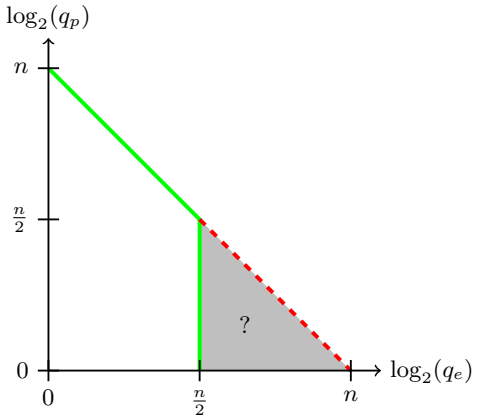

Fig. 5. The XRKA security of the (1-round) Even-Mansour construction with a nonlinear key-schedule. All parameters below the solid line are secure by Theorem 3, while all parameters above the dashed line (which merges with the solid line for $q_e \leq 2^{n/2}$) are insecure by the single-key attacks of [Dae91, DKS12]. The security for parameters between these two lines remains unknown.

It is well-known that it is impossible to rigorously define a notion of resistance to chosen-key attacks for block ciphers in the standard model (i.e., for block ciphers not relying on an underlying ideal primitive) without running into impossibility results similar to the one of [CGH98] about random oracles. However, it is possible to avoid such pitfalls in idealized models, as we explain now.

For this, we introduce the concept of evasive relation which, informally, refers to a relation such that it is hard for an algorithm with oracle access to an ideal cipher E to come with a tuple of inputs $(\alpha_1, \ldots, \alpha_m)$ such that $((\alpha_1, \ldots, \alpha_m), (E(\alpha_1), \ldots, E(\alpha_m)))$ satisfies this relation.

Definition 2 (Evasive Relation). *An m-ary relation $\mathcal{R}$ is said (q, ε)-evasive (with respect to an ideal cipher) if for any oracle Turing machine $\mathcal{M}$ making at most q oracle queries, one has*

$$\Pr\left[E \leftarrow_{\$} \mathsf{BC}(\kappa, n), (\alpha_1, \ldots, \alpha_m) \leftarrow \mathcal{M}^E : ((\alpha_1, \ldots, \alpha_m), (E(\alpha_1), \ldots, E(\alpha_m))) \in \mathcal{R}\right] \leq \varepsilon,$$

where the probability is taken over the random draw of E and the random coins of $\mathcal{M}$.

Example 1. Consider the problem of finding a preimage of zero for a compression function $f(k, x) := E(k, x) \oplus x$ built from a block cipher E in Davies-Meyer mode, i.e., finding a pair (k, x) such that $E(k, x) \oplus x = 0$. This corresponds to the unary relation $\mathcal{R} = \{((+, k, x), y) \in \mathsf{Dom} \times \mathsf{Rng} : x \oplus y = 0\}$. A celebrated result by Winternitz [Win84], generalized by Black *et al.* [BRS02], says that this relation is $(q, \mathcal{O}(q/2^n))$-evasive with respect to an ideal cipher. Similarly, the collision resistance of the Davies-Meyer mode [BRS02] can be recast as a binary $(q, \mathcal{O}(q^2/2^n))$-evasive relation for the underlying block cipher.

Definition 3 (Correlation Intractable Block Cipher). *Let $\mathcal{C}$ be a block cipher construction using (in a black-box way) an underlying primitive F, and let $\mathcal{R}$ be an m-ary relation. $\mathcal{C}^F$ is said to be* (q, ε)-correlation intractable *with respect to $\mathcal{R}$ if for any oracle Turing machine*

$\mathcal{M}$ making at most q oracle queries, one has

$$\Pr\left[(\alpha_1,\ldots,\alpha_m) \leftarrow \mathcal{M}^F : ((\alpha_1,\ldots,\alpha_m),(\mathcal{C}^F(\alpha_1),\ldots,\mathcal{C}^F(\alpha_m))) \in \mathcal{R}\right] \leq \varepsilon,$$

where the probability is taken over the random draw of F (in some well-understood set) and the random coins of $\mathcal{M}$.

Informally, a block cipher construction $\mathcal{C}^F$ can be deemed resistant to chosen-key attacks if for any (q,ε)-evasive relation $\mathcal{R}$, $\mathcal{C}^F$ is (q',ε')-correlation intractable with respect to $\mathcal{R}$ with $q' \simeq q$ and $\varepsilon' \simeq \varepsilon$. Note that our definitions above are information-theoretic, since later we will be able to prove information-theoretic security for the 4-round IEM cipher. There is no obstacle in providing corresponding computational definitions by taking the running time of the algorithms into account.

4.2 Sequential Indifferentiability

We define here the notion of *sequential indifferentiability* (*seq-indifferentiability* for short), introduced by [MPS12], which is a weakened variant of (full) indifferentiability as introduced by [MRH04], and then explain how it is related to correlation intractability. We use the definition of sequential indifferentiability given in [MPS12], tailored to the case of block ciphers.

We start with some definitions. Let $\mathcal{C}$ be a block cipher construction using in a black-box way an underlying primitive F. Let $\mathcal{D}$ be a distinguisher accessing a pair of oracles that we denote generically (E,F), which can be either the construction together with the underlying primitive F, i.e., $(\mathcal{C}^F,F)$, or $(E,\mathcal{S}^E)$ where E is an ideal cipher and $\mathcal{S}$ is an oracle Turing machine with oracle access to E called a *simulator*. We will refer informally to E as the *left* oracle and F as the *right* oracle. A distinguisher is said to be *sequential* if after its first query to its left (construction/ideal cipher) oracle, it does not query its right (primitive/simulator) oracle any more. Hence, such a distinguisher works in two phases: first it queries only its right oracle, and then only its left oracle (see Figure 6). We define the *total oracle query cost* of $\mathcal{D}$ as the total number of queries received by F (from $\mathcal{D}$ *or* $\mathcal{C}$) when $\mathcal{D}$ interacts with $(\mathcal{C}^F,F)$. In particular, if $\mathcal{C}$ makes c queries to F to answer any query it receives, and if $\mathcal{D}$ makes q_e queries to its left oracle and q_f queries to its right oracle, then the total oracle query cost of $\mathcal{D}$ is at most $q_f + cq_e$.

Definition 4 (Seq-indifferentiability). *Let $q,\sigma,t \in \mathbb{N}$ and $\varepsilon \in \mathbb{R}^+$. A block cipher construction $\mathcal{C}$ with black-box access to an ideal primitive F is said to be (q,σ,t,ε)-seq-indifferentiable from an ideal cipher if there exists an oracle algorithm $\mathcal{S}$ such that for any* sequential *distinguisher $\mathcal{D}$ of total oracle query cost at most q, $\mathcal{S}$ makes at most σ oracle queries, runs in time at most t, and one has*

$$\left|\Pr\left[\mathcal{D}^{E,\mathcal{S}^E}=1\right] - \Pr\left[\mathcal{D}^{\mathcal{C}^F,F}=1\right]\right| \leq \varepsilon,$$

where the first probability is taken over the random draw of the ideal cipher E and the random coins of $\mathcal{S}$, and the second probability is taken over the random draw of F (from some well understood set).

Note that this definition is information-theoretic (the distinguisher might be computationally unbounded), and demands the existence of a *universal* simulator (this is sometimes called

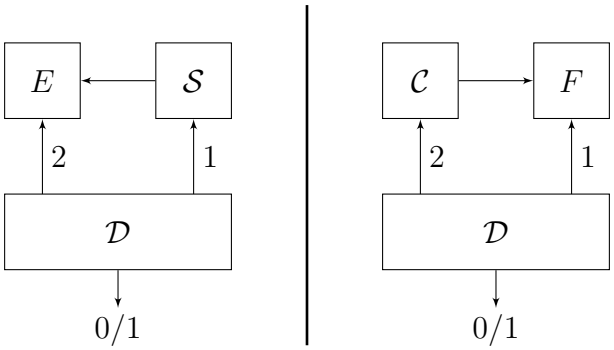

Fig. 6. The sequential indifferentiability notion. The numbers next to query arrows indicate in which order the distinguisher accesses both oracles. After its first query to the left oracle, the distinguisher cannot query the right oracle any more.

strong indifferentiability; when the simulator is allowed to depend on the distinguisher, this is called *weak* indifferentiability).

The usefulness of seq-indifferentiability in the context of CKAs comes from the following theorem (the proof is essentially similar to the proof of [MPS12, Theorem 3], but we make the relation between the various parameters explicit).

Theorem 4. *Let $\mathcal{C}$ be a block cipher construction using (in a black-box way) an underlying primitive F such that $\mathcal{C}$ makes at most c queries to F on any input. Assume that $\mathcal{C}^F$ is $(q + cm, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher. Then for any m-ary relation $\mathcal{R}$, if $\mathcal{R}$ is $(\sigma + m, \varepsilon_{\mathcal{R}})$-evasive with respect to an ideal cipher, then $\mathcal{C}^F$ is $(q, \varepsilon + \varepsilon_{\mathcal{R}})$-correlation intractable with respect to $\mathcal{R}$.*

Proof. Assume that there exists an m-ary relation $\mathcal{R}$ which is $(\sigma + m, \varepsilon_{\mathcal{R}})$-evasive but such that $\mathcal{C}^F$ is not $(q, \varepsilon + \varepsilon_{\mathcal{R}})$-correlation intractable with respect to $\mathcal{R}$. Then there exists an oracle machine $\mathcal{M}$ making at most q oracle queries such that $\mathcal{M}^F$ outputs with probability $\varepsilon' > \varepsilon_{\mathcal{R}} + \varepsilon$ a sequence $(\alpha_1, \ldots, \alpha_m)$ such that

$$((\alpha_1, \ldots, \alpha_m), (\mathcal{C}^F(\alpha_1), \ldots, \mathcal{C}^F(\alpha_m))) \in \mathcal{R}.$$

Consider the following sequential distinguisher $\mathcal{D}$ accessing a pair of oracles (E, F): it runs $\mathcal{M}$, answering $\mathcal{M}$'s oracle queries with its own oracle F, until $\mathcal{M}$ returns a tuple $(\alpha_1, \ldots, \alpha_m)$. $\mathcal{D}$ then makes oracle queries $E(\alpha_1), \ldots, E(\alpha_m)$ and checks[7] whether

$$((\alpha_1, \ldots, \alpha_m), (E(\alpha_1), \ldots, E(\alpha_m))) \in \mathcal{R}.$$

If this is the case it returns 1, otherwise it returns 0. Note that the total oracle query cost of $\mathcal{D}$ is at most $q + cm$.

When the distinguisher is interacting with $(\mathcal{C}^F, F)$, the probability that it returns 1 is exactly $\varepsilon' > \varepsilon_{\mathcal{R}} + \varepsilon$. On the other hand, when it interacts with $(E, \mathcal{S}^E)$, then the union of $\mathcal{D}$ and $\mathcal{S}$ is an oracle machine with oracle access to E making at most $\sigma + m$ oracle queries, so

[7] Note that we are working in the information-theoretic framework, so that the running time of $\mathcal{D}$ is irrelevant. In the computational framework, one should take into account the time necessary to recognize relation $\mathcal{R}$.

that, by definition of a $(\sigma + m, \varepsilon_{\mathcal{R}})$-evasive relation, $\mathcal{D}$ outputs 1 with probability at most $\varepsilon_{\mathcal{R}}$. Hence, the advantage of the distinguisher is $\varepsilon' - \varepsilon_{\mathcal{R}} > \varepsilon$, which contradicts the $(q + cm, \sigma, \varepsilon)$-seq-indifferentiability of $\mathcal{C}$. □

INTERPRETATION. Assuming c and m are constants which are negligible compared with q and σ, Theorem 4 can be paraphrased as follows: if $\mathcal{C}$ is $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher, and if a relation $\mathcal{R}$ cannot be found with probability better than $\varepsilon_{\mathcal{R}}$ with σ queries to an ideal cipher, then $\mathcal{R}$ cannot be found for $\mathcal{C}^F$ with probability better than $\varepsilon + \varepsilon_{\mathcal{R}}$ with q queries to F. (Note that the running time of the simulator is irrelevant here since we used an information-theoretic definition of correlation intractability.) Hence, seq-indifferentiability measures how much easier it is to find some relation $\mathcal{R}$ for a block cipher construction $\mathcal{C}^F$ than for an ideal cipher. In a sense, Theorem 4 can be seen as the analogue in the case of sequential indifferentiability of the composition theorem of [MRH04, RSS11] for full indifferentiability.

If one is only concerned with asymptotic security, then seq-indifferentiability implies correlation intractability in the following sense. Let $(\mathcal{C}_n^F)_{n\in\mathbb{N}}$ be a block cipher construction family indexed by a security parameter n. We simply say that $\mathcal{C}_n^F$ is seq-indifferentiable from an ideal cipher if for any $q \in \mathtt{poly}(n)$, $\mathcal{C}_n^F$ is $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher with $\sigma, t \in \mathtt{poly}(n)$ and $\varepsilon \in \mathtt{negl}(n)$. We simply say that $\mathcal{C}_n^F$ is correlation intractable if for any (q, ε)-evasive relation $\mathcal{R}$ (with respect to an ideal cipher) where $q \in \mathtt{poly}(n)$ and $\varepsilon \in \mathtt{negl}(n)$, $\mathcal{C}_n^F$ is (q', ε')-correlation intractable with respect to $\mathcal{R}$ for some $q' \in \mathtt{poly}(n)$ and $\varepsilon' \in \mathtt{negl}(n)$. Then a direct corollary of Theorem 4 is that if $\mathcal{C}_n^F$ is (asymptotically) seq-indifferentiable from an ideal cipher, then it is also (asymptotically) correlation intractable.

However, if we adopt the "concrete" security viewpoint, then the exact seq-indifferentiability parameters are important to quantify how well exactly the construction withstands chosen-key attacks. Consider Example 1 of preimage resistance of the Davies-Meyer compression function, which can be phrased as a $(q, \mathcal{O}(q/2^n))$-evasive relation $\mathcal{R}$ for the underlying (ideal) cipher. Assume that a block cipher construction $\mathcal{C}^F$ is $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher with, e.g., $\sigma = \mathcal{O}(q^2)$ and $\varepsilon = \mathcal{O}(q^2/2^n)$. Then Theorem 4 implies that $\mathcal{C}^F$ is $(q, \mathcal{O}(q^2/2^n))$-correlation intractable with respect to $\mathcal{R}$, or in other words, that the Davies-Meyer compression function based on $\mathcal{C}^F$ is $(q, \mathcal{O}(q^2/2^n))$-preimage resistant (in the ideal-F model). Hence, the quadratic query complexity of the simulator implies a security loss for correlation intractability. This motivates to look for block cipher constructions that are $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher with $\sigma = \mathcal{O}(q)$ and $\varepsilon = \mathcal{O}(q/2^n)$, which we leave for future work.

4.3 Proof of Sequential Indifferentiability for Four Rounds

FOUR ROUNDS ARE NECESSARY. We first recall that Lampe and Seurin gave an attack against full indifferentiability of the 3-round IEM cipher [LS13] (a different attack has been independently described by Andreeva *et al.* [ABD+13]). A closer look at their attack shows that their distinguisher is in fact sequential (we refer to [LS13] for a detailed description of the attack for reasons of space), so that the 3-round IEM cipher cannot even be seq-indifferentiable from an ideal cipher. Hence, at least four rounds are necessary (and, as we will see now, sufficient) to achieve seq-indifferentiability from an ideal cipher.

MAIN RESULT. We now state and prove the main result of this section regarding the seq-indifferentiability of the 4-round IEM cipher. The proof essentially follows the same lines as

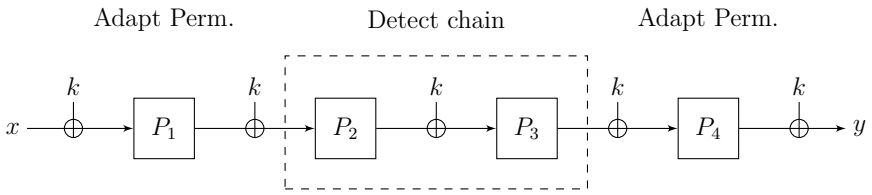

Fig. 7. The 4-round iterated Even-Mansour cipher with independent permutations and identical round keys. The detection and adaptations zones used by the simulator for proving seq-indifferentiability from an ideal cipher are also depicted.

the proof of full indifferentiability of [LS13] for twelve rounds, but is quite simpler since the simulator does not recurse when completing chains.

Theorem 5. *Let $N = 2^n$. For any integer q such that $q^2 \leq N/4$, the 4-round IEM construction (with independent permutations and identical round keys) is $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher with n-bit blocks and n-bit keys, with*

$$\sigma = q^2, \quad t = \mathcal{O}(q^2), \quad \text{and} \quad \varepsilon = \frac{68q^4}{N}.$$

Remark 1. It was shown in [MPS12] that for stateless ideal primitives (i.e., primitives whose answers do not depend on the order of the queries it receives), seq-indifferentiability implies public indifferentiability [YMO09, DRS09], a variant of indifferentiability where the simulator gets to know all queries of the distinguisher to E. Since an ideal cipher is stateless, Theorem 5 implies that the 4-round IEM construction is also publicly indifferentiable from an ideal cipher.

In order to prove this theorem, we will first define a simulator $\mathcal{S}$, then prove that it runs in polynomial time and makes a polynomial number of queries (Lemma 4), and finally prove that the two systems $\Sigma_1 = (E, \mathcal{S}^E)$ and $\Sigma_3 = (\mathsf{EM}^P, P)$ are indistinguishable, using an intermediate system Σ_2 that we will describe later (Lemmas 6 and 7).

Informal Description of the Simulator and Notation. We start with an informal description of the simulator (a formal description in pseudocode is given in Appendix A). The simulator offers an interface $\mathsf{Query}(i, \delta, w)$ to the distinguisher for querying the internal permutations, where $i \in \{1, \ldots, 4\}$ names the permutation, $\delta \in \{+, -\}$ indicates whether this a direct or inverse query, and $w \in \{0,1\}^n$ is the actual value queried. For each $i = 1, \ldots, 4$, the simulator internally maintains a table Π_i mapping entries $(\delta, w) \in \{+, -\} \times \{0,1\}^n$ to values $w' \in \{0,1\}^n$, initially undefined for all entries. We denote Π_i^+, resp. Π_i^-, the (time-dependent) sets of strings $w \in \{0,1\}^n$ such that $\Pi_i(+, w)$, resp. $\Pi_i(-, w)$, is defined. When the simulator receives a query (i, δ, w), it looks in table Π_i to see whether the corresponding answer $\Pi_i(\delta, w)$ is already defined. When this is the case, it outputs the answer and waits for the next query. Otherwise, it randomly draws an answer $w' \in \{0,1\}^n$ and defines $\Pi_i(\delta, w) := w'$ as well as the answer to the opposite query $\Pi_i(\bar{\delta}, w') := w$. In order to handily describe how the answer w' is drawn, we make the randomness used by the simulator explicit through a tuple of random permutations $P = (P_1, \ldots, P_4)$. As for the ideal cipher E, we formally let each P_i have a single interface, namely $P_i := \{+, -\} \times \{0,1\}^n \to \{0,1\}^n$, and for any $u, v \in \{0,1\}^n$, $P_i(+, u) = v \Leftrightarrow P_i(-, v) = u$. We assume that the tuple $(P_1, \ldots, P_4)$ is drawn uniformly at

random at the beginning of the experiment, but we note that $\mathcal{S}$ could equivalently lazily sample these permutations throughout its execution. Then w' is simply defined by the simulator as $w' := P_i(\delta, w)$. (For reasons that will become clear later, this is not equivalent to drawing w' uniformly from $\{0,1\}^n \setminus \Pi_i^{\bar{\delta}}$, see Remark 2.)

After this random choice of the answer w', and before returning it to the distinguisher, the simulator takes additional steps to ensure consistency with the ideal cipher E by running a *chain completion* mechanism. Namely, if the distinguisher called $\mathsf{Query}(i, \delta, w)$ with $i = 2$ or 3, the simulator completes all newly created "chains" (v_2, u_3), where $v_2 \in \Pi_2^-$ and $u_3 \in \Pi_3^+$ by executing a procedure $\mathsf{CompleteChain}(v_2, u_3, \ell)$, where ℓ indicates where the chain will be "adapted". For example, assume that the distinguisher called $\mathsf{Query}(2, +, u_2)$ and that the answer randomly chosen by the simulator was v_2 (or the backward counterpart, namely the distinguisher called $\mathsf{Query}(2, -, v_2)$ and the answer randomly chosen by the simulator was u_2). Then for each $u_3 \in \Pi_3^+$, the simulator computes the corresponding key $k := v_2 \oplus u_3$, and evaluates the IEM construction backward, letting $u_2 := \Pi_2(-, v_2)$ and $v_1 := u_2 \oplus k$, and forward, letting $v_3 := \Pi_3(+, u_3)$, $u_4 := v_3 \oplus k$, $v_4 := \Pi_4(+, u_4)$ (setting this value at random in case it was not in Π_4), $y := v_4 \oplus k$, $x := E(-, k, y)$ (hence making a query to E to "wrap around"), and $u_1 := x \oplus k$, until the corresponding input/output values (u_1, v_1) for the first permutation are defined. It then "adapts" (rather than setting randomly) table Π_1 by calling procedure $\mathsf{ForceVal}(u_1, v_1, 1)$ which sets $\Pi_1(+, u_1) := v_1$ and $\Pi_1(-, v_1) := u_1$ in order to ensure consistency of the simulated IEM construction with E. (A crucial point of the proof will be to show that this does not cause an overwrite, i.e., that these two values are undefined before the adaptation occurs.) In case the query was to $\mathsf{Query}(3, \cdot, \cdot)$, the behavior of the simulator is symmetric, namely adaptation of the chain takes place in table Π_4.

In all the following, we define the *size* of each table Π_i as $|\Pi_i| = \max\{|\Pi_i^+|, |\Pi_i^-|\}$. (Note that as long as no value is overwritten in the tables, $|\Pi_i^+| = |\Pi_i^-|$.)

Remark 2. As already noted, we could have easily described an equivalent simulator that lazily samples the random permutations $(P_1, \ldots, P_4)$ throughout its execution. However, we remark that this is not equivalent to replacing line (6) of the formal description of the simulator in Appendix A by $w' \leftarrow_{\$} \{0,1\}^n \setminus \Pi_i^{\bar{\delta}}$ for $i = 1$ and $i = 4$ since the simulator sometimes adapts the value of these tables, so that the tables Π_i and the permutations P_i will differ in general on the adapted entries.

Complexity of the Simulator. We start by proving that the simulator runs in polynomial time and makes a polynomial number of queries to the ideal cipher. More precisely, we have the following lemma.

Lemma 4. *Consider an execution of the simulator $\mathcal{S}^E$ where the simulator receives at most q queries in total. Then:*

(i) the size of Π_2 and Π_3 is at most q, and the size of Π_1 and Π_4 is at most $q^2 + q$;
(ii) the simulator executes $\mathsf{CompleteChain}$ *at most q^2 times, makes at most q^2 queries to E, and runs in time $\mathcal{O}(q^2)$.*

Proof. The size of Π_2, resp. Π_3, can only increase by one when the distinguisher makes a direct call to $\mathsf{Query}(2, \delta, w)$, resp. $\mathsf{Query}(3, \delta, w)$, so that the size of Π_2 and Π_3 is at most q. Procedure $\mathsf{CompleteChain}$ is called once for each pair $(v_2, u_3) \in \Pi_2^- \times \Pi_3^+$, hence at most q^2 times in

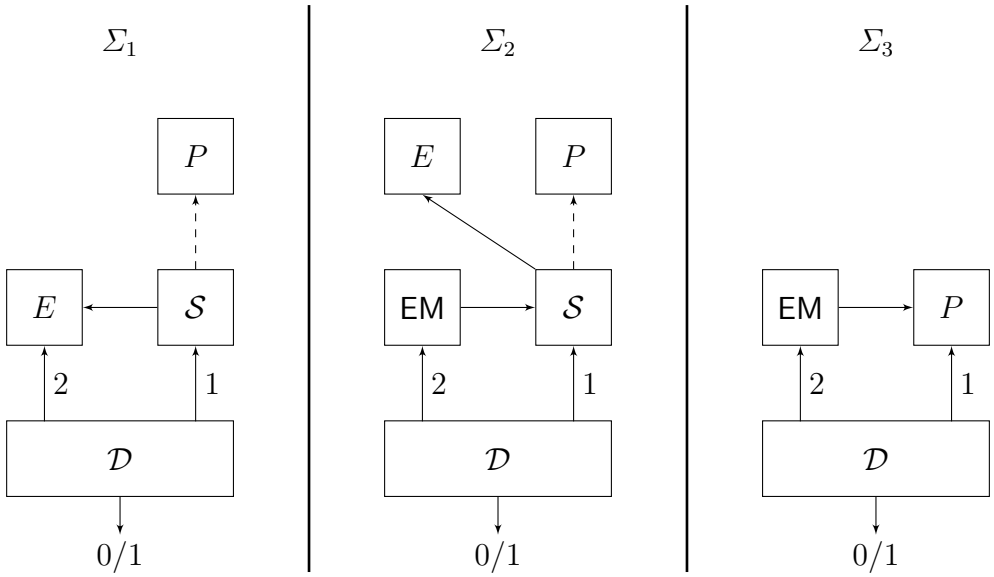

Fig. 8. Systems used in the seq-indifferentiability proof.

total. Since the simulator makes exactly one query to E per execution of CompleteChain, the total number of queries made by the simulator to E is at most q^2. The size of Π_1, resp. Π_4, can only increase by one when the distinguisher calls $\mathsf{Query}(1, \delta, w)$, resp. $\mathsf{Query}(4, \delta, w)$, or when CompleteChain is called, hence the size of Π_1 and Π_4 is at most $q^2 + q$. Clearly, the simulator running time is dominated by the executions of CompleteChain, hence the simulator runs in time $\mathcal{O}(q^2)$. □

Intermediate System. In all the following, we consider some fixed distinguisher $\mathcal{D}$, and assume that it is deterministic (this is *wlog* since we consider computationally unbounded distinguishers). We will denote $\mathcal{S}(E, P)$ rather than $\mathcal{S}(P)^E$ the simulator with oracle access to the ideal cipher E and using random permutations P as source of randomness. In order to prove the indistinguishability of the two systems $(E, \mathcal{S}(E, P))$ and (EM^P, P), we will use an intermediate system.[8] Let Σ_1 be the "ideal" world where the distinguisher interacts with $(E, \mathcal{S}(E, P))$. Note that all the randomness of system Σ_1 is captured by the pair (E, P). Let also Σ_3 be the "real" world where the distinguisher interacts with (EM^P, P). All the randomness of system Σ_3 is captured by P. In the intermediate system Σ_2, the distinguisher interacts with $(\mathsf{EM}^{\mathcal{S}(E,P)}, \mathcal{S}(E, P))$ (see Figure 8). In words, the right oracle is the simulator $\mathcal{S}(E, P)$ with oracle access to an ideal cipher E as in Σ_1, but now the left oracle is the 4-round IEM construction with oracle access to $\mathcal{S}(E, P)$ (rather than random permutations). As for Σ_1, all the randomness of system Σ_2 is captured by (E, P).

Transition from Σ_1 to Σ_2 and Good Executions. We first consider the transition from the first to the second system.

[8] We warn that this intermediate system is different from the one used in [LS13] to prove full indifferentiability of the 12-round IEM cipher, namely $(\mathsf{EM}^P, \mathcal{S}(\mathsf{EM}^P, P))$. It is in fact analogue to the one used by [MPS12] to prove the seq-indifferentiability of the 6-round Feistel construction.

Definition 5. *A pair (E, P) is said* good *if the simulator never overwrites an entry of its tables Π_i during an execution of $\mathcal{D}^{\Sigma_2(E,P)}$. Otherwise the pair is said* bad.

An overwrite may happen either during a random assignment (line (8) of the formal description of the simulator in Appendix A), or when adapting a chain (lines (48) and (49)). Note that whether a pair (E, P) is good or not depends on the distinguisher $\mathcal{D}$. We first upper bound the probability that a random pair (E, P) is bad.

Lemma 5. *Consider a distinguisher $\mathcal{D}$ of total oracle query cost at most q, with $q^2 \leq N/4$. Then a uniformly random pair (E, P), where $E \leftarrow_{\$} \mathsf{BC}(n, n)$ and $P \leftarrow_{\$} (\mathcal{P}_n)^4$, is bad (with respect to $\mathcal{D}$) with probability at most $\frac{16q^4}{N}$.*

Proof. First, note that the total number of queries received by the simulator in Σ_2 (either from $\mathcal{D}$ or from the construction EM) is exactly the total oracle query cost q of the distinguisher. Since entries in Π_2 and Π_3 are never adapted, they can never be overwritten either. Hence, we only need to consider the probability of an overwrite in Π_1 or Π_4. Let $\mathsf{BadRand}$ be the event that an overwrite occurs during a random assignment (i.e., at line (8)) and $\mathsf{BadAdapt}$ be the event that an overwrite occurs when adapting a chain (v_2, u_3) (i.e., at line (48) or (49)).

We first consider the probability of $\mathsf{BadRand}$. Consider a random assignment in Π_i, for $i = 1$ or 4, namely $\Pi_i(\delta, w) := w'$, $\Pi_i(\bar{\delta}, w') := w$, with w' randomly defined as $w' := P_i(\delta, w)$. By Lemma 4 (i), there are at most $q^2 + q$ random assignments in Π_1 and Π_4, so that w' is uniformly random in a set of size at least $N - (q^2 + q)$. Moreover, this random assignment cannot overwrite a value that was previously added during a random assignment, but only a value that was added by $\mathsf{ForceVal}$ (i.e., when adapting a chain), and by Lemma 4 (ii) there are at most q^2 such values. Hence, the probability that w' is equal to one of the at most q^2 values previously added in table Π_i by a call to $\mathsf{ForceVal}$ is at most $\frac{q^2}{N-q^2-q}$. Summing over the at most $q^2 + q$ random assignments in Π_1 and Π_4, we get

$$\Pr\left[\mathsf{BadRand}\right] \leq 2(q^2 + q) \times \frac{q^2}{N - q^2 - q} \leq \frac{8q^4}{N}. \tag{10}$$

We now consider the probability of $\mathsf{BadAdapt}$, conditioned on $\mathsf{BadRand}$ not happening. Let $\mathsf{BadAdapt}_i$ be the event that a value is overwritten by the i-th call to $\mathsf{ForceVal}$. We will upper bound the probability

$$\Pr\left[\mathsf{BadAdapt}_i \middle| \neg\mathsf{BadRand} \wedge \neg\mathsf{BadAdapt}_j, j = 1, \ldots, i-1\right].$$

Consider the i-th execution of $\mathsf{CompleteChain}(v_2, u_3, \ell)$, and assume that $\mathsf{BadRand}$ does not occur and $\mathsf{BadAdapt}_j$ does not occur for $1 \leq j \leq i - 1$. This means that no value was overwritten before this i-th call to $\mathsf{CompleteChain}$. For concreteness, suppose that this chain completion was triggered by a call to $\mathsf{Query}(2, \cdot, \cdot)$ from the distinguisher, so that $\ell = 1$ (the reasoning is symmetric for a call to $\mathsf{Query}(3, \cdot, \cdot)$ for which $\ell = 4$). The simulator will eventually call $\mathsf{ForceVal}(u_1, v_1, 1)$, and we must show that with high probability, the values $\Pi_1(+, u_1)$ and $\Pi_1(-, v_1)$ are undefined previously to this call. We first consider the case of v_1. This value is defined by the simulator by setting $k := v_2 \oplus u_3$ and $v_1 := u_2 \oplus k$, hence $v_1 = u_2 \oplus v_2 \oplus u_3$. Independently of the direction of the query of the distinguisher, and since there are at most q random assignments in Π_2, the value $u_2 \oplus v_2$ comes at random from a set

of size at least $N - q$ (if the distinguisher called $\mathsf{Query}(2, +, u_2)$ then v_2 is random, whereas if it called $\mathsf{Query}(2, -, v_2)$ then u_2 is random). Hence, the probability that v_1 is equal to one of the at most $q^2 + q$ values already in Π_1 is at most $\frac{q^2+q}{N-q}$. We now argue that $\Pi_1(+, u_1)$ is also undefined with high probability. For this, we show that the query $E(-, k, y)$ made by the simulator to wrap around when evaluating the IEM construction forward is fresh, i.e., it never made this query before nor received y as answer to a previous query $E(+, k, x)$. Assume that this does not hold. Then this means that such a query previously occurred when completing another chain (v_2', u_3'). But since we assumed that no value was overwritten in the tables before this call to $\mathsf{CompleteChain}(v_2, u_3, 1)$, it can easily be seen that this implies that $(v_2', u_3') = (v_2, u_3)$, which cannot be since the simulator completes any chain at most once by construction. This implies that the value x returned by E comes at random from a set of size at least $N - q^2$ (since by Lemma 4 the simulator makes at most q^2 queries to E), so that $u_1 := x \oplus k$ is equal to one of the at most $q^2 + q$ values already in table Π_1 with probability at most $\frac{q^2+q}{N-q^2}$. Hence, summing over the at most q^2 calls to $\mathsf{CompleteChain}$, we obtain

$$\begin{aligned}\Pr\left[\mathsf{BadAdapt}|\neg\mathsf{BadRand}\right] &\leq \sum_{i=1}^{q^2} \Pr\left[\mathsf{BadAdapt}_i \middle| \neg\mathsf{BadRand} \wedge \neg\mathsf{BadAdapt}_j, j = 1, \ldots, i-1\right] \\ &\leq q^2 \left(\frac{q^2+q}{N-q} + \frac{q^2+q}{N-q^2}\right) \leq \frac{8q^4}{N}. \end{aligned} \tag{11}$$

Combining (10) and (11) yields the result. □

Lemma 6. *For any distinguisher $\mathcal{D}$ of total oracle query cost at most q, one has*

$$\left|\Pr\left[\mathcal{D}^{\Sigma_1(E,P)} = 1\right] - \Pr\left[\mathcal{D}^{\Sigma_2(E,P)} = 1\right]\right| \leq \frac{16q^4}{N},$$

where both probabilities are taken over $E \leftarrow_\$ \mathsf{BC}(n, n), P \leftarrow_\$ (\mathcal{P}_n)^4$.

Proof. Recall that the distinguisher is sequential, i.e., it first queries only its right oracle (which for both Σ_1 and Σ_2 is $\mathcal{S}(E, P)$) and then only its left oracle (which is E in Σ_1 and $\mathsf{EM}^{\mathcal{S}(E,P)}$ in Σ_2). We show that for any good pair (E, P), the transcript of the interaction of $\mathcal{D}$ with $\Sigma_1(E, P)$ and $\Sigma_2(E, P)$ is *exactly* the same. This is clear for the transcript of the first phase of the interaction, i.e., for the queries of $\mathcal{D}$ to $\mathcal{S}$, since in both cases they are answered by $\mathcal{S}$ using the same pair (E, P).[9] For the second phase of the interaction (i.e., queries of $\mathcal{D}$ to its left oracle), it directly follows from the adaptation mechanism and the fact that the simulator never overwrites values in its tables Π_i that for any forward query of the distinguisher, $\mathsf{EM}^{\mathcal{S}(E,P)}(+, k, x) = E(+, k, x)$, and similarly for any backward query, $\mathsf{EM}^{\mathcal{S}(E,P)}(-, k, y) = E(-, k, y)$. Hence, the transcripts of the interaction of $\mathcal{D}$ with $\Sigma_1(E, P)$ and $\Sigma_2(E, P)$ are the same for any good pair (E, P). Consequently,

$$\left|\Pr\left[\mathcal{D}^{\Sigma_1(E,P)} = 1\right] - \Pr\left[\mathcal{D}^{\Sigma_2(E,P)} = 1\right]\right| \leq \Pr\left[(E, P) \text{ is bad}\right],$$

from which the result follows by Lemma 5. □

[9] Note that the fact that the distinguisher is sequential is used precisely here: for a non-sequential distinguisher, the behavior of the simulator would be different in Σ_1 and Σ_2 since in Σ_2 the simulator would receive queries from the IEM construction that it does not receive in Σ_1.

TRANSITION FROM Σ_2 TO Σ_3 AND RANDOMNESS MAPPING. We now consider the transition from the second to the third system, using a randomness mapping argument similar to the one of [HKT11, LS13]. For this, we define a map Λ mapping pairs (E, P) either to the special symbol $\perp$ when (E, P) is bad, or to a tuple of *partial permutations* $P' = (P'_1, \ldots, P'_4)$ when (E, P) is good. A partial permutation is a function $P'_i : \{+, -\} \times \{0,1\}^n \to \{0,1\}^n \cup \{*\}$ such that for all $u, v \in \{0,1\}^n$, $P'_i(+, u) = v \neq * \Leftrightarrow P'_i(-, v) = u \neq *$.

The map Λ is defined for good pairs (E, P) as follows: run $\mathcal{D}^{\Sigma_2(E,P)}$, and consider the tables Π_i of the simulator at the end of the execution; then fill all undefined entries of the Π_i's with the special symbol $*$. The result is exactly $\Lambda(E, P)$. Since for a good pair (E, P), the simulator never overwrites an entry in its tables, it follows that $\Lambda(E, P)$ is a tuple of partial permutations as just defined above. We say that a tuple of partial permutations $P' = (P'_1, \ldots, P'_4)$ is good if it has a good preimage by Λ. We say that a tuple of permutations $P = (P_1, \ldots, P_4)$ extends a tuple of partial permutations $P' = (P'_1, \ldots, P'_4)$, denoted $P \vdash P'$, if for each $1 \leq i \leq 4$, P_i and P'_i agree on all entries such that $P'_i(\delta, w) \neq *$.

Lemma 7. *For any distinguisher $\mathcal{D}$ of total oracle query cost at most q, one has*

$$\left|\Pr\left[\mathcal{D}^{\Sigma_2(E,P)} = 1\right] - \Pr\left[\mathcal{D}^{\Sigma_3(P)} = 1\right]\right| \leq \frac{52q^4}{N},$$

where the first probability is taken over $E \leftarrow_\$ \mathsf{BC}(n, n), P \leftarrow_\$ (\mathcal{P}_n)^4$, and the second over $P \leftarrow_\$ (\mathcal{P}_n)^4$.

Proof. Let

$$\varepsilon \stackrel{\text{def}}{=} \left|\Pr\left[\mathcal{D}^{\Sigma_2(E,P)} = 1\right] - \Pr\left[\mathcal{D}^{\Sigma_3(P)} = 1\right]\right|$$

and assume *w.l.o.g.* that $\Pr\left[\mathcal{D}^{\Sigma_2(E,P)} = 1\right] \geq \Pr\left[\mathcal{D}^{\Sigma_3(P)} = 1\right]$.

By definition of the randomness mapping, for any good tuple of partial permutations P', the outputs of $\mathcal{D}^{\Sigma_2(E,P)}$ and $\mathcal{D}^{\Sigma_3(P)}$ are equal for any pair (E, P) such that $\Lambda(E, P) = P'$ and any tuple of permutations P such that $P \vdash P'$. Let Θ_1 be the set of tuple of partial permutations P' such that $\mathcal{D}^{\Sigma_2(E,P)}$ outputs 1 for any pair (E, P) such that $\Lambda(E, P) = P'$. Then

$$\varepsilon \leq \Pr\left[(E, P) \text{ is bad}\right] + \sum_{P' \in \Theta_1} \Pr[\Lambda(E, P) = P'] - \sum_{P' \in \Theta_1} \Pr\left[P \vdash P'\right]. \tag{12}$$

Fix a good tuple of partial permutations $P' = (P'_1, \ldots, P'_4)$, and let $|P'_i| = |\{u \in \{0,1\}^n : P'_i(+, u) \neq *\}| = |\{v \in \{0,1\}^n : P'_i(-, v) \neq *\}|$. Then, clearly,

$$\Pr\left[P \leftarrow_\$ (\mathcal{P}_n)^4 : P \vdash P'\right] = \frac{1}{\prod_{i=1}^4 (N)_{|P'_i|}}.$$

Fix now any good preimage $(\tilde{E}, \tilde{P})$ of P', where $\tilde{P} = (\tilde{P}_1, \ldots, \tilde{P}_4)$, and let q_e and q_i $(1 \leq i \leq 4)$ be the number of queries made by the simulator respectively to $\tilde{E}$ and $\tilde{P}_i$ in the execution of $\mathcal{D}^{\Sigma_2(\tilde{E},\tilde{P})}$. One can check that for any pair (E, P), $\Lambda(E, P) = P'$ *iff* the transcript of the interaction of $\mathcal{S}$ with (E, P) in $\mathcal{D}^{\Sigma_2(E,P)}$ is the same as the transcript of the interaction of $\mathcal{S}$ with $(\tilde{E}, \tilde{P})$ in $\mathcal{D}^{\Sigma_2(\tilde{E},\tilde{P})}$. It follows that

$$\Pr\left[E \leftarrow_\$ \mathsf{BC}(n, n), P \leftarrow_\$ (\mathcal{P}_n)^4 : \Lambda(E, P) = P'\right] \leq \frac{1}{(N)_{q_e} \prod_{i=1}^4 (N)_{q_i}}.$$

(The exact value of this probability depend on the number of queries per key made to E, but clearly it is maximal when all q_e queries are made for the same key.) Moreover, since the number of executions of ForceVal made by the simulator (i.e., the number of chain adaptations) is equal to the number of queries made by the simulator to E, one has

$$\sum_{i=1}^{4} |P'_i| = q_e + \sum_{i=1}^{4} q_i \leq 2q^2 + 4q, \tag{13}$$

where the inequality follows by Lemma 4 (i) on the final size of the tables Π_i maintained by the simulator. Hence, we have

$$\begin{aligned}
\frac{\Pr[P \vdash P']}{\Pr[\Lambda(E,P) = P']} &= \frac{(N)_{q_e} \prod_{i=1}^{4} (N)_{q_i}}{\prod_{i=1}^{4} (N)_{|P'_i|}} \\
&\geq \underbrace{\frac{N^{q_e + \sum_{i=1}^{4} q_i}}{N^{\sum_{i=1}^{4} |P'_i|}}}_{=1 \text{ by } (13)} \times \prod_{j=1}^{q_e - 1} \left(1 - \frac{j}{N}\right) \prod_{i=1}^{4} \prod_{j=1}^{q_i - 1} \left(1 - \frac{j}{N}\right) \\
&\geq 1 - \frac{q_e^2 + \sum_{i=1}^{4} q_i^2}{N} \\
&\geq 1 - \frac{(2q^2 + 4q)^2}{N} && \text{by (13)} \\
&\geq 1 - \frac{36q^4}{N}.
\end{aligned}$$

Combining this lower bound with (12), we obtain

$$\begin{aligned}
\varepsilon &\leq \Pr[(E,P) \text{ is bad}] + \sum_{P' \in \Theta_1} \Pr[\Lambda(E,P) = P'] \left(1 - \frac{\Pr[P \vdash P']}{\Pr[\Lambda(E,P) = P']}\right) \\
&\leq \Pr[(E,P) \text{ is bad}] + \frac{36q^4}{N} \sum_{P' \in \Theta_1} \Pr[\Lambda(E,P) = P'] \\
&\leq \Pr[(E,P) \text{ is bad}] + \frac{36q^4}{N}.
\end{aligned}$$

The result follows from Lemma 5. □

Concluding. The proof of Theorem 5 directly follows by combining Lemmas 4, 6, and 7. As a corollary, we obtain from Theorem 4 that for any (q^2, ε)-evasive relation $\mathcal{R}$, the 4-round IEM cipher is $(q, \varepsilon + \mathcal{O}(q^4/2^n))$-correlation intractable with respect to $\mathcal{R}$. Using again Example 1, the Davies-Meyer compression function based on the 4-round IEM cipher is $(q, \mathcal{O}(q^4/2^n))$-preimage resistant in the Random Permutation Model. This is quite a weak security guarantee, and as already explained, this motivates the search for a block cipher construction (potentially the IEM cipher with a sufficient number of rounds) which is $(q, \sigma, t, \varepsilon)$-seq-indifferentiable from an ideal cipher with $\sigma = \mathcal{O}(q)$ and $\varepsilon = \mathcal{O}(q/2^n)$.

References

[ABD+13] Elena Andreeva, Andrey Bogdanov, Yevgeniy Dodis, Bart Mennink, and John P. Steinberger. On the Indifferentiability of Key-Alternating Ciphers. In Ran Canetti and Juan A. Garay, editors,

Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I), volume 8042 of *LNCS*, pages 531–550. Springer, 2013. Full version available at http://eprint.iacr.org/2013/061.

[ABM13] Elena Andreeva, Andrey Bogdanov, and Bart Mennink. Towards Understanding the Known-Key Security of Block Ciphers. In Shiho Moriai, editor, *Fast Software Encryption - FSE 2013*, volume 8424 of *LNCS*, pages 348–366. Springer, 2013.

[BC10] Mihir Bellare and David Cash. Pseudorandom Functions and Permutations Provably Secure against Related-Key Attacks. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010*, volume 6223 of *LNCS*, pages 666–684. Springer, 2010.

[BDK05] Eli Biham, Orr Dunkelman, and Nathan Keller. Related-Key Boomerang and Rectangle Attacks. In Ronald Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005*, volume 3494 of *LNCS*, pages 507–525. Springer, 2005.

[BK03] Mihir Bellare and Tadayoshi Kohno. A Theoretical Treatment of Related-Key Attacks: RKA-PRPs, RKA-PRFs, and Applications. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 491–506. Springer, 2003.

[BKL+12] Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations - (Extended Abstract). In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62. Springer, 2012.

[BKN09] Alex Biryukov, Dmitry Khovratovich, and Ivica Nikolić. Distinguisher and Related-Key Attack on the Full AES-256. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *LNCS*, pages 231–249. Springer, 2009.

[BRS02] John Black, Phillip Rogaway, and Thomas Shrimpton. Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *LNCS*, pages 320–335. Springer, 2002.

[CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The Random Oracle Methodology, Revisited (Preliminary Version). In *Symposium on Theory of Computing - STOC '98*, pages 209–218. ACM, 1998. Full version available at http://arxiv.org/abs/cs.CR/0010019.

[CLL+14] Shan Chen, Rodolphe Lampe, Jooyoung Lee, Yannick Seurin, and John P. Steinberger. Minimizing the Two-Round Even-Mansour Cipher. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 (Proceedings, Part I)*, volume 8616 of *LNCS*, pages 39–56. Springer, 2014. Full version available at http://eprint.iacr.org/2014/443.

[CS14] Shan Chen and John Steinberger. Tight Security Bounds for Key-Alternating Ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, 2014. Full version available at http://eprint.iacr.org/2013/222.

[Dae91] Joan Daemen. Limitations of the Even-Mansour Construction. In Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors, *Advances in Cryptology - ASIACRYPT '91*, volume 739 of *LNCS*, pages 495–498. Springer, 1991.

[DKS12] Orr Dunkelman, Nathan Keller, and Adi Shamir. Minimalism in Cryptography: The Even-Mansour Scheme Revisited. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 336–354. Springer, 2012.

[DRS09] Yevgeniy Dodis, Thomas Ristenpart, and Thomas Shrimpton. Salvaging Merkle-Damgård for Practical Applications. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 371–388. Springer, 2009.

[EM97] Shimon Even and Yishay Mansour. A Construction of a Cipher from a Single Pseudorandom Permutation. *Journal of Cryptology*, 10(3):151–162, 1997.

[FP15] Pooya Farshim and Gordon Procter. The Related-Key Security of Iterated Even-Mansour Ciphers. In *Fast Software Encryption - FSE 2015*, 2015. To appear. Full version available at http://eprint.iacr.org/2014/953.

[Gaz13] Peter Gazi. Plain versus Randomized Cascading-Based Key-Length Extension for Block Ciphers. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I)*, volume 8042 of *LNCS*, pages 551–570. Springer, 2013.

[GL10] David Goldenberg and Moses Liskov. On Related-Secret Pseudorandomness. In Daniele Micciancio, editor, *Theory of Cryptography - TCC 2010*, volume 5978 of *LNCS*, pages 255–272. Springer, 2010.

[HKT11] Thomas Holenstein, Robin Künzler, and Stefano Tessaro. The Equivalence of the Random Oracle Model and the Ideal Cipher Model, Revisited. In Lance Fortnow and Salil P. Vadhan, editors, *Symposium on Theory of Computing - STOC 2011*, pages 89–98. ACM, 2011. Full version available at http://arxiv.org/abs/1011.1264.

[IK04] Tetsu Iwata and Tadayoshi Kohno. New Security Proofs for the 3GPP Confidentiality and Integrity Algorithms. In Bimal K. Roy and Willi Meier, editors, *Fast Software Encryption - FSE 2004*, volume 3017 of *LNCS*, pages 427–445. Springer, 2004.

[JNP14] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - Proceedings, Part II*, volume 8874 of *LNCS*, pages 274–288. Springer, 2014.

[KPS13] Eike Kiltz, Krzysztof Pietrzak, and Mario Szegedy. Digital Signatures with Minimal Overhead from Indifferentiable Random Invertible Functions. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I)*, volume 8042 of *LNCS*, pages 571–588. Springer, 2013.

[KR07] Lars R. Knudsen and Vincent Rijmen. Known-Key Distinguishers for Some Block Ciphers. In Kaoru Kurosawa, editor, *Advances in Cryptology - ASIACRYPT 2007*, volume 4833 of *LNCS*, pages 315–324. Springer, 2007.

[LPS12] Rodolphe Lampe, Jacques Patarin, and Yannick Seurin. An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 278–295. Springer, 2012.

[LRW02] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable Block Ciphers. In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *LNCS*, pages 31–46. Springer, 2002.

[LS13] Rodolphe Lampe and Yannick Seurin. How to Construct an Ideal Cipher from a Small Set of Public Permutations. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 (Proceedings, Part I)*, volume 8269 of *LNCS*, pages 444–463. Springer, 2013. Full version available at http://eprint.iacr.org/2013/255.

[Luc04] Stefan Lucks. Ciphers Secure against Related-Key Attacks. In Bimal K. Roy and Willi Meier, editors, *Fast Software Encryption - FSE 2004*, volume 3017 of *LNCS*, pages 359–370. Springer, 2004.

[MPS12] Avradip Mandal, Jacques Patarin, and Yannick Seurin. On the Public Indifferentiability and Correlation Intractability of the 6-Round Feistel Construction. In Ronald Cramer, editor, *Theory of Cryptography Conference - TCC 2012*, volume 7194 of *LNCS*, pages 285–302. Springer, 2012. Full version available at http://eprint.iacr.org/2011/496.

[MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology. In Moni Naor, editor, *Theory of Cryptography Conference- TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, 2004.

[NK92] Kaisa Nyberg and Lars R. Knudsen. Provable Security Against Differential Cryptanalysis. In Ernest F. Brickell, editor, *Advances in Cryptology - CRYPTO '92*, volume 740 of *LNCS*, pages 566–574. Springer, 1992.

[Pat08] Jacques Patarin. The "Coefficients H" Technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography - SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, 2008.

[Rog06] Phillip Rogaway. Formalizing Human Ignorance. In Phong Q. Nguyen, editor, *Progress in Cryptology - VIETCRYPT 2006*, volume 4341 of *LNCS*, pages 211–228. Springer, 2006.

[RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with Composition: Limitations of the Indifferentiability Framework. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 487–506. Springer, 2011.

[Ste12] John Steinberger. Improved Security Bounds for Key-Alternating Ciphers via Hellinger Distance. IACR Cryptology ePrint Archive, Report 2012/481, 2012. Available at http://eprint.iacr.org/2012/481.

[Win84] Robert S. Winternitz. A Secure One-Way Hash Function Built from DES. In *IEEE Symposium on Security and Privacy*, pages 88–90, 1984.

[YMO09] Kazuki Yoneyama, Satoshi Miyagawa, and Kazuo Ohta. Leaky Random Oracle. *IEICE Transactions*, 92-A(8):1795–1807, 2009.

A Formal Description of the Simulator

Simulator $\mathcal{S}(P)$:
Variables:
  tables $\Pi_1, \ldots, \Pi_4$, initially empty

public procedure Query(i, δ, w):
  if $(\delta, w) \notin \Pi_i$ **then**
    $w' := P_i(\delta, w)$
    $\Pi_i(\delta, w) := w'$
    $\Pi_i(\bar{\delta}, w') := w$ \\ may overwrite an entry
    \\ complete newly created chains (v_2, u_3) if any
    if $i = 2$ **then**
      if $\delta = +$ **then** $v_2 := w'$ **else** $v_2 := w$
      forall $u_3 \in \Pi_3^+$ **do**
        CompleteChain$(v_2, u_3, 1)$
    else if $i = 3$ **then**
      if $\delta = +$ **then** $u_3 := w$ **else** $u_3 := w'$
      forall $v_2 \in \Pi_2^-$ **do**
        CompleteChain$(v_2, u_3, 4)$
  return $\Pi_i(\delta, w)$

private procedure CompleteChain(v_2, u_3, ℓ):
  $k := v_2 \oplus u_3$
  case $\ell = 1$:
    \\ evaluate the chain bw. up to v_1
    $u_2 := \Pi_2(-, v_2)$
    $v_1 := u_2 \oplus k$
    \\ evaluate the chain fw. up to u_1
    $v_3 := \Pi_3(+, u_3)$
    $u_4 := v_3 \oplus k$
    $v_4 := \text{Query}(4, +, u_4)$
    $y := v_4 \oplus k$
    $x := E(-, k, y)$
    $u_1 := x \oplus k$
    \\ adapt the chain
    ForceVal$(u_1, v_1, 1)$
  case $\ell = 4$:
    \\ evaluate the chain fw. up to u_4
    $v_3 := \Pi_3(+, u_3)$
    $u_4 := v_3 \oplus k$
    \\ evaluate the chain bw. up to v_4
    $u_2 := \Pi_2(-, v_2)$
    $v_1 := u_2 \oplus k$
    $u_1 := \text{Query}(1, -, v_1)$
    $x := u_1 \oplus k$
    $y := E(+, k, x)$
    $v_4 := y \oplus k$
    \\ adapt the chain
    ForceVal$(u_4, v_4, 4)$

private procedure ForceVal(u_i, v_i, i):
  $\Pi_i(+, u_i) := v_i$ \\ may overwrite an entry
  $\Pi_i(-, v_i) := u_i$ \\ may overwrite an entry

B Proof of Theorem 3

Theorem 3 is a direct consequence of Lemma 1 and Lemmas 8 and 9 proven below. Following the H-coefficient technique, we start by defining bad transcripts.

Definition 6. *Let $\tau = (\mathcal{Q}_E, \mathcal{Q}_P, k)$ be an attainable transcript. We say that τ is* bad *if*

$$k \in \mathsf{BadK} = \bigcup_{1 \le i \le 4} \mathsf{BadK}_i$$

where:

$$\begin{aligned}
k \in \mathsf{BadK}_1 &\Leftrightarrow \text{there exists } (\Delta, x, y) \in \mathcal{Q}_E \text{ and } (u, v) \in \mathcal{Q}_P \text{ such that } \gamma_0(k \oplus \Delta) = x \oplus u \\
k \in \mathsf{BadK}_2 &\Leftrightarrow \text{there exists } (\Delta, x, y) \in \mathcal{Q}_E \text{ and } (u, v) \in \mathcal{Q}_P \text{ such that } \gamma_1(k \oplus \Delta) = v \oplus y \\
k \in \mathsf{BadK}_3 &\Leftrightarrow \text{there exists } (\Delta, x, y), (\Delta', x', y') \in \mathcal{Q}_E \text{ with } \Delta \neq \Delta' \text{ such that} \\
&\qquad \gamma_0(k \oplus \Delta) \oplus \gamma_0(k \oplus \Delta') = x \oplus x' \\
k \in \mathsf{BadK}_4 &\Leftrightarrow \text{there exists } (\Delta, x, y), (\Delta', x', y') \in \mathcal{Q}_E \text{ with } \Delta \neq \Delta' \text{ such that} \\
&\qquad \gamma_1(k \oplus \Delta) \oplus \gamma_1(k \oplus \Delta') = y \oplus y'.
\end{aligned}$$

Otherwise, τ is said good. *We denote $\mathcal{T}_{\mathrm{bad}}$ the set of bad transcripts, and $\mathcal{T}_{\mathrm{good}} = \mathcal{T} \setminus \mathcal{T}_{\mathrm{bad}}$ the set of good transcripts.*

First, we upper bound the probability to get a bad transcript in the ideal world.

Lemma 8.

$$\Pr[T_{\mathrm{id}} \in \mathcal{T}_{\mathrm{bad}}] \le \frac{2q_e q_p}{N} + \frac{(\delta(\gamma_0) + \delta(\gamma_1))q_e^2}{2N}.$$

Proof. In the ideal world, the key k is drawn uniformly at random at the end of the query phase. Hence, we simply have to upper bound the size of BadK_i for $i = 1, \ldots, 4$, for any query transcript $(\mathcal{Q}_E, \mathcal{Q}_P)$. Fix any query transcript $(\mathcal{Q}_E, \mathcal{Q}_P)$. For any pair $(\Delta, x, y) \in \mathcal{Q}_E$, $(u, v) \in \mathcal{Q}_P$, there is exactly one key such that $\gamma_0(k \oplus \Delta) = x \oplus u$ since γ_0 is a permutation. Hence, we have $|\mathsf{BadK}_1| \le q_e q_p$. Similarly, $|\mathsf{BadK}_2| \le q_e q_p$. For any pair $(\Delta, x, y), (\Delta', x', y') \in \mathcal{Q}_E$ with $\Delta \neq \Delta'$, there are at most $\delta(\gamma_0)$ keys satisfying $\gamma_0(k \oplus \Delta) \oplus \gamma_0(k \oplus \Delta') = x \oplus x'$. Hence $|\mathsf{BadK}_3| \le \delta(\gamma_0)q_e^2/2$. Similarly, $|\mathsf{BadK}_4| \le \delta(\gamma_1)q_e^2/2$. Hence the result. □

Then, we consider good transcripts.

Lemma 9. *For any good transcript $\tau \in \mathcal{T}_{\mathrm{good}}$, one has $\Pr[T_{\mathrm{re}} = \tau] \ge \Pr[T_{\mathrm{id}} = \tau]$.*

Proof. If $\mathcal{T}_{\mathrm{good}} = \emptyset$, there is nothing to prove. Otherwise, fix a good transcript $\tau = (\mathcal{Q}_E, \mathcal{Q}_P, k)$. Let m denote the number of distinct offsets Δ appearing in the query transcript $\mathcal{Q}_E$, and for $i = 1, \ldots, m$, let q_i denote the number of queries with the i-th offset (ordering the offsets arbitrarily). Note that $\sum_{i=1}^m q_i = q_e$. Then, in the ideal world, we simply have

$$\begin{aligned}
\Pr[T_{\mathrm{id}} = \tau] &= \Pr[k' \leftarrow_{\$} \{0,1\}^n : k' = k] \times \Pr[P \leftarrow_{\$} \mathcal{P}_n : P \vdash \mathcal{Q}_P] \\
&\qquad \times \Pr[E \leftarrow_{\$} \mathsf{BC}(n,n) : (E,k) \vdash \mathcal{Q}_E] \\
&= \frac{1}{N} \cdot \frac{1}{(N)_{q_p}} \cdot \frac{1}{\prod_{i=1}^m (N)_{q_i}}.
\end{aligned}$$

On the other hand, since the transcript is good, all values $x \oplus \gamma_0(k \oplus \Delta)$ for (Δ, x, y) ranging over $\mathcal{Q}_E$ are pairwise distinct, and also distinct from all values u for $(u, v) \in \mathcal{Q}_P$, and similarly all values $y \oplus \gamma_1(k \oplus \Delta)$ for (Δ, x, y) ranging over $\mathcal{Q}_E$ are pairwise distinct, and also distinct from all values v for $(u, v) \in \mathcal{Q}_P$. Hence

$$\Pr[T_{\mathrm{re}} = \tau] = \frac{1}{N} \cdot \frac{1}{(N)_{q_p+q_e}}.$$

Since $(N)_{q_p+q_e} \leq (N)_{q_p}(N)_{q_e} \leq (N)_{q_p} \prod_{i=1}^{m} (N)_{q_i}$, we get the result. □

C Known-Key Attacks

Andreeva *et al.* [ABM13], in an attempt to formalize known-key attacks, have introduced the notion of known-key indifferentiability (KK-indifferentiability), and shown that the 1-round Even-Mansour cipher is KK-indifferentiable from an ideal cipher. KK-indifferentiability for a block cipher construction $\mathcal{C}^F$ is defined in a similar way as (full) indifferentiability, except that a random key k is drawn at the beginning of the security experiment, and the distinguisher is restricted to querying the construction $\mathcal{C}^F$ in the real world or the ideal cipher E in the ideal world with the key k. Moreover, in the ideal world, the simulator is given the key k as input.

We argue however that the notion of [ABM13] is slightly too restrictive to fully capture known-key attacks, because their definition involves only one single random key. If one tries to consider attacks with larger key arity, then the 1-round Even-Mansour cipher is *not* secure against known-key attacks. Consider the following simple example of a known-key attack against the 1-round Even-Mansour cipher (with identical round keys) involving two random keys. The adversary receives two random keys $k \neq k'$. It picks an arbitrary $x \in \{0,1\}^n$ and defines $x' = x \oplus k \oplus k'$. Let $y = \mathsf{EM}_k^P(x)$ and $y' = \mathsf{EM}_{k'}^P(x')$. Then one can easily check that $x \oplus x' = y \oplus y'$. Yet for an ideal cipher E, given two random keys $k \neq k'$, finding two pairs (x, y) and (x', y') such that $E_k(x) = y$, $E_{k'}(x') = y'$, and $x \oplus x' = y \oplus y'$ can be shown to be hard: more precisely, an adversary making at most q queries to E finds such pairs with probability $\mathcal{O}(\frac{q^2}{2^n})$. In other words, for the 1-round EM construction, the adversary can very easily find a binary relation which is $(q, \mathcal{O}(\frac{q^2}{2^n}))$-evasive with respect to an ideal cipher and involves the two "challenge" keys k, k'.

It is straightforward to extend the KK-indifferentiability definition of [ABM13] to handle larger key arity, by restricting the distinguisher to query its left oracle $(\mathcal{C}^F/E)$ on a set of at most m keys $k_1, \ldots, k_m$ randomly drawn at the beginning of the experiment. Then, for $m > 1$, the 1-round IEM cipher is not KK-indifferentiable from an ideal cipher under this definition, as shown by the attack outlined above.

Similarly, one could easily modify the definition of correlation intractability (cf. Definition 3) in order to better capture the known-key setting, by simply drawing m' random keys $k_1, \ldots, k_{m'}$ given as input to $\mathcal{M}^F$, and imposing to $\mathcal{M}$ that its output $(\alpha_1, \ldots, \alpha_m)$ only involves the "challenge" keys $k_1, \ldots, k_{m'}$.

We leave the study of these new notions to future work.