Strengthening the Known-Key Security Notion for Block Ciphers*

Benoît Cogliati and Yannick Seurin

University of Versailles, France
benoitcogliati@hotmail.fr

ANSSI, Paris, France
yannick.seurin@m4x.org

April 20, 2016

Abstract. We reconsider the formalization of known-key attacks against ideal primitive-based block ciphers. This was previously tackled by Andreeva, Bogdanov, and Mennink (FSE 2013), who introduced the notion of *known-key indifferentiability*. Our starting point is the observation, previously made by Cogliati and Seurin (EUROCRYPT 2015), that this notion, which considers only a single known key available to the attacker, is too weak in some settings to fully capture what one might expect from a block cipher informally deemed resistant to known-key attacks. Hence, we introduce a stronger variant of known-key indifferentiability, where the adversary is given *multiple* known keys to "play" with, the informal goal being that the block cipher construction must behave as an independent random permutation for each of these known keys. Our main result is that the 9-round iterated Even-Mansour construction (with the trivial key-schedule, i.e., the same round key xored between permutations) achieves our new "multiple" known-keys indifferentiability notion, which contrasts with the previous result of Andreeva *et al.* that one single round is sufficient when only a single known key is considered. We also show that the 3-round iterated Even-Mansour construction achieves the weaker notion of multiple known-keys *sequential* indifferentiability, which implies in particular that it is *correlation intractable* with respect to relations involving any (polynomial) number of known keys.

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

1 Introduction

Background on Known-Key Attacks. Informally, a known-key attack against a block cipher E consists in the following: the adversary is given a key k from the key space of E, and must find a "non-trivial" property of the permutation E_k associated with k faster than what it would cost given only black-box access to a truly random permutation. An example of such a non-trivial property would be

a plaintext/ciphertext pair (x, y) under the key k such that, say, the first half of x and the first half of y seen as bit strings are both zero (for a random permutation P over n-bit strings, it is easy to see that this requires roughly $2^{n/2}$ queries to P). Known-key attacks against block ciphers were first introduced by Knudsen and Rijmen [KR07], who exhibited such attacks against a reduced-round version of AES and against certain kinds of Feistel ciphers. These attacks were extended in a number of follow-up papers, e.g. [MPP09, GP10, NPSS10, SY11, Gil14].

Even though the informal idea underlying known-key security might intuitively seem clear (given a key k, the permutation E_k associated with k must "look random"), how to put known-key attacks on theoretical sound grounds has remained elusive. Indeed, any attempt to rigorously formalize what is a known-attack against a fixed block cipher runs into impossibility results similar to those undermining a sound definition of what a "good" hash function should be [CGH98]. In particular, seeing a block cipher as a family of permutations indexed by the key, the fact that the key-length is similar to the input-length of the permutations (i.e., the block-length of the block cipher) leads to the following "diagonal" problem: consider the set of pairs $(k, E_k(k))$ for k ranging over the key space (we assume that the block-length and the key-length are equal for ease of exposition); then it is hard, given oracle access to a random permutation, to find an input/output pair in this set, whereas given any key k for E it is very easy to find an input/output pair for E_k in this set.

A way to circumvent these impossibilities is to consider block cipher constructions based on some ideal primitive (for example, a Feistel cipher based on public random round functions or (iterated) Even-Mansour ciphers based on public permutations). In that case, even though the adversary is given the known key, it only has oracle access to the underlying primitive, which effectively acts as an (exponentially long) seed indexing the permutation associated with the key. A first step towards formalizing known-key attacks for ideal primitive-based block ciphers was taken by Andreeva, Bogdanov, and Mennink (ABM) [ABM13] through what they called *known-key indifferentiability* (KK-indifferentiability for short), a variant of the standard indifferentiability notion [MRH04]. A block cipher construction $\mathcal{C}^F$ from some underlying primitive F is said indifferentiable from an ideal cipher E if there exists an efficient simulator $\mathcal{S}$ with black box access to E such that the two pairs of oracles $(\mathcal{C}^F, F)$ and $(E, \mathcal{S}^E)$ are indistinguishable. Hence the simulator must make E "look like" $\mathcal{C}^F$ by returning answers that are coherent with the distinguisher's queries to E (without, in general, knowing these E-queries) and that are statistically close to answers of a real F oracle.

The KK-indifferentiability notion of ABM modifies the security experiment as follows: a key k is drawn at random and made available to the distinguisher and the simulator; the distinguisher is then allowed to query its left oracle (construction/ideal cipher) *only for this specific key* k. Hence the simulator's job is somehow made simpler since it has a "hint" about which queries the distinguisher can make to its left oracle. Note that in the ideal (simulated) world, the distinguisher effectively has access to a single random permutation (since an ideal cipher behaves as an independent random permutation for each key).

Hence this KK-indifferentiability notion intuitively captures the requirement that for each key k, the block cipher construction $\mathcal{C}^F$ must "look like" a random permutation. In contrast, the standard indifferentiability notion is related with *chosen-key* attacks, since the distinguisher is allowed to freely choose the keys it examines.

Shortcoming of the ABM Security Notion. The starting point of this paper is an observation, previously made by Cogliati and Seurin (Appendix C of the full version of [CS15]) that the ABM security notion might be too restrictive in some situations because it considers *one single* known-key. This might be problematic in some cryptosystems where intuitively resistance to known-key attacks should be sufficient to provide security, but where the ABM security notion fails because the cryptosystem uses *multiple* known keys. Think for example of the permutation-based hashed functions by Rogaway and Steinberger [RS08a, RS08b]: these constructions are based on a few (typically 3 to 6) public permutations, which would typically be instantiated by a block cipher used with distinct publicly known keys. A crucial requirement for the security proof of these constructions to hold (in the ideal permutation model) is that the permutations are independent. Since this is not ensured by the ABM security notion, it is not applicable here, even though one would like to say that a block cipher which is secure against known-key attacks can safely be used in the Rogaway-Steinberger constructions. (Jumping ahead, our new KK-indifferentiability notion will be sufficient to safely instantiate the block cipher in the same constructions.)

To better emphasize this gap between a single known-key notion and a multiple known-key notion, consider the case of the 1-round Even-Mansour (EM) [EM97, DKS12] construction based on a permutation P on $\{0,1\}^n$, which maps a key $k \in \{0,1\}^n$ and a plaintext $x \in \{0,1\}^n$ to the ciphertext defined as

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

ABM showed that when the permutation P is ideal, this construction is KK-indifferentiable from an ideal cipher in the single known-key setting. However, if the adversary is given any pair of distinct keys (k_1, k_2), it can pick any $x_1 \in \{0,1\}^n$, define $x_2 = x_1 \oplus k_1 \oplus k_2$, and compute $y_1 = \mathsf{EM}^P_{k_1}(x_1)$ and $y_2 = \mathsf{EM}^P_{k_2}(x_2)$. Then one can easily check that $x_1 \oplus x_2 = y_1 \oplus y_2$. Yet for an ideal cipher E, given two distinct keys $k_1 \neq k_2$, finding two pairs (x_1, y_1) and (x_2, y_2) such that $E_{k_1}(x_1) = y_1$, $E_{k_2}(x_2) = y_2$, and $x_1 \oplus x_2 = y_1 \oplus y_2$ can be shown to be hard: more precisely, an adversary making at most q queries to E can find such pairs with probability at most $\mathcal{O}(\frac{q^2}{2^n})$. In other words, the permutations associated with distinct keys for the 1-round EM construction do not "behave" independently.

Our Contribution. Our first contribution is definitional: in order to remedy the limitation that we just pointed out, we extend and strengthen the known-key security definition of [ABM13], by allowing the distinguisher to be given multiple known keys. Our new notion is parameterized by an integer μ, the number of known keys that the adversary is given. For $\mu = 1$, one recovers the ABM

definition. If one lets $\mu = |\mathcal{K}|$, where $\mathcal{K}$ is the key space of the block cipher, one recovers the standard indifferentiability notion. In fact, our KK-indifferentiability notion will emerge as a special case of a more general notion that we name *restricted-input*-indifferentiability, which might be of independent interest. We also formulate our KK-indifferentiability notion in a "worst-case" fashion (it must hold for *any* subset of keys of size μ), whereas the ABM notion was in the "average-case" style (the known key being randomly drawn). In addition, we define a weaker "sequential" variant [MPS12, CS15] of our new μ-KK-indifferentiability notion, called μ-KK-seq-indifferentiability, where the adversary must query its two oracles in a specific order. This notion is useful since it implies the weaker notion of correlation intractability.

Our second contribution is about constructions: we show that KK-indifferentiability is a meaningful notion by proving that the iterated Even-Mansour (IEM) construction with nine rounds is μ-KK-indifferentiable from an ideal cipher for any $\mu = \mathtt{poly}(n)$ (where n is a security parameter indexing the construction), which contrasts with the fact that one round is sufficient when considering one single known-key, and also with the best number of rounds known to be sufficient to achieve full indifferentiability from an ideal cipher, namely twelve [LS13]. We also show that three rounds are necessary and sufficient to achieve the weaker μ-KK-seq-indifferentiability notion, which again contrast with the fact that four rounds are necessary and sufficient to achieve (full) seq-indifferentiability from an ideal cipher [CS15]. See Table 1 for a summary of known results on the IEM construction.

More Related Work. A number of papers have studied the indifferentiability of variants of the IEM construction. In particular, Andreeva *et al.* [ABD+13] have studied the case where the key-schedule is modeled as a random oracle, and Guo and Lin have studied the case of Even-Mansour ciphers with two interleaved keys [GL15a] and of key-alternating Feistel ciphers [GL15b].

Organization. We start with some general definitions in Section 2. Then we define precisely our strengthened KK-indifferentiability notion (as well as the more general notion of *restricted-input*-indifferentiability, of which KK-indifferentiability is a special case) in Section 3. In Section 4, we give a known-key attack (using two known keys) against the 2-round IEM construction. Finally, we prove that the 3-round, resp. 9-round, IEM construction achieves μ-KK-seq-indifferentiability, resp. μ-KK-indifferentiability, in Sections 5 and 6.

2 Preliminaries

General Notation. In all the following, we fix an integer $n \geq 1$ and denote $N = 2^n$. Given a non-empty set $\mathcal{M}$, the set of all permutations of $\mathcal{M}$ will be denoted $\mathsf{Perm}(\mathcal{M})$. We simply denote $\mathsf{Perm}(n)$ the set of all permutations over $\{0,1\}^n$. A block cipher with key space $\mathcal{K}$ and message space $\mathcal{M}$ is a mapping $E : \mathcal{K} \times \mathcal{M} \to \mathcal{M}$ such that for any key $k \in \mathcal{K}$, $x \mapsto E(k,x)$ is a permutation.

Table 1. Summary of provable security results for the iterated Even-Mansour cipher with independent inner permutations and the trivial key-schedule. The first two notions are secret-key notions, the other ones are indifferentiability-based.

Sec. notion	# rounds	Sec. bound	Sim. complexity (query / time)	Ref.
Single-key (pseudorandomness)	1	$q^2/2^n$	—	[EM97, DKS12]
	2	$q^{3/2}/2^n$	—	[CLL+14]
XOR Related-Key	3	$q^2/2^n$	—	[CS15, FP15]
1-KK-indiff.	1	0	q / q	[ABM13]
μ-KK-Seq-indiff., $\mu > 1$	3	$\mu^2 q^2/2^n$	μq / μq	this paper
Full Seq-indiff.	4	$q^4/2^n$	q^2 / q^2	[CS15]
μ-KK-indiff., $\mu > 1$	9	$\mu^6 q^6/2^n$	$\mu^2 q$ / $\mu^2 q$	this paper
Full indiff.	12	$q^{12}/2^n$	q^4 / q^6	[LS13]

We interchangeably use the notations $E(k,x)$ and $E_k(x)$. We denote $\mathsf{BC}(\mathcal{K},\mathcal{M})$ the set of all block ciphers with key space $\mathcal{K}$ and message space $\mathcal{M}$, and $\mathsf{BC}(n,n)$ the set of block ciphers with key space and message space $\{0,1\}^n$. For integers $1 \le s \le t$, we will write $(t)_s = t(t-1)\cdots(t-s+1)$ and $(t)_0 = 1$ by convention.

Ideal Primitives. An *ideal primitive* F is a triplet $(\mathsf{F.Dom}, \mathsf{F.Rng}, \mathsf{F.Inst})$: the domain $\mathsf{F.Dom}$ and the range $\mathsf{F.Rng}$ are two non-empty sets, and the instance space $\mathsf{F.Inst}$ is a set of functions $F : \mathsf{F.Dom} \to \mathsf{F.Rng}$.

The two main ideal primitives we will be interested in are ideal permutations and ideal ciphers. Given a non-empty set $\mathcal{M}$, the ideal permutation P over $\mathcal{M}$ is defined as follows. Let $\mathsf{P.Dom} = \{+,-\} \times \mathcal{M}$ and $\mathsf{P.Rng} = \mathcal{M}$, and define

$$\mathsf{P.Inst} \stackrel{\text{def}}{=} \left\{P : \exists \pi \in \mathsf{Perm}(\mathcal{M}), P(+,x) = \pi(x) \text{ and } P(-,y) = \pi^{-1}(y)\right\}.$$

Clearly, there is a one-to-one correspondence between $\mathsf{P.Inst}$ and $\mathsf{Perm}(\mathcal{M})$.

Similarly, given two non-empty sets $\mathcal{K}$ and $\mathcal{M}$, the ideal cipher with key space $\mathcal{K}$ and message space $\mathcal{M}$ is defined as follows. Let $\mathsf{E.Dom} = \{+,-\} \times \mathcal{K} \times \mathcal{M}$, $\mathsf{E.Rng} = \mathcal{M}$, and define

$$\mathsf{E.Inst} \stackrel{\text{def}}{=} \left\{E : \exists \eta \in \mathsf{BC}(\mathcal{K},\mathcal{M}), E(+,k,x) = \eta_k(x) \text{ and } E(-,k,y) = \eta_k^{-1}(y)\right\}.$$

Again, there is a one-to-one correspondence between $\mathsf{E.Inst}$ and $\mathsf{BC}(\mathcal{K},\mathcal{M})$.

The Iterated Even-Mansour Cipher. Fix integers $n, r \ge 1$. Let $\mathbf{f} = (f_0, \ldots, f_r)$ be a $(r+1)$-tuple of permutations of $\{0,1\}^n$. The r-round iterated Even-Mansour construction $\mathsf{EM}[n,r,\mathbf{f}]$ specifies, from any r-tuple $\mathbf{P} =$

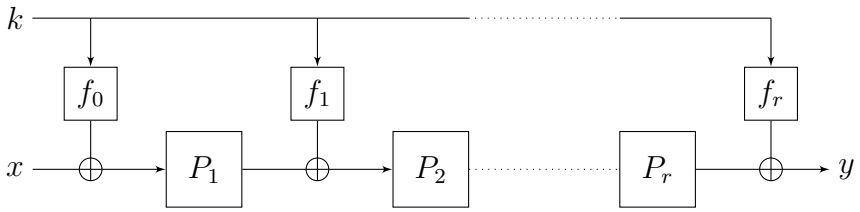

Fig. 1. The r-round iterated Even-Mansour cipher.

$(P_1, \ldots, P_r)$ of permutations of $\{0,1\}^n$, a block cipher with n-bit keys and n-bit messages, simply denoted $\mathsf{EM}^{\mathbf{P}}$ in all the following (parameters $[n, r, \mathbf{f}]$ 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 Fig. 1):

$$\mathsf{EM}^{\mathbf{P}}(k,x) = f_r(k) \oplus P_r(f_{r-1}(k) \oplus P_{r-1}(\cdots P_2(f_1(k) \oplus P_1(f_0(k) \oplus x))\cdots)).$$

We say that the key-schedule is *trivial* when all f_i's are the identity.

While the pseudorandomness of the IEM cipher was mostly studied with *independent* round keys [BKL+12, LPS12, CS14] (with the notable exception of [CLL+14]), it is well known that independent round keys cannot, in general, provide any security in the setting where the adversary has some control over the master key (related-, known-, or chosen-key attacks) [LS13]. Hence, in this paper, we focus on the case where the round keys are derived from an n-bit master key (actually, all our results deal with the case of the trivial key-schedule).

3 Restricted-Input Indifferentiability and Variants

We introduce the notion of restricted-input indifferentiability (*RI-indifferentiability*), and explain how known-key indifferentiability is a special case of it. Let E and F be two ideal primitives.[1] A *construction* implementing E from F is a deterministic algorithm $\mathcal{C}$ with oracle access to an instance F of F, which we denote $\mathcal{C}^F$, such that for any $F \in \mathsf{F.Inst}$, $\mathcal{C}^F \in \mathsf{E.Inst}$. A *simulator* for F is a randomized algorithm with oracle access to an instance E of E, which we denote $\mathcal{S}^E$, such that for any $E \in \mathsf{E.Inst}$, $\mathcal{S}^E : \mathsf{F.Dom} \to \mathsf{F.Rng}$. A distinguisher $\mathcal{D}$ is a deterministic[2] algorithm with oracle access to two oracles, the first one with signature $\mathsf{E.Dom} \to \mathsf{E.Rng}$, the second one with signature $\mathsf{F.Dom} \to \mathsf{F.Rng}$, and which returns a bit b, which we denote $\mathcal{D}(\mathcal{O}_1, \mathcal{O}_2) = b$. We will call $\mathcal{O}_1$ the *left* oracle and $\mathcal{O}_2$ the *right* oracle. Following [MPS12], we define the *total oracle query cost* of $\mathcal{D}$ as the maximum, over $F \in \mathsf{F.Inst}$, of the total number of queries

[1] This might be any ideal primitives, in particular E might not be an ideal cipher.
[2] Since we will consider computationally unbounded distinguishers, this is without loss of generality.

received by F (from $\mathcal{D}$ or $\mathcal{C}$) when $\mathcal{D}$ interacts with $(\mathcal{C}^F, F)$. The indifferentiability advantage of $\mathcal{D}$ against $(\mathcal{C}, \mathcal{S})$ is defined by

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathcal{C},\mathcal{S}}(\mathcal{D}) = \Big| \Pr\left[E \leftarrow_{\$} \mathsf{E.Inst} : \mathcal{D}(E, \mathcal{S}^E) = 1\right] - \Pr\left[F \leftarrow_{\$} \mathsf{F.Inst} : \mathcal{D}(\mathcal{C}^F, F) = 1\right] \Big|. \quad (1)$$

(Note that the first probability is also taken over the randomness of $\mathcal{S}$).

For any subset of X of $\mathsf{E.Dom}$, $\mathcal{D}$ is said X-restricted if it only makes queries to its left oracle (E or $\mathcal{C}^F$) from the set X.

Definition 1 (Restricted-Input Indifferentiability). *Let* E *and* F *be two ideal primitives and* $\mathcal{C}$ *be a construction implementing* E *from* F*. Let* $q, \sigma, t \in \mathbb{N}$ *and* $\varepsilon \in \mathbb{R}^+$*. Let* $\mathcal{X}$ *be a family of subsets of* $\mathsf{E.Dom}$*. Construction* $\mathcal{C}$ *is said* $(\mathcal{X}, q, \sigma, t, \varepsilon)$*-RI-indifferentiable from* E *if for any* $X \in \mathcal{X}$*, there exists a simulator* $\mathcal{S}$ *such that for any* X*-restricted 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*

$$\mathbf{Adv}^{\mathrm{indiff}}_{\mathcal{C},\mathcal{S}}(\mathcal{D}) \leq \varepsilon.$$

Informally, we simply say that $\mathcal{C}$ is $\mathcal{X}$-RI-indifferentiable from E if it is $(\mathcal{X}, q, \sigma, t, \varepsilon)$-RI-indifferentiable for "reasonable" values of σ, t, and ε expressed as functions of q (in particular, when $\mathcal{C}$ is indexed by some security parameter $n \in \mathbb{N}$, if $\sigma, t \in \mathtt{poly}(n)$ and $\varepsilon \in \mathtt{negl}(n)$ for any $q \in \mathtt{poly}(n)$).

As is standard in works on indifferentiability, this definition is information-theoretic, i.e., the distinguisher is allowed to be computationally unbounded (this is sometimes called *statistical indifferentiability*), and demands the existence of a *universal* simulator which does not depend on the distinguisher (this is sometimes called *strong* indifferentiability; when the simulator is allowed to depend on the distinguisher, this is called *weak* indifferentiability).

Note also the following points:

- by letting $\mathcal{X} = \{\mathsf{E.Dom}\}$ in the definition above, one recovers the standard definition of indifferentiability [MRH04];
- when $\mathcal{X} = \{X\}$ is reduced to a single subset of $\mathsf{E.Dom}$, the definition is equivalent to the standard definition of indifferentiability of the restriction of $\mathcal{C}^F$ to X from the restriction of E to X; hence this definition is only "new" when considering at least two distinct subsets X and X' such that $X \not\subseteq X'$ and $X' \not\subseteq X$ (since a X-restricted distinguisher is also a X'-restricted distinguisher when $X \subseteq X'$), and can be equivalently rephrased as the indifferentiability of the family of restrictions of $\mathcal{C}$ to sets in $\mathcal{X}$, with a uniform upper bound on the simulator's complexity and the distinguisher's advantage;
- the simulator is allowed to depend on the specific set $X \in \mathcal{X}$ considered;
- the upper bound on the advantage of the distinguisher must hold for any $X \in \mathcal{X}$ (not, say, on average on the random draw of X from $\mathcal{X}$).

The RI version of indifferentiability can be combined with other flavors of indifferentiability, in particular with public indifferentiability [DRS09, YMO09] and sequential indifferentiability [MPS12, CS15]. Let us elaborate for the case of sequential indifferentiability. A distinguisher is called *sequential* if after its first query to its left ($\mathsf{E}/\mathcal{C}^F$) oracle, it does not make any query to its right ($\mathcal{S}^E/F$) oracle any more. In other words, it works in two phases: first it only queries its right oracle, and then only its left oracle. Then we can define *RI-seq-indifferentiability* exactly as in Definition 1, except that we quantify over X-restricted *sequential* distinguishers only. (Hence this is a weaker definition since for each subset $X \in \mathcal{X}$, the simulator has to be effective only against a smaller class of distinguishers, namely sequential ones.)

Composition Theorem. The meaningfulness of the indifferentiability notion comes from the following composition theorem [MRH04]: if a cryptosystem is proven secure when implemented with ideal primitive E, then it remains provably secure when E is replaced with $\mathcal{C}$ based on ideal primitive F, assuming $\mathcal{C}$ is indifferentiable from E. (For this theorem to hold, the security of the cryptosystem must be defined with respect to a class of adversaries which "supports" the simulator used to prove that $\mathcal{C}$ is indifferentiable from E [RSS11, DGHM13].) This theorem straightforwardly translates to $\mathcal{X}$-RI-indifferentiability as follows: if a cryptosystem is proven secure when implemented with ideal primitive E and *if for any adversary $\mathcal{A}$, there is $X \in \mathcal{X}$ such that the challenger of the security game only queries E on inputs $x \in X$ when interacting with $\mathcal{A}$*, then it remains provably secure when E is replaced with $\mathcal{C}$ based on ideal primitive F, assuming $\mathcal{C}$ is $\mathcal{X}$-RI-indifferentiable from E. The short proof is as follows: denote Γ the challenger for the security game, which has access to an instance of E, and fix an adversary $\mathcal{A}$ against the cryptosystem implemented with $\mathcal{C}^F$ (hence $\mathcal{A}$ has oracle access to the instance F of the ideal primitive F); see the combination of Γ and $\mathcal{A}$ as a single X-restricted distinguisher $\mathcal{D}$; by the $\mathcal{X}$-RI-indifferentiability assumption, there is a simulator $\mathcal{S}$ such that $(\mathcal{C}^F, F)$ cannot be distinguished from $(E, \mathcal{S}^E)$; then the combination of $\mathcal{A}$ and $\mathcal{S}$ constitutes an attacker against the cryptosystem implemented with E, and the winning probability of $\mathcal{A}'$ is small by the assumption that the cryptosystem is secure when implemented with E; hence the winning probability of $\mathcal{A}$ is small as well.

Known-Key Indifferentiability. We now explain how to formalize resistance to known-key attacks using RI-indifferentiability. Fix non-empty sets $\mathcal{K}$ and $\mathcal{M}$, and let E be the ideal cipher with key space $\mathcal{K}$ and message space $\mathcal{M}$. Recall that $\mathsf{E.Dom} = \{+,-\} \times \mathcal{K} \times \mathcal{M}$. For any integer $1 \leq \mu \leq |\mathcal{K}|$, let $\mathcal{X}_\mu$ be the family of subsets of $\mathsf{E.Dom}$ consisting of queries whose key is in $\mathcal{K}'$, for $\mathcal{K}'$ ranging over all subsets of $\mathcal{K}$ of size μ; more formally,

$$\mathcal{X}_\mu = \{\{(+,k,x) : k \in \mathcal{K}'\} \cup \{(-,k,y) : k \in \mathcal{K}'\} : \mathcal{K}' \subseteq \mathcal{K}, |\mathcal{K}'| = \mu\}.$$

Note that $\mathcal{X}_{|\mathcal{K}|} = \{\mathsf{E.Dom}\}$.

Definition 2 (μ-Known-Key Indifferentiability). *Let $\mathcal{C}$ be a construction of a block cipher with key space $\mathcal{K}$ and message space $\mathcal{M}$ from an ideal primitive*

F. *Let* $\mu, q, \sigma, t \in \mathbb{N}$ *and* $\varepsilon \in \mathbb{R}^+$. *Construction* $\mathcal{C}$ *is said to be* $(\mu, q, \sigma, t, \varepsilon)$-*KK-indifferentiable from an ideal cipher if and only if it is* $(\mathcal{X}_\mu, q, \sigma, t, \varepsilon)$-*RI-indifferentiable from an ideal cipher, with* $\mathcal{X}_\mu$ *defined as above.*

Unfolding the definition, this is equivalent to the following: for any subset $\mathcal{K}' \subseteq \mathcal{K}$ *of size* μ, *there exists a simulator* $\mathcal{S}$ *such that for any distinguisher* $\mathcal{D}$ *whose queries to its first (construction/ideal cipher) oracle use only keys* $k \in \mathcal{K}'$ *and of total oracle query cost at most* q, $\mathcal{S}$ *makes at most* σ *oracle queries, runs in time at most* t, *and*

$$\mathbf{Adv}^{\text{indiff}}_{\mathcal{C},\mathcal{S}}(\mathcal{D}) \leq \varepsilon.$$

The KK-indifferentiability notion of Andreeva *et al.* [ABM13] corresponds to the definition above for $\mu = 1$. In fact, this is slightly more subtle. Their variant is rather an "average" version of this definition over the random draw of the known key, resulting from the following changes: the security experiment starts by drawing a random key k which is given as input to both the distinguisher and the simulator, and the two probabilities involved in the definition (1) of the advantage of the distinguisher are also taken over the random draw of the challenge key $k \leftarrow_\$ \mathcal{K}$. It is not hard to see that our "worst-case" variant of the definition is stronger (i.e., implies) the average-case version (the average-case simulator simply has a copy of each worst-case simulator $\mathcal{S}_{\mathcal{K}'}$ for each possible subset $\mathcal{K}' \subseteq \mathcal{K}$ of size μ, and on input the challenge subset of keys runs the corresponding worst-case simulator).

The standard indifferentiability notion [MRH04] is recovered by letting $\mu = |\mathcal{K}|$ in the definition above. The composition theorem specializes to the case of μ-KK-indifferentiability as follows: if a cryptosystem is proven secure when implemented with an ideal cipher E with key space $\mathcal{K}$ and if for any adversary $\mathcal{A}$, there is a subset of keys $\mathcal{K}'$ of size μ such that the challenger of the security game only queries E with keys $k \in \mathcal{K}'$ when interacting with $\mathcal{A}$, then it remains provably secure when E is replaced with $\mathcal{C}$ based on ideal primitive F, assuming $\mathcal{C}$ is μ-KK-indifferentiable from an ideal cipher.

Known-Key Correlation Intractability. As for the general notion of RI-indifferentiability, KK-indifferentiability can be combined with the notion of sequential indifferentiability. Hence, if we restrict Definition 2 by quantifying only over sequential distinguishers, we obtain the notion of KK-seq-indifferentiability (see also Fig. 2). This notion is interesting because it implies the (arguably more natural) notion of known-key *correlation intractability*, as we explain now.

For this, we first recall the concept of evasive relation and correlation intractability [CGH98, MPS12, CS15]. Let E be an ideal primitive. For an integer $m \geq 1$, an m-ary *relation* $\mathcal{R}$ (for E) is simply a subset $\mathcal{R} \subset (\mathsf{E.Dom})^m \times (\mathsf{E.Rng})^m$. Informally, a relation is *evasive* with respect to E if it is hard, on average, for an adversary with oracle access to a random instance E of E to find 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. The definition below is very general and applies to any ideal primitive.

Definition 3 (Evasive Relation). *Let* E *be an ideal primitive. An* m*-ary relation* $\mathcal{R}$ *for* E *is said* (q, ε)-evasive *if for any adversary* $\mathcal{A}$ *with oracle access to an*

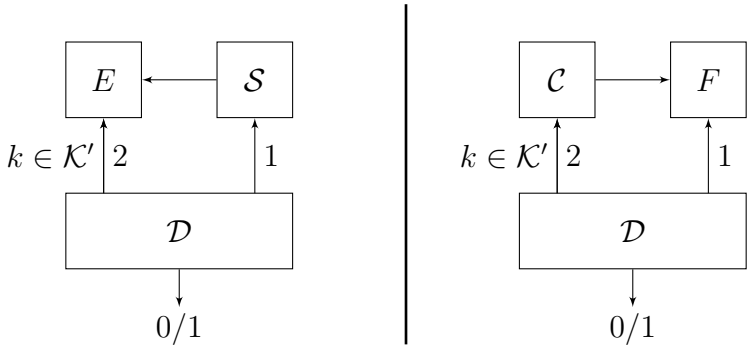

Fig. 2. Various flavors of the indifferentiability notion. For full indifferentiability, the queries of the distinguisher are completely unrestricted. For μ-known-key indifferentiability, queries to the left oracle (ideal cipher/construction) can only be made for keys $k \in \mathcal{K}'$ for some subset $\mathcal{K}'$ of size μ of the key space $\mathcal{K}$ (the simulator being allowed to depend on $\mathcal{K}'$). For sequential indifferentiability, 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. Combining the two constraints results in the KK-seq-indifferentiability notion.

instance E of E*, making at most q oracle queries, one has*

$$\Pr\left[E \leftarrow_{\$} \mathsf{E.Inst}, (\alpha_1, \ldots, \alpha_m) \leftarrow \mathcal{A}^E : ((\alpha_1, \ldots, \alpha_m), (E(\alpha_1), \ldots, E(\alpha_m))) \in \mathcal{R}\right] \leq \varepsilon,$$

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

Recall that the domain and the range of an ideal cipher E with key space $\mathcal{K}$ and message space $\mathcal{M}$ are $\mathsf{E.Dom} = \{+,-\} \times \mathcal{K} \times \mathcal{M}$ and $\mathsf{E.Rng} = \mathcal{M}$ so that, if we particularize the definition above for an ideal cipher, each α_i is a triplet in $\mathsf{E.Dom}$, and $E(\alpha_i) \in \mathcal{M}$.

If we now consider a construction $\mathcal{C}$ implementing E from some other ideal primitive F, a natural thing to ask is that any relation which is evasive with respect to E remains hard to find for $\mathcal{C}^F$, on average over the random draw of F, for any adversary with oracle access to F. This is formalized by the following definition.

Definition 4 (Correlation Intractability). *Let* E *and* F *be two ideal primitives, and let $\mathcal{C}$ be a construction implementing* E *from* F*. Let $\mathcal{R}$ be an m-ary relation for* E*. Then $\mathcal{C}$ is said to be* (q, ε)-correlation intractable *with respect to $\mathcal{R}$ if for any adversary $\mathcal{A}$ with oracle access to an instance of* F*, making at most*

q oracle queries, one has

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

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

A theorem by Mandal *et al.* [MPS12] (see also [CS15, Theorem 4]) establishes that seq-indifferentiability allows, for any relation $\mathcal{R}$, to "reduce" the correlation intractability of $\mathcal{C}$ with respect to $\mathcal{R}$ to the evasiveness of $\mathcal{R}$ (with respect to E). More precisely, if $\mathcal{C}$ is seq-indifferentiable from E and if a relation $\mathcal{R}$ is (q, ε)-evasive with respect to E, then $\mathcal{C}$ is (q', ε')-correlation intractable with respect to $\mathcal{R}$, and the "degradation" of security parameters (q', ε') compared with (q, ε) depends on the seq-indifferentiability parameters. In other words, if $\mathcal{C}$ is seq-indifferentiable from E, then any relation which is hard to find for E remains hard to find for $\mathcal{C}^F$ (on average over the random draw of F).

This result can be straightforwardly declined for the case of KK-seq-indifferentiability (and more generally RI-seq-indifferentiability): if $\mathcal{C}$ is $\mathcal{X}$-RI-seq-indifferentiable from E for some family $\mathcal{X}$ of subsets of $\mathsf{E.Dom}$, then a similar result holds, but only for relations $\mathcal{R}$ such that all inputs involved in $\mathcal{R}$ belong to some subset $X \in \mathcal{X}$; similarly, if $\mathcal{C}$ is μ-KK-seq-indifferentiable from an ideal cipher E with key space $\mathcal{K}$, then the result holds for relations $\mathcal{R}$ such that all inputs involved in $\mathcal{R}$ use the same μ keys.

Concretely we have the following theorem. The proof is similar to the proof of [CS15, Theorem 4] and therefore deferred to Appendix A. First we give two preliminary definitions. Let E be an ideal primitive, and X be a subset of $\mathsf{E.Dom}$; then an m-ary relation $\mathcal{R}$ for E is said X-restricted if

$$\forall((\alpha_1, \ldots, \alpha_m), (\beta_1, \ldots, \beta_m)) \in \mathcal{R},\ \forall i = 1, \ldots, m,\ \alpha_i \in X.$$

Similarly, let E be an ideal cipher with key space $\mathcal{K}$, and $\mu \geq 1$; then an m-ary relation $\mathcal{R}$ for E is said μ-restricted if there exists a subset $\mathcal{K}'$ of $\mathcal{K}$ of size μ such that

$$\forall((\delta_i, k_i, z_i), \ldots, (\delta_m, k_m, z_m)), (z'_1, \ldots, z'_m)) \in \mathcal{R},\ \forall i = 1, \ldots, m,\ k_i \in \mathcal{K}'.$$

Theorem 1. *Let E and F be two ideal primitives, and let $\mathcal{C}$ be a construction implementing E from F such that $\mathcal{C}$ makes at most c queries to its oracle on any input. Let $\mathcal{X}$ be a family of subsets of $\mathsf{E.Dom}$. Assume that $\mathcal{C}$ is $(\mathcal{X}, q+cm, \sigma, t, \varepsilon)$-RI-seq-indifferentiable from E. Then for any m-ary relation $\mathcal{R}$ which is X-restricted for some $X \in \mathcal{X}$, if $\mathcal{R}$ is $(\sigma + m, \varepsilon_{\mathcal{R}})$-evasive with respect to E, then $\mathcal{C}$ is $(q, \varepsilon + \varepsilon_{\mathcal{R}})$-correlation intractable with respect to $\mathcal{R}$.*

In particular, let E be an ideal cipher with key space $\mathcal{K}$, and assume that $\mathcal{C}$ is $(\mu, q+cm, \sigma, t, \varepsilon)$-KK-seq-indifferentiable from E. Then for any μ-restricted m-ary relation $\mathcal{R}$, if $\mathcal{R}$ is $(\sigma + m, \varepsilon_{\mathcal{R}})$-evasive with respect to E, then $\mathcal{C}$ is $(q, \varepsilon + \varepsilon_{\mathcal{R}})$-correlation intractable with respect to $\mathcal{R}$.

Remark 1. We need to dispel some confusion that might be created by the following observation (this will also help illustrate all definitions above with a concrete example): Lampe and Seurin [LS13] have exhibited an attacker against the 3-round IEM construction which, given oracle access to the inner permutations, finds four tuples (k_i, x_i, y_i), $i = 1, \ldots, 4$, satisfying the following evasive relation:

$$\begin{cases} k_1 \oplus k_2 \oplus k_3 \oplus k_4 = 0 \\ x_1 \oplus x_2 \oplus x_3 \oplus x_4 = 0 \\ y_1 \oplus y_2 \oplus y_3 \oplus y_4 = 0. \end{cases}$$

Since we will later prove that the 3-round IEM construction is μ-KK-seq-indifferentiable from an ideal cipher for any polynomial μ, this might seem contradictory with Theorem 1. The catch is that two of the four keys involved in the relation and obtained at the end of the attack are not controlled by the adversary and in fact range over the entire key space when the inner permutations range over $\mathsf{Perm}(n)$. Hence, the evasive relation actually involves keys from the entire key space (not just a small subset of it).

4 KK-Attack on the Two-Round IEM Construction

We explained in Section 1 that the 1-round EM construction is not resistant to μ-known-key attacks for $\mu \geq 2$. We show here that this extends to the 2-round IEM construction (with independent inner permutations and the trivial key-schedule), more formally, that this construction is not μ-KK-seq-indifferentiable from an ideal cipher for $\mu \geq 2$. Our attack shares some similarities with the related-key attack against the same construction of [CS15]. Formally, we prove the following theorem.

Theorem 2. *The 2-round IEM construction* $\mathsf{EM}[n, 2, \mathbf{f}]$ *with independent inner permutations and the trivial key schedule*[3] $\mathbf{f}$ *is not* 2*-KK-seq-indifferentiable from an ideal cipher. More precisely, for any pair of distinct keys* (k_1, k_2)*, there is an adversary which distinguishes the construction from an ideal cipher with advantage close to 1 by making only queries to its left (construction/ideal cipher) oracle involving these two keys. The adversary makes no queries to its right (inner permutations/simulator) oracle.*

Proof. We denote generically (E, F) the oracles to which the adversary has access and (k_1, k_2) two distinct keys the attacker is allowed to use. Consider the following distinguisher (see Fig. 3 for a diagram of the attack):

(1) choose an arbitrary value $x_1 \in \{0,1\}^n$, and query $y_1 := E(+, k_1, x_1)$;
(2) compute $x_2 := x_1 \oplus k_2 \oplus k_1$, and query $y_2 := E(+, k_2, x_2)$;
(3) compute $y_3 := y_1 \oplus k_1 \oplus k_2$, and query $x_3 := E(-, k_2, y_3)$;
(4) compute $y_4 := y_2 \oplus k_2 \oplus k_1$, and query $x_4 := E(-, k_1, y_4)$;
(5) check whether $x_4 = x_3 \oplus k_1 \oplus k_2$.

[3] In fact, the attack applies whenever the key-schedule is linear.

When the distinguisher is interacting with an ideal cipher E, two cases can occur. Either $y_4 = y_1$, or $y_4 \neq y_1$. In the first case, this means that $y_1 \oplus y_2 = k_1 \oplus k_2$, which happens with probability 2^{-n} since x_1 and x_2 are the first queries to the uniformly random and independent permutations E_{k_1} and E_{k_2}. If $y_4 \neq y_1$, then y_4 is the second query to the uniformly random permutation E_{k_1}, thus x_4 is uniformly random and this equality happens with probability at most $1/(2^n - 1)$. Moreover one has $y_2 \neq y_1 \oplus k_1 \oplus k_2$ which happens with probability $1 - 2^{-n}$ since x_2 is the first query to E_{k_2}. Since E is a uniformly randomly drawn blockcipher, E_{k_1} and E_{k_2} are independent permutations and this case happens with probability at most 2^{-n}. Overall, when E is an ideal cipher, this relation is satisfied with a probability at most 2^{n-1}.

Now we show that when the distinguisher is interacting with the two round Even-Mansour construction, it always returns 1, independently of k, and the inner permutations, which we denote P_1 and P_2. Noting that, by definition, $x_2 = x_1 \oplus k_2 \oplus k_1$, we denote u_1 the common value

$$u_1 \stackrel{\text{def}}{=} x_1 \oplus k_1 = x_2 \oplus k_2,$$

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

$$\begin{aligned} u_2 &= v_1 \oplus k_1 && (2)\\ v_2 &= P_2(u_2) \\ u'_2 &= v_1 \oplus k_2 && (3)\\ v'_2 &= P_2(u'_2). \end{aligned}$$

Hence, one has

$$\begin{aligned} y_1 &= v_2 \oplus k_1 && (4)\\ y_2 &= v'_2 \oplus k_2. && (5) \end{aligned}$$

Since $y_3 = y_1 \oplus k_1 \oplus k_2$, we can see, using (4), that

$$y_3 \oplus k_2 = y_1 \oplus k_1 = v_2.$$

Define

$$\begin{aligned} v'_1 &= u_2 \oplus k_2 && (6)\\ u'_1 &= P_1^{-1}(v'_1). \end{aligned}$$

This implies that

$$x_3 = u'_1 \oplus k_2. \quad (7)$$

Since $y_4 = y_2 \oplus k_2 \oplus k_1$, we see by (5) that

$$y_4 \oplus k_1 = y_2 \oplus k_2 = v'_2.$$

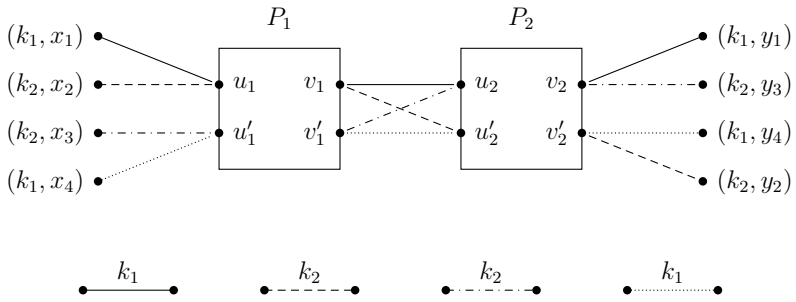

Fig. 3. A 2-known-key attack on the iterated Even-Mansour cipher with two rounds and the trivial key-schedule.

Moreover, we have

$$\begin{aligned} u'_2 \oplus k_1 &= u'_2 \oplus k_2 \oplus k_1 \oplus k_2 \\ &= v_1 \oplus k_1 \oplus k_2 && \text{by (3)} \\ &= u_2 \oplus k_2 && \text{by (2)} \\ &= v'_1 && \text{by (6)}. \end{aligned}$$

This finally implies by (7) that

$$x_4 \oplus k_1 = u'_1 = x_3 \oplus k_2,$$

which concludes the proof. □

5 KK-Seq-Indifferentiability for Three Rounds

We have just given a 2-known-keys attack against the 2-round IEM cipher. This implies that the 2-round IEM construction cannot be μ-KK-seq-indifferentiable from an ideal cipher as soon as $\mu \geq 2$. (Remember on the other hand that the 1-round EM construction is 1-KK-indifferentiable from an ideal cipher [ABM13].) Hence, at least three rounds are necessary (and, as we will see now, sufficient) to achieve μ-KK-seq-indifferentiability from an ideal cipher for $\mu \geq 2$.

Concretely, the main result of this section regarding the KK-seq-indifferentiability of the 3-round IEM cipher is as follows.

Theorem 3. *Let $N = 2^n$. For any integers μ and q such that $\mu q \leq N/4$, the 3-round IEM construction $\mathsf{EM}[n, 3, \mathbf{f}]$ with independent permutations and the trivial key-schedule $\mathbf{f}$ is $(\mu, q, \sigma, t, \varepsilon)$-KK-seq-indifferentiable from an ideal cipher with n-bit blocks and n-bit keys, with*

$$\sigma = \mu q, \quad t = \mathcal{O}(\mu q), \quad \text{and} \quad \varepsilon = \frac{57\mu^2 q^2}{N}.$$

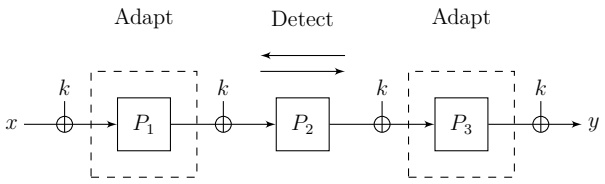

Fig. 4. Detection and adaptations zones used by the simulator for proving KK-seq-indifferentiability of the 3-round iterated Even-Mansour construction from an ideal cipher.

As a corollary, we obtain from Theorem 1 that for any m-ary relation $\mathcal{R}$ which is μ-restricted and $(\mu q, \varepsilon)$-evasive w.r.t. an ideal cipher (and assuming q is large compared with $c = 3$ and m), the 3-round IEM cipher is $\left(q, \varepsilon + \mathcal{O}\left(\mu^2 q^2 / 2^n\right)\right)$-correlation intractable with respect to $\mathcal{R}$.

It is also known [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 the ideal primitive E. Since an ideal cipher is stateless, Theorem 3 implies that the 3-round IEM construction is also KK-publicly indifferentiable from an ideal cipher.

Proof Idea. The proof of Theorem 3 is very similar to the proof of (full, not KK) seq-indifferentiability for the 4-round IEM construction of [CS15]. The main difference in the simulation strategy is the following: in the full seq-indifferentiability setting, the simulator has no hint about which key(s) the adversary is using to try to distinguish the real world from the ideal (simulated) world. Hence, it uses a 2-round "detection" zone in the middle made of permutations P_2 and P_3, which allows, given a query to P_2 (say, $P_2(u_2) = v_2$) and a query to P_3 (say, $P_3(u_3) = v_3$), to deduce the key associated to this "chain" of queries (namely, $k = v_2 \oplus u_3$). Permutations P_1 and P_4 are then used to "adapt" these detected chains and make them match the ideal cipher E. In the KK-setting, the simulator knows the set $\mathcal{K}'$ of keys that the distinguisher is allowed to use in its ideal cipher queries. Hence, the detection zone can be reduced to one single round (the middle one, i.e. P_2 for the 3-round IEM): each time the distinguisher makes a query to P_2, the simulator completes the μ chains corresponding to this query and *each key* $k \in \mathcal{K}'$, again using extremal round P_1 and P_3 to adapt the chains (see Fig. 4).

We only give an informal description of the simulator here and defer the formal description in pseudocode and the full proof of Theorem 3 to Appendix B. The simulator is given the subset $\mathcal{K}'$ of keys that the distinguisher is bound to use. It offers an interface $\mathsf{Query}(i, \delta, w)$ to the distinguisher for querying the internal permutations, where $i \in \{1, 2, 3\}$ 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 $i = 1, \ldots, 3$, the simulator internally maintains a table Π_i reflecting which values have been already internally set for each simulated permutation. Each table maps 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 checks in table Π_i whether the corresponding answer $\Pi_i(\delta, w)$ is already defined. When this is the case, it returns the answer to the distinguisher 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 $\Pi_i(\bar{\delta}, w') := w$. The randomness used by the simulator is made explicit through a tuple of random permutations $\mathbf{P} = (P_1, P_2, P_3)$ with $P_i := \{+,-\} \times \{0,1\}^n \to \{0,1\}^n$, 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, P_2, P_3) 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 $w' := P_i(\delta, w)$.[4]

Before returning w' to the distinguisher, the simulator takes additional steps to ensure that the whole IEM construction matches the ideal cipher E by running a *chain completion* mechanism. Namely, if the distinguisher called $\mathsf{Query}(i, \delta, w)$ with $i = 2$, the simulator completes the "chains" for each known key $k \in \mathcal{K}'$ by executing a procedure $\mathsf{CompleteChain}(u_2, v_2, k, \ell)$, where ℓ indicates where the chain will be "adapted" and (u_2, v_2) is the pair of values that was just added to Π_2. For example, assume that the distinguisher called $\mathsf{Query}(2, +, u_2)$ and that the answer randomly chosen by the simulator was v_2. Then for each $k \in \mathcal{K}'$, the simulator computes the corresponding value $u_3 = v_2 \oplus k$, and evaluates the IEM construction backward, letting $v_1 := u_2 \oplus k$, $u_1 := \Pi_1(-, v_1)$ (setting this value at random in case it was not in Π_1), $x := u_1 \oplus k$, $y := E(+, k, x)$ (hence making a query to E to "wrap around"), and $v_3 := y \oplus k$, until the corresponding input/output values (u_3, v_3) for the third permutation are defined. It then "adapts" (rather than setting randomly) table Π_3 by calling procedure $\mathsf{ForceVal}(u_3, v_3, 3)$ which sets $\Pi_3(+, u_3) := v_3$ and $\Pi_3(-, v_3) := u_3$ 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 $\mathsf{Query}(2, -, \cdot)$, the behavior of the simulator is symmetric, namely adaptation of the chain takes place in table Π_1.

6 KK-indifferentiability for Nine Rounds

In this section, we show that nine rounds of the IEM construction are sufficient to achieve μ-KK-indifferentiability from an ideal cipher. Note that this is less than

[4] Note that for $i = 1$ and $i = 3$, this is not equivalent to letting $w' \leftarrow_{\$} \{0,1\}^n \setminus \Pi_i^{\bar{\delta}}$ since the simulator sometimes "adapts" the value of these tables, so that the tables Π_i and the permutations P_i will differ (with overwhelming probability) on adapted entries.

what is currently known to be sufficient to achieve full indifferentiability from an ideal cipher, namely twelve rounds, as shown by Lampe and Seurin [LS13]. We conjecture that four rounds are actually sufficient.

We use the same technique as in Section 5 for going from four rounds for seq-indifferentiability to three rounds for KK-seq-indifferentiability: we start from the 12-round simulator of [LS13], and shorten the detection zones using the fact that the simulator knows the subset of keys used by the distinguisher.

We only give an informal description of the simulator and sketch how to modify the indifferentiability proof of [LS13], so that the result should rather be considered as a (substantiated) conjecture. (Given that nine is unlikely to be the minimal number of rounds needed to achieve μ-KK-indifferentiability, and that we already known that twelve rounds are sufficient to achieve full indifferentiability and hence μ-KK-indifferentiability, the benefit of writing down the full proof is rather low.) The high-level principle of how the simulator works is similar to Section 5 except that there are now additional detection zones besides the middle one preventing the distinguisher from creating "wrap around" chains (remember that the distinguisher is not bound to be sequential here, so it can make an ideal cipher query $y := E(+, k, x)$ and evaluate the IEM construction from both extremities by making permutation queries until the simulator is trapped into a contradiction). Moreover, since the simulator can now recurse (i.e., completing a chain can create new chains to be completed), it uses a queue of chains detected and to be completed as in [LS13].

As before, the simulator reacts on any query to P_5, and completes the chains for any key $k \in \mathcal{K}'$ by adapting at P_7 if this is a direct query and adapting at P_3 if this is an inverse query. Moreover, the simulator also reacts on direct queries to P_1 or inverse queries to P_9. Let us consider the case of a query $P_1(+, u_1)$. Then for each key $k \in \mathcal{K}'$, the simulator computes $x := u_1 \oplus k$, queries $y := E(+, k, x)$, lets $v_9 := y \oplus k$, and checks if $v_9 \in \Pi_9^-$. If this is the case, then the chain (u_1, k) is enqueued to be completed and adapted at P_3. For an inverse query to P_9, adaptation takes place at P_7. As in [LS13], the four "buffer" rounds P_2, P_4, P_6 and P_8 surrounding adaptation rounds ensure that no collision can occur when adapting distinct chains.

The analysis of this simulator then follows the same lines as in [LS13]. Its complexity can be upper bounded as follows: first, one applies the standard argument that the number of wrap-around chains that will be detected is upper bounded (with very high probability) by the number of ideal cipher queries of the distinguisher, hence by q. This implies that the size of table Π_5 is always at most $2q$ (since it increases only because of a distinguisher's query or when completing a wrap-around chain). It follows that the number of middle chains completed is at most $2\mu q$, and the size of all tables Π_i for $i \neq 5$ is at most $q + q + 2\mu q = 2(\mu + 1)q$. Also, the number of calls made by the simulator to the ideal cipher can be upper bounded by $2\mu q$ (number of middle chains that are completed), plus $4\mu(\mu + 1)q$ (number of wrap-around chains that are checked), hence it is $O(\mu^2 q)$ (the running time is similar).

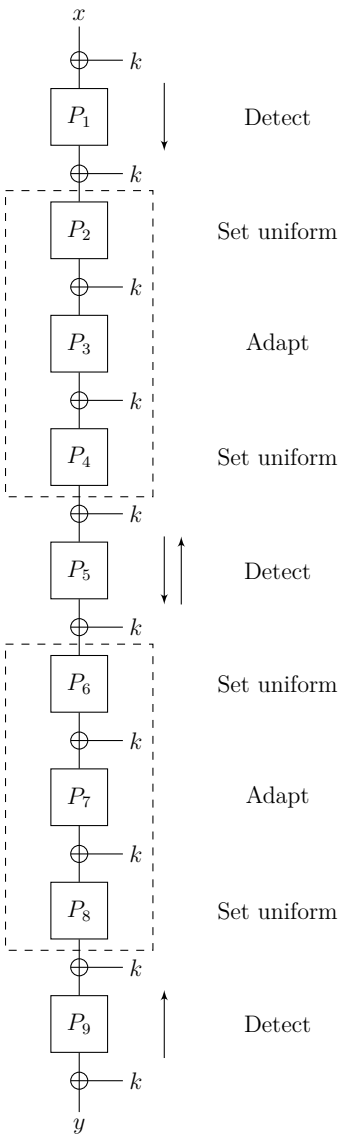

Fig. 5. Detection and adaptation zones used by the simulator for proving KK-indifferentiability of the 9-round iterated Even-Mansour construction from an ideal cipher.

Finally, proving a rigorous upper bound on the distinguishing advantage is a cumbersome task that remains to be done. A rough estimation following the lines of [LS13] would be that bad events that would make the simulator to overwrite a value when adapting chains (which is what dominates the security bound) happen with probability at most $(\max|\Pi_i|)^6/2^n$, hence $O(\mu^6 q^6)$.

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.

[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.

[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.

[CS15] Benoît Cogliati and Yannick Seurin. On the Provable Security of the Iterated Even-Mansour Cipher against Related-Key and Chosen-Key Attacks. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 (Proceedings, Part I)*, volume 9056 of *LNCS*, pages 584–613. Springer, 2015. Full version available at http://eprint.iacr.org/2015/069.

[DGHM13] Grégory Demay, Peter Gazi, Martin Hirt, and Ueli Maurer. Resource-Restricted Indifferentiability. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 664–683. Springer, 2013. Full version available at http://eprint.iacr.org/2012/613.

[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 Gregor Leander, editor, *Fast Software Encryption - FSE 2015*, volume 9054 of *LNCS*, pages 342–363. Springer, 2015. Full version available at http://eprint.iacr.org/2014/953.

[Gil14] Henri Gilbert. A Simplified Representation of AES. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 (Proceedings, Part I)*, volume 8873 of *LNCS*, pages 200–222. Springer, 2014.

[GL15a] Chun Guo and Dongdai Lin. A Synthetic Indifferentiability Analysis of Interleaved Double-Key Even-Mansour Ciphers. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 (Proceedings, Part II)*, volume 9453 of *LNCS*, pages 389–410. Springer, 2015.

[GL15b] Chun Guo and Dongdai Lin. On the Indifferentiability of Key-Alternating Feistel Ciphers with No Key Derivation. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - TCC 2015 (Proceedings, Part I)*, volume 9014 of *LNCS*, pages 110–133. Springer, 2015.

[GP10] Henri Gilbert and Thomas Peyrin. Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In Seokhie Hong and Tetsu Iwata, editors, *Fast Software Encryption - FSE 2010*, volume 6147 of *LNCS*, pages 365–383. 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.

[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.

[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.

[MPP09] Marine Minier, Raphael C.-W. Phan, and Benjamin Pousse. Distinguishers for Ciphers and Known Key Attack against Rijndael with Large Blocks. In Bart Preneel, editor, *Progress in Cryptology - AFRICACRYPT 2009*, volume 5580 of *LNCS*, pages 60–76. Springer, 2009.

[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.

[NPSS10] Ivica Nikolic, Josef Pieprzyk, Przemyslaw Sokolowski, and Ron Steinfeld. Known and Chosen Key Differential Distinguishers for Block Ciphers. In Kyung Hyune Rhee and DaeHun Nyang, editors, *Information Security and Cryptology - ICISC 2010*, volume 6829 of *LNCS*, pages 29–48. Springer, 2010.

[RS08a] Phillip Rogaway and John P. Steinberger. Constructing Cryptographic Hash Functions from Fixed-Key Blockciphers. In David Wagner, editor, *Advances in Cryptology - CRYPTO 2008*, volume 5157 of *LNCS*, pages 433–450. Springer, 2008.

[RS08b] Phillip Rogaway and John P. Steinberger. Security/Efficiency Tradeoffs for Permutation-Based Hashing. In Nigel P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 220–236. Springer, 2008.

[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.

[SY11] Yu Sasaki and Kan Yasuda. Known-Key Distinguishers on 11-Round Feistel and Collision Attacks on Its Hashing Modes. In Antoine Joux, editor, *Fast Software Encryption - FSE 2011*, volume 6733 of *LNCS*, pages 397–415. Springer, 2011.

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

A Proof of Theorem 1

Assume that there exists $X \in \mathcal{X}$ and an m-ary X-restricted relation $\mathcal{R}$ which is $(\sigma+m, \varepsilon_{\mathcal{R}})$-evasive with respect to E 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 adversary $\mathcal{A}$ making at most q oracle queries such that $\mathcal{A}^F$ 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 X-restricted sequential distinguisher $\mathcal{D}$ accessing a pair of oracles (E, F): it runs $\mathcal{A}$, answering $\mathcal{A}$'s oracle queries with its own oracle F,

until $\mathcal{A}$ returns a tuple $(\alpha_1, \ldots, \alpha_m)$. If $(\alpha_1, \ldots, \alpha_m) \notin X^m$, then $\mathcal{D}$ returns 0. Otherwise, if $(\alpha_1, \ldots, \alpha_m) \in X^m$, then $\mathcal{D}$ makes oracle queries $E(\alpha_1), \ldots, E(\alpha_m)$ and checks[5] 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 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 $(\mathcal{X}, q + cm, \sigma, t, \varepsilon)$-RI-seq-indifferentiability of $\mathcal{C}$.

B Proof of Theorem 3

In order to prove Theorem 3, we first define a simulator $\mathcal{S}$ in pseudocode in Fig. 6 (the simulator depends on the subset of keys $\mathcal{K}'$, but we do not denote it explicitly further in the proof), then prove that it runs in polynomial time and makes a polynomial number of queries (Lemma 1), and finally prove that the two systems $\Sigma_1 = (E, \mathcal{S}^E)$ and $\Sigma_3 = (\mathsf{EM}^{\mathbf{P}}, \mathbf{P})$ are indistinguishable, using an intermediate system Σ_2 that we will describe later (Lemmas 3 and 4).

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

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 1. *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 is at most q, and the size of Π_1 and Π_3 is at most $\mu q + q$;
(ii) the simulator executes CompleteChain *at most μq times, makes at most μq queries to E, and runs in time $\mathcal{O}(\mu q)$.*

Proof. The size of Π_2 can only increase by one when the distinguisher makes a direct call to $\mathsf{Query}(2, \delta, w)$, so that the size of Π_2 is at most q. Procedure CompleteChain is called once for each pair in $(u_2, k) \in \Pi_2^+ \times \mathcal{K}'$, hence at most μq times in total. Since the simulator makes exactly one query to E per execution of

[5] 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 $\mathcal{R}$.

```
Simulator S_{K'}(P):
Variables:
  tables Π_1, Π_2, Π_3, initially empty

public procedure Query(i, δ, w):
  if (δ, w) ∉ Π_i then
    w' := P_i(δ, w)
    Π_i(δ, w) := w'
    Π_i(δ̄, w') := w            \\ may overwrite an entry
    \\ complete all new chains
    if (i, δ) = (2, +) then
      u_2 := w; v_2 := w'
      forall k ∈ K' do
        CompleteChain(u_2, v_2, k, 3)
    if (i, δ) = (2, −) then
      u_2 := w'; v_2 := w
      forall k ∈ K' do
        CompleteChain(u_2, v_2, k, 1)
  return Π_i(δ, w)

private procedure CompleteChain(u_2, v_2, k, ℓ):
  case ℓ = 1:
    v_1 := u_2 ⊕ k
    \\ evaluate the chain fw. up to u_1
    u_3 := v_2 ⊕ k
    v_3 := Query(3, +, u_3)
    y := v_3 ⊕ k
    x := E(−, k, y)
    u_1 := x ⊕ k
    \\ adapt the chain
    ForceVal(u_1, v_1, 1)
  case ℓ = 3:
    u_3 := v_2 ⊕ k
    \\ evaluate the chain bw. up to v_3
    v_1 := u_2 ⊕ k
    u_1 := Query(1, −, v_1)
    x := u_1 ⊕ k
    y := E(+, k, x)
    v_3 := y ⊕ k
    \\ adapt the chain
    ForceVal(u_3, v_3, 3)

private procedure ForceVal(u_i, v_i, i):
  Π_i(+, u_i) := v_i           \\ may overwrite an entry
  Π_i(−, v_i) := u_i           \\ may overwrite an entry
```

Fig. 6. The 3-round KK-seq-indifferentiability simulator in pseudocode.

CompleteChain, the total number of queries made by the simulator to E is at most μq. The size of Π_1, resp. Π_3, can only increase by one when the distinguisher calls $\mathsf{Query}(1, \delta, w)$, resp. $\mathsf{Query}(3, \delta, w)$, or when CompleteChain is called, hence the size of Π_1 and Π_3 is at most $\mu q + q$. Clearly, the simulator running time is dominated by the executions of CompleteChain, hence the simulator runs in time $\mathcal{O}(\mu q)$. □

Intermediate System. In all the following, we consider some distinguisher $\mathcal{D}$, and assume that it is deterministic (this is *wlog* since we consider computationally unbounded distinguishers). We will denote $\mathcal{S}(E, \mathbf{P})$ rather than $\mathcal{S}(\mathbf{P})^E$ the simulator with oracle access to the ideal cipher E and using random permutations $\mathbf{P}$ as source of randomness. In order to prove the indistinguishability of the two systems $(E, \mathcal{S}(E, \mathbf{P}))$ and $(\mathsf{EM}^{\mathbf{P}}, \mathbf{P})$, we will use an intermediate system.[6] Let Σ_1 be the "ideal" world where the distinguisher interacts with $(E, \mathcal{S}(E, \mathbf{P}))$. Note that all the randomness of system Σ_1 is captured by the pair $(E, \mathbf{P})$. Let also Σ_3 be the "real" world where the distinguisher interacts with $(\mathsf{EM}^{\mathbf{P}}, \mathbf{P})$. All the randomness of system Σ_3 is captured by $\mathbf{P}$. In the intermediate system Σ_2, the distinguisher interacts with $(\mathsf{EM}^{\mathcal{S}(E,\mathbf{P})}, \mathcal{S}(E, \mathbf{P}))$ (see Fig. 7). In words, the right oracle is the simulator $\mathcal{S}(E, \mathbf{P})$ with oracle access to an ideal cipher E as in Σ_1, but now the left oracle is the 3-round IEM construction with oracle access to $\mathcal{S}(E, \mathbf{P})$ (rather than random permutations). As for Σ_1, all the randomness of system Σ_2 is captured by $(E, \mathbf{P})$.

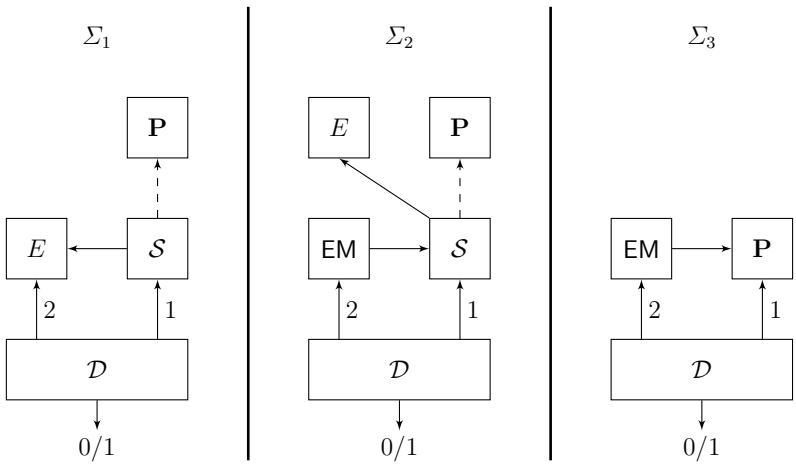

Fig. 7. Systems used in the KK-seq-indifferentiability proof.

[6] Note that this intermediate system is the same as the one used in [CS15].

TRANSITION FROM Σ_1 TO Σ_2 AND GOOD EXECUTIONS. We first consider the transition from the first to the second system.

Definition 5. *A pair* $(E, \mathbf{P})$ *is said* good *if the simulator never overwrites an entry of its tables* Π_i *during an execution of* $\mathcal{D}^{\Sigma_2(E,\mathbf{P})}$*. 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 Fig. 6), or when adapting a chain (lines (41) and (42)). Note that whether a pair $(E, \mathbf{P})$ is good or not depends on the distinguisher $\mathcal{D}$. We first upper bound the probability that a random pair $(E, \mathbf{P})$ is bad.

Lemma 2. *Consider a distinguisher* $\mathcal{D}$ *of total oracle query cost at most* q*, with* $\mu q \leq N/4$*. Then a uniformly random pair* $(E, \mathbf{P})$*, where* $E \leftarrow_{\$} \mathsf{BC}(n, n)$ *and* $\mathbf{P} \leftarrow_{\$} (\mathsf{Perm}(n))^3$*, is bad (with respect to* $\mathcal{D}$*) with probability at most* $\frac{16\mu^2 q^2}{N}$*.*

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 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_3. 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 (i.e., at line (41) or (42)).

We first consider the probability of $\mathsf{BadRand}$. Consider a random assignment in Π_i, for $i = 1$ or 3, namely $\Pi_i(\delta, w) := w'$, $\Pi_i(\bar{\delta}, w') := w$, with w' randomly defined as $w' := P_i(\delta, w)$. By Lemma 1 (i), there are at most $(\mu + 1)q$ random assignments in Π_1 and Π_3, so that w' is uniformly random in a set of size at least $N - (\mu + 1)q$. 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 1 (ii) there are at most μq such values. Hence, the probability that w' is equal to one of the at most μq values previously added in table Π_i by a call to $\mathsf{ForceVal}$ is at most $\frac{\mu q}{N-(\mu+1)q}$. Summing over the at most $(\mu + 1)q$ random assignments in Π_1 and Π_3, we get

$$\Pr\left[\mathsf{BadRand}\right] \leq 2(\mu + 1)q \times \frac{\mu q}{N - (\mu + 1)q} \leq \frac{8\mu^2 q^2}{N}. \tag{8}$$

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 $\mathsf{CompleteChain}(u_2, v_2, k, \ell)$, and assume that event $\mathsf{BadRand}$ 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 $\mathsf{Query}(2,+,\cdot)$ from the distinguisher, so that $\ell = 3$ (the reasoning is symmetric for a call to $\mathsf{Query}(2,-,\cdot)$ for which $\ell = 1$). The simulator will eventually call $\mathsf{ForceVal}(u_3, v_3, 3)$, and we must show that with high probability, the values $\Pi_3(+,u_3)$ and $\Pi_3(-,v_3)$ are undefined previously to this call. We first consider the case of u_3. This value is defined by the simulator by setting $u_3 := v_2 \oplus k$. Since the distinguisher called $\mathsf{Query}(2,+,\cdot)$ and since there are at most q random assignments in Π_2, then v_2 comes at random from a set of size at least $N-q$. Hence, the probability that u_3 is equal to one of the at most $(\mu+1)q$ values already in Π_3 is at most $\frac{(\mu+1)q}{N-q}$. We now argue that $\Pi_3(-,v_3)$ is also undefined with high probability. For this, we show that the query $E(+,k,x)$ 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 x as answer to a previous query $E(-,k,y)$. Assume that this does not hold. Then this means that such a query previously occurred when completing another chain. But since we assumed that no value was overwritten in the tables before this call to $\mathsf{CompleteChain}(u_2,v_2,k,3)$, it can easily be seen that this implies that $(u_2',v_2',k') = (u_2,v_2,k)$, which cannot be since the simulator completes any chain at most once by construction. This implies that the value y returned by E comes at random from a set of size at least $N-\mu q$ (since by Lemma 1 the simulator makes at most μq queries to E), so that $v_3 := y \oplus k$ is equal to one of the at most $(\mu+1)q$ values already in table Π_1 with probability at most $\frac{(\mu+1)q}{N-\mu q}$. Hence, summing over the at most μq calls to $\mathsf{CompleteChain}$, we obtain

$$\begin{aligned}\Pr\left[\mathsf{BadAdapt}|\neg\mathsf{BadRand}\right] &\leq \sum_{i=1}^{\mu q} \Pr\Big[\mathsf{BadAdapt}_i\Big| \\ &\qquad \neg\mathsf{BadRand} \wedge \neg\mathsf{BadAdapt}_j, j=1,\ldots,i-1\Big] \\ &\leq \mu q\left(\frac{(\mu+1)q}{N-q} + \frac{(\mu+1)q}{N-\mu q}\right) \leq \frac{8\mu^2q^2}{N}. \qquad (9)\end{aligned}$$

Combining (8) and (9) yields the result. □

Lemma 3. *For any distinguisher $\mathcal{D}$ of total oracle query cost at most q, one has*

$$\left|\Pr\left[\mathcal{D}^{\Sigma_1(E,\mathbf{P})}=1\right] - \Pr\left[\mathcal{D}^{\Sigma_2(E,\mathbf{P})}=1\right]\right| \leq \frac{16\mu^2q^2}{N},$$

where both probabilities are taken over $E \leftarrow_\$ \mathsf{BC}(n,n)$ and $\mathbf{P} \leftarrow_\$ (\mathsf{Perm}(n))^3$.

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 $\mathcal{S}(E,\mathbf{P})$) and then only its left oracle (which is E in Σ_1 and $\mathsf{EM}^{\mathcal{S}(E,\mathbf{P})}$ in Σ_2). We show that for any good pair $(E,\mathbf{P})$, the transcript of the interaction of $\mathcal{D}$ with $\Sigma_1(E,\mathbf{P})$ and $\Sigma_2(E,\mathbf{P})$ 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 $(E, \mathbf{P})$.[7] 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, $\mathsf{EM}^{\mathcal{S}(E,\mathbf{P})}(+,k,x) = E(+,k,x)$, and similarly for any backward query, $\mathsf{EM}^{\mathcal{S}(E,\mathbf{P})}(-,k,y) = E(-,k,y)$. Hence, the transcripts of the interaction of $\mathcal{D}$ with $\Sigma_1(E,\mathbf{P})$ and $\Sigma_2(E,\mathbf{P})$ are the same for any good pair $(E,\mathbf{P})$. Consequently,

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

from which the result follows by Lemma 2. □

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 $(E,\mathbf{P})$ either to the special symbol $\perp$ when $(E,\mathbf{P})$ is bad, or to a tuple of *partial permutations* $\mathbf{P'} = (P'_1, P'_2, P'_3)$ when $(E,\mathbf{P})$ is good. A partial permutation is a function $\mathbf{P}'_i : \{+,-\} \times \{0,1\}^n \to \{0,1\}^n \cup \{*\}$ such that for all $u, v \in \{0,1\}^n$, $\mathbf{P}'_i(+,u) = v \neq * \Leftrightarrow \mathbf{P}'_i(-,v) = u \neq *$.

The map Λ is defined for good pairs $(E,\mathbf{P})$ as follows: run $\mathcal{D}^{\Sigma_2(E,\mathbf{P})}$, 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 $\Lambda(E,\mathbf{P})$. Since for a good pair $(E,\mathbf{P})$, the simulator never overwrites an entry in its tables, it follows that $\Lambda(E,\mathbf{P})$ is a tuple of partial permutations as just defined above. We say that a tuple of partial permutations $\mathbf{P'} = (P'_1, P'_2, P'_3)$ is good if it has a good preimage by Λ. We say that a tuple of permutations $\mathbf{P} = (P_1, P_2, P_3)$ extends a tuple of partial permutations $\mathbf{P'} = (P'_1, P'_2, P'_3)$, denoted $\mathbf{P} \vdash \mathbf{P'}$, if for each $1 \leq i \leq 3$, P_i and P'_i agree on all entries such that $P'_i(\delta, w) \neq *$.

Lemma 4. *For any distinguisher $\mathcal{D}$ of total oracle query cost at most q, one has*

$$\left|\Pr\left[\mathcal{D}^{\Sigma_2(E,\mathbf{P})} = 1\right] - \Pr\left[\mathcal{D}^{\Sigma_3(\mathbf{P})} = 1\right]\right| \leq \frac{41\mu^2 q^2}{N},$$

where the first probability is taken over $E \leftarrow_\$ \mathsf{BC}(n,n)$ and $\mathbf{P} \leftarrow_\$ (\mathsf{Perm}(n))^3$, and the second over $\mathbf{P} \leftarrow_\$ (\mathsf{Perm}(n))^3$.

Proof. Let

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

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

By definition of the randomness mapping, for any good tuple of partial permutations $\mathbf{P'}$, the outputs of $\mathcal{D}^{\Sigma_2(E,\mathbf{P})}$ and $\mathcal{D}^{\Sigma_3(\mathbf{P})}$ are equal for any pair

[7] 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.

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

$$\varepsilon \leq \Pr\left[(E, \mathbf{P}) \text{ is bad}\right] + \sum_{\mathbf{P}' \in \Theta_1} \Pr[\Lambda(E, \mathbf{P}) = \mathbf{P}'] - \sum_{\mathbf{P}' \in \Theta_1} \Pr\left[\mathbf{P} \vdash \mathbf{P}'\right]. \tag{10}$$

Fix a good tuple of partial permutations $\mathbf{P}' = (P'_1, P'_2, P'_3)$, and let

$$|P'_i| = |\{u \in \{0,1\}^n : P'_i(+, u) \neq *\}| = |\{v \in \{0,1\}^n : P'_i(-, v) \neq *\}|.$$

Then, clearly,

$$\Pr\left[\mathbf{P} \leftarrow_{\$} (\mathsf{Perm}(n))^3 : \mathbf{P} \vdash \mathbf{P}'\right] = \frac{1}{\prod_{i=1}^{3} (N)_{|P'_i|}}.$$

Fix now any good preimage $(\tilde{E}, \tilde{\mathbf{P}})$ of $\mathbf{P}'$, where $\tilde{\mathbf{P}} = (\tilde{P}_1, \tilde{P}_2, \tilde{P}_3)$, and let q_e and q_i $(1 \leq i \leq 3)$ be the number of queries made by the simulator respectively to $\tilde{E}$ and $\tilde{P}_i$ in the execution of $\mathcal{D}^{\Sigma_2(\tilde{E},\tilde{\mathbf{P}})}$. One can check that for any pair $(E, \mathbf{P})$, $\Lambda(E, \mathbf{P}) = \mathbf{P}'$ *iff* the transcript of the interaction of $\mathcal{S}$ with $(E, \mathbf{P})$ in $\mathcal{D}^{\Sigma_2(E,\mathbf{P})}$ is the same as the transcript of the interaction of $\mathcal{S}$ with $(\tilde{E}, \tilde{\mathbf{P}})$ in $\mathcal{D}^{\Sigma_2(\tilde{E},\tilde{\mathbf{P}})}$. It follows that

$$\Pr\left[E \leftarrow_{\$} \mathsf{BC}(n, n), \mathbf{P} \leftarrow_{\$} (\mathsf{Perm}(n))^3 : \Lambda(E, \mathbf{P}) = \mathbf{P}'\right] \leq \frac{1}{(N)_{q_e} \prod_{i=1}^{3} (N)_{q_i}}.$$

(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 $\mathsf{ForceVal}$ 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

$$\sum_{i=1}^{3} |P'_i| = q_e + \sum_{i=1}^{3} q_i \leq 2\mu q + 3q, \tag{11}$$

where the inequality follows by Lemma 1 (i) on the final size of the tables Π_i maintained by the simulator. Hence, we have

$$\begin{aligned}
\frac{\Pr\left[\mathbf{P} \vdash \mathbf{P}'\right]}{\Pr\left[\Lambda(E, \mathbf{P}) = \mathbf{P}'\right]} &= \frac{(N)_{q_e} \prod_{i=1}^{3} (N)_{q_i}}{\prod_{i=1}^{3} (N)_{|P'_i|}} \\
&\geq \underbrace{\frac{N^{q_e + \sum_{i=1}^{3} q_i}}{N^{\sum_{i=1}^{3} |P'_i|}}}_{=1 \text{ by } (11)} \times \prod_{j=1}^{q_e - 1} \left(1 - \frac{j}{N}\right) \prod_{i=1}^{3} \prod_{j=1}^{q_i - 1} \left(1 - \frac{j}{N}\right) \\
&\geq 1 - \frac{q_e^2 + \sum_{i=1}^{3} q_i^2}{N}
\end{aligned}$$

$$\geq 1 - \frac{(2\mu q + 3q)^2}{N} \qquad \text{by (11)}$$
$$\geq 1 - \frac{25(\mu q)^2}{N}.$$

Combining this lower bound with (10), we obtain

$$\begin{aligned}
\varepsilon &\leq \Pr\left[(E, \mathbf{P}) \text{ is bad}\right] + \sum_{\mathbf{P}' \in \Theta_1} \Pr\left[\Lambda(E, \mathbf{P}) = \mathbf{P}'\right] \left(1 - \frac{\Pr\left[\mathbf{P} \vdash \mathbf{P}'\right]}{\Pr\left[\Lambda(E, \mathbf{P}) = \mathbf{P}'\right]}\right) \\
&\leq \Pr\left[(E, \mathbf{P}) \text{ is bad}\right] + \frac{25(\mu q)^2}{N} \sum_{\mathbf{P}' \in \Theta_1} \Pr\left[\Lambda(E, \mathbf{P}) = \mathbf{P}'\right] \\
&\leq \Pr\left[(E, \mathbf{P}) \text{ is bad}\right] + \frac{25(\mu q)^2}{N}.
\end{aligned}$$

The result follows from Lemma 2. □

Concluding. The proof of Theorem 3 directly follows by combining Lemmas 1, 3, and 4.

As a corollary, we obtain from Theorem 1 that for any m-ary, μ-restricted, $(\mu q, \varepsilon)$-evasive relation $\mathcal{R}$, the 3-round IEM cipher is $\left(q, \varepsilon + \mathcal{O}\left(\mu^2 q^2 / 2^n\right)\right)$-correlation intractable with respect to $\mathcal{R}$.