Minimizing the Two-Round Even-Mansour Cipher

Shan Chen[⋆], Rodolphe Lampe[⋆⋆], Jooyoung Lee[⋆⋆⋆],
Yannick Seurin[†], and John Steinberger[‡]

Abstract. The r-round (iterated) *Even-Mansour cipher* (also known as *key-alternating cipher*) defines a block cipher from r fixed public n-bit permutations $P_1, \ldots, P_r$ as follows: given a sequence of n-bit round keys $k_0, \ldots, k_r$, an n-bit plaintext x is encrypted by xoring round key k_0, applying permutation P_1, xoring round key k_1, etc. The (strong) pseudorandomness of this construction in the random permutation model (i.e., when the permutations $P_1, \ldots, P_r$ are public random permutation oracles that the adversary can query in a black-box way) was studied in a number of recent papers, culminating with the work of Chen and Steinberger (EUROCRYPT 2014), who proved that the r-round Even-Mansour cipher is indistinguishable from a truly random permutation up to $\mathcal{O}(2^{\frac{rn}{r+1}})$ queries of any adaptive adversary (which is an optimal security bound since it matches a simple distinguishing attack). All results in this entire line of work share the common restriction that they only hold under the assumption that *the round keys $k_0, \ldots, k_r$ and the permutations $P_1, \ldots, P_r$ are independent*. In particular, for two rounds, the current state of knowledge is that the block cipher $E(x) = k_2 \oplus P_2(k_1 \oplus P_1(k_0 \oplus x))$ is provably secure up to $\mathcal{O}(2^{2n/3})$ queries of the adversary, when k_0, k_1, and k_2 are three independent n-bit keys, and P_1 and P_2 are two independent random n-bit permutations. In this paper, we ask whether one can obtain a similar bound for the two-round Even-Mansour cipher *from just one n-bit key and one n-bit permutation*. Our answer is positive: when the three n-bit round keys k_0, k_1, and k_2 are adequately derived from an n-bit master key k, and the same permutation P is used in place of P_1 and P_2, we prove a qualitatively similar $\widetilde{\mathcal{O}}(2^{2n/3})$ security bound (in the random permutation model). To the best of our knowledge, this is the first "beyond the birthday bound" security result for AES-like ciphers that does not assume independent round keys.

Keywords: generalized Even-Mansour cipher, key-alternating cipher, indistinguishability, pseudorandom permutation, random permutation model, sum-capture problem

[⋆] Tsinghua University, P.R. China. E-mail: dragoncs16@gmail.com.
[⋆⋆] University of Versailles, France. E-mail: rodolphe.lampe@gmail.com.
[⋆⋆⋆] Sejong University, Seoul, Korea. E-mail: jlee05@sejong.ac.kr.
[†] ANSSI, Paris, France. E-mail: yannick.seurin@m4x.org.
[‡] Tsinghua University, P.R. China. E-mail: jpsteinb@gmail.com

1 Introduction

Background. An elementary way to construct a block cipher with message space $\{0,1\}^n$ from r fixed and public n-bit permutations $P_1, \dots P_r$ is to encrypt a plaintext x by computing

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

where $(k_0, \dots, k_r)$ is a sequence of n-bit round keys which are usually derived from some master key K. This construction, which captures the high-level structure of (most) block cipher designs known as Substitution-Permutation Networks (SPNs), such as AES [12], PRESENT [7], or LED [20] to name a few, was coined a *key-alternating cipher* by Daemen and Rijmen [13].

In the random permutation model (i.e., when permutations $P_1, \dots, P_r$ are modeled as public random permutation oracles), provable security results for this construction were first obtained for $r = 1$ round by Even and Mansour [16], who showed that the block cipher encrypting x into $k_1 \oplus P_1(k_0 \oplus x)$, where k_0 and k_1 are independent n-bit keys, and P_1 is a random permutation oracle, is secure up to $\mathcal{O}(2^{n/2})$ queries of the adversary.[1] For this reason, this construction is often referred to as the *Even-Mansour cipher*. Curiously, the general construction with $r > 1$ remained unstudied for a long while until a paper by Bogdanov *et al.* [8], who showed that for $r \geq 2$, security is guaranteed up to $\mathcal{O}(2^{2n/3})$ queries of the adversary. They also conjectured that the security should be $\mathcal{O}(2^{\frac{rn}{r+1}})$ for general r, which matches a simple distinguishing attack. Progress towards solving this conjecture was rather quick: Steinberger [32] proved security up to $\mathcal{O}(2^{3n/4})$ queries for $r \geq 3$, Lampe *et al.* [26] proved security up to $\mathcal{O}(2^{\frac{rn}{r+2}})$ queries for any even r, and finally Chen and Steinberger [9] resolved the conjecture and proved the $\mathcal{O}(2^{\frac{rn}{r+1}})$-security bound for any r. We stress that *all these results* only hold assuming that the $r + 1$ round keys and the r permutations are independent.[2]

Our Problem. Let us quickly recapitulate existing provable security results on the Even-Mansour cipher for a low number of rounds. For $r = 1$, we know that the single-key Even-Mansour cipher $x \mapsto k \oplus P(k \oplus x)$ ensures security up to $\mathcal{O}(2^{n/2})$ queries of the adversary. As pointed out by Dunkelman *et al.* [15], this construction is "minimal" in the sense that if one removes any component (either the addition of one of the keys, or the permutation P), the construction becomes trivially breakable. For the two-round Even-Mansour cipher, the best provable security result we have so far requires two independent n-bit permutations P_1 and P_2, and two independent n-bit keys (k, k') to construct three pairwise independent round keys, for example $(k, k' \oplus k, k')$. Concretely, the block cipher

[1] Actually it is not very hard to prove that a similar result holds when using $k_0 = k_1$.

[2] Actually, this is not perfectly accurate: one only needs the $r + 1$ round keys $(k_0, \dots, k_r)$ to be r-wise independent [9], which can be obtained from only an rn-bit long master key, the most simple example being round keys of the form $(k'_1, k'_1 \oplus k'_2, k'_2 \oplus k'_3, \dots, k'_{r-1} \oplus k'_r, k'_r)$, in which case the resulting iterated Even-Mansour cipher is exactly the cascade of r single-key one-round Even-Mansour ciphers $x \mapsto k'_i \oplus P_i(k'_i \oplus x)$.

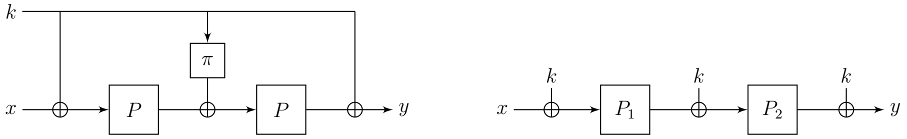

Fig. 1. Two constructions of “minimal” two-round Even-Mansour ciphers provably secure up to $\widetilde{\mathcal{O}}(2^{\frac{2n}{3}})$ queries of any (adaptive) adversary. Left: π is a (fixed) linear orthomorphism of $\mathbb{F}_2^n$, and P is a public random permutation oracle. Right: P_1 and P_2 are two independent public random permutation oracles.

$x \mapsto k' \oplus P_2((k' \oplus k) \oplus P_1(k \oplus x))$ ensures security up to $\mathcal{O}(2^{2n/3})$ queries of the adversary. In this paper, we tackle the following question:

> *Can we obtain a $\mathcal{O}(2^{2n/3})$-security bound similar to the one proven for the two-round Even-Mansour cipher with (pairwise) independent round keys and independent permutations, from just one n-bit key k and one n-bit random permutation P?*

This question is natural since in most (if not all) SPN block ciphers, round keys are derived from an n-bit master key (or more generally an ℓ-bit master key, where $\ell \in [n, 2n]$ is small compared with the total length of the round keys), and the same permutation, or very similar ones, are used at each round. It is therefore fundamental to determine whether security can actually benefit from the iterative structure and increase beyond the birthday bound, even though one does not use more key material nor more permutations than in the single-key one-round Even-Mansour cipher.

Our Results. We answer positively to the question above. Our main theorem states sufficient conditions on the way to derive three n-bit round keys (k_0, k_1, k_2) from one n-bit master key k so that the two-round Even-Mansour cipher defined from a single permutation $x \mapsto k_2 \oplus P(k_1 \oplus P(k_0 \oplus x))$ is secure up to $\widetilde{\mathcal{O}}(2^{2n/3})$ queries of the adversary, where the $\widetilde{\mathcal{O}}(\cdot)$ notation hides logarithmic (in $N = 2^n$) factors. In particular, such a good key-schedule $k \mapsto (k_0, k_1, k_2)$ can be constructed from any (fixed) linear orthomorphism of $\mathbb{F}_2^n$. A permutation π of $\{0,1\}^n$ is called an orthomorphism if $x \mapsto x \oplus \pi(x)$ is also a permutation. The good cryptographic properties of orthomorphisms have already been noticed in a number of papers [29, 19], and are in particular used in Lai-Massey schemes [25, 34] such as the block ciphers IDEA [25] and FOX [22]. Our main theorem is as follows.

Theorem (Informal). *Let π be any (fixed) linear orthomorphism of $\mathbb{F}_2^n$, and let P be a public random n-bit permutation oracle. Then the block cipher with message space and key space $\{0,1\}^n$ defined as (see Figure 1, left)*

$$\mathsf{EM}_k^P(x) = k \oplus P(\pi(k) \oplus P(k \oplus x)) \qquad (\star)$$

is secure against any adversary making up to $\widetilde{\mathcal{O}}(2^{\frac{2n}{3}})$ queries to EM_k^P and P. (Queries can be adaptive and are allowed in both directions for EM_k^P and P).

We remark that if one omits π in construction ($\star$), i.e., if one adds the same round key k each time, security drops back to $\mathcal{O}(2^{n/2})$ queries. More generally, if round keys are all equal and the same permutation P is used at each round of the iterated Even-Mansour cipher, security caps at $\mathcal{O}(2^{n/2})$ queries of the adversary, independently of the number r of rounds. This seems to be known as a folklore result about slide attacks [5, 6], but since we could not find a detailed exposition in the literature, we precisely describe and analyze this attack (as well as a simple extension for two rounds when the key-schedule simply consists in xoring constants to the master key) in this paper. Hence, construction ($\star$) can be regarded as a "minimal" two-round Even-Mansour cipher delivering security beyond the birthday bound, since removing any component causes security to drop back to $\mathcal{O}(2^{n/2})$ queries at best (for π this follows from the slide attack just mentioned, while removing any instance of permutation P brings us back to a one-round Even-Mansour cipher). Additionally, we show that when using two independent public random permutations P_1 and P_2, the trivial key-schedule is sufficient: adding the same round key k at each round (see Figure 1, right) also yields a $\widetilde{\mathcal{O}}(2^{2n/3})$-security bound.

To the best of our knowledge, these are the first results proving "beyond the birthday bound" security for key-alternating ciphers such as AES that do not rely on the assumption that round keys are independent. This sheds some light on which exact properties are required from the key-schedule in order to lift the round keys independence assumption in provable security results. In particular, this seems to point out that a *pseudorandom* key-schedule is not needed (we remind the reader that our results come with the usual caveat that they are only proved in the very strong Random Permutation Model, and hence can only be taken as a heuristic security insurance once the inner permutation(s) are instantiated).

Overview of Our Techniques. In order to prove our results, we use the indistinguishability framework, namely we consider a distinguisher which must tell apart two worlds: the "real" world where it interacts with (EM_k^P, P), where EM_k^P is the Even-Mansour cipher instantiated with permutation P and a random key k, and the "ideal" world where it interacts with (E, P) where E is a random permutation independent from P. The distinguisher can make at most q_e queries to EM_k^P/E and at most q_p queries to P (all queries are adaptive and can be forward or backward, and we work in the information-theoretic setting, i.e., the adversary is computationally unbounded). In order to upper bound the distinguishing advantage of this attacker, we use, as already done in [9], the H-coefficient method of Patarin [31]. In a nutshell, this technique consists in partitioning the set of all possible transcripts of the interaction between the distinguisher and the tuple of permutations into a set $\mathcal{T}_1$ of "good" transcripts and a set $\mathcal{T}_2$ of "bad" transcripts. Good transcripts $\tau \in \mathcal{T}_1$ have the property that the ratio of the probabilities to obtain τ in the real and in the ideal world is greater that $1 - \varepsilon_1$ for some small $\varepsilon_1 > 0$, while the probability to obtain any bad transcript $\tau \in \mathcal{T}_2$ (in the ideal world) is less than some small $\varepsilon_2 > 0$. Then the advantage of the distinguisher can be upper bounded by $\varepsilon_1 + \varepsilon_2$.

In order to get intuition about what hides behind good and bad transcripts, it helps to first look at an example of how an adversary might "get lucky" during an attack. Specifically, we focus on the following attack scenario (we assume that $q_e = q_p = q$ for simplicity). The distinguisher (adversary) $\mathcal{D}$ starts by making q arbitrary queries to EM_k^P/E, resulting in a set of q pairs $\mathcal{Q}_E = \{(x_1, y_1), \ldots, (x_q, y_q)\}$; then $\mathcal{D}$ determines the pair of sets (U, V) with $|U| = |V| = q$ and $U, V \subseteq \{0,1\}^n$, that maximizes the size of the set

$$\mathcal{K}(\mathcal{Q}_E, U, V) \overset{\text{def}}{=} \{k' \in \{0,1\}^n : \exists (x_i, y_i) \in \mathcal{Q}_E \text{ s.t. } x_i \oplus k' \in U, y_i \oplus k' \in V\}, \quad (1)$$

and $\mathcal{D}$ queries $P(u)$, $P^{-1}(v)$ for all $u \in U$, $v \in V$. (This makes $2q$ queries to P instead of q, but this small constant factor is unimportant for the sake of intuition.) Note that if $\mathcal{D}$ is in the real world and if the real key k is in the set $\mathcal{K}(\mathcal{Q}_E, U, V)$ defined in (1), then $\mathcal{D}$ can see that one of its EM_k^P/E-queries is compatible with two of its P-queries with respect to k (in more detail, there exists a value i and queries (u, v), (u', v') to P such that $x_i \oplus k = u$, $v \oplus \pi(k) = u'$, and $v' \oplus k = y_i$). Elementary probabilistic considerations show that such a "complete cycle" will occur for at most a handful of keys in $\mathcal{K}(\mathcal{Q}_E, U, V)$, so that "false alerts" can be quickly weeded out and the correct key k validated in a few extra queries, all assuming $k \in \mathcal{K}(\mathcal{Q}_E, U, V)$. Moreover, heuristic considerations indicate that k will be in $\mathcal{K}(\mathcal{Q}_E, U, V)$ with probability $|\mathcal{K}(\mathcal{Q}_E, U, V)|/2^n$. In particular, thus, it becomes necessary to show that $|\mathcal{K}(\mathcal{Q}_E, U, V)|$ is significantly smaller than 2^n with high probability over $\mathcal{Q}_E$, i.e., that

$$\max_{\substack{U,V \subseteq \{0,1\}^n \\ |U|=|V|=q}} |\{k' \in \{0,1\}^n : \exists (x_i, y_i) \in \mathcal{Q}_E \text{ s.t. } x_i \oplus k' \in U, y_i \oplus k' \in V\}| \quad (2)$$

is significantly smaller than 2^n with high probability over $\mathcal{Q}_E$, in order to show that $\mathcal{D}$ has small advantage at q queries. One of the criteria that can make a transcript "bad" in our proof happens to be, precisely, if the set of queries $\mathcal{Q}_E$ to EM_k^P/E contained within the transcript is such that (2) is larger than desirable. (Jumping ahead, $\mathcal{K}(\mathcal{Q}_E, U, V)$ will be re-baptized BadK_1 in Definition 1 of a bad transcript.)

To elaborate a little more on this, note that

$$\begin{aligned}|\mathcal{K}(\mathcal{Q}_E, U, V)| \le |\{(k', u, v) \in \{0,1\}^n \times U \times V :& \\ k' \oplus u = x_i, k' \oplus v = y_i \text{ for some } 1 \le i \le q\}|& \\ = |\{(i, u, v) \in \{1, \ldots, q\} \times U \times V : x_i \oplus y_i = u \oplus v\}|.&\end{aligned}$$

Also note that the set of values $\{x_i \oplus y_i : (x_i, y_i) \in \mathcal{Q}_E\}$ is essentially a random set since if the i-th query to EM_k^P/E is forward then y_i comes at random from a large set, whereas otherwise x_i comes at random from a large set. Moreover, as a matter of fact, the problem of upper bounding

$$\mu(A) \overset{\text{def}}{=} \max_{\substack{U,V \subseteq \{0,1\}^n \\ |U|=|V|=q}} |\{(a,u,v) \in A \times U \times V : a = u \oplus v\}$$

for a *truly random* set $A \subseteq \{0,1\}^n$ of size q has already been studied before [3, 21, 1, 24, 33], being dubbed[3] the *sum-capture problem* in [33]. One of the main known results [3, 33] on the sum-capture problem is that $\mu(A)$ is upper bounded by roughly $q^{3/2}$ for $q \leq 2^{2n/3}$. Surprisingly enough, this bound is exactly sufficient for our application, since $q^{3/2} \ll 2^n$ for $q \ll 2^{2n/3}$. (Implying, thus, that (2) is far from 2^n as long as q remains beneath $2^{2n/3}$, as desired.) Our own setting is, of course, slightly different, since the set $\{x_i \oplus y_i : (x_i, y_i) \in \mathcal{Q}_E\}$ isn't, unlike A, a purely random set of size q. Other complications also arise: in the general case where the three round keys (k_0, k_1, k_2) are derived from the n-bit master key k using non-trivial (bijective) key derivation functions $\gamma_i : k \mapsto k_i$, $\mathcal{K}(\mathcal{Q}_E, U, V)$ takes the more complicated form

$$\{k' \in \{0,1\}^n : \exists (x_i, y_i) \in \mathcal{Q}_E \text{ s.t. } x_i \oplus \gamma_0(k') \in U, y_i \oplus \gamma_2(k') \in V\},$$

so that we have to upper bound

$$|\{(i,u,v) \in \{1,\ldots,q\} \times U \times V : x_i \oplus u = \gamma_0 \circ \gamma_2^{-1}(y_i \oplus v)\}|.$$

All this means that we have to carefully adapt (and to some degree significantly extend) the Fourier-analytic techniques used in [3, 33].

Once the probability to obtain a bad transcript has been upper bounded, the second part of the proof is to show that the ratio between the probabilities to obtain any good transcript in the real and the ideal world is close to 1. This part is in essence a permutation counting argument. When the two permutations are independent (Figure 1, right), the counting argument is not overly complicated. While we could, in principle, re-use the general results of [9], we expose it in the full version of this paper [10] since it constitutes a good warm-up for the reader before the more complicated counting in the subsequent section. For the single-permutation case, things become much more involved: first, we need to consider more conditions defining bad transcripts; and second, the permutation counting itself becomes much more intricate. Interestingly, this part is related to the following simple to state (yet to the best of our knowledge unexplored) problem: how many queries are needed to distinguish a random squared permutation $P \circ P$ (where P is uniformly random) from a uniformly random permutation E?

Related Work. Two recent papers analyzed a stronger security property of the iterated Even-Mansour cipher than mere pseudorandomness, namely indifferentiability from an ideal cipher [2, 27]. Aside with provable security results already mentioned, a number of papers explored attacks on the (iterated) Even-Mansour cipher for one round [11, 6, 15], two rounds [30], three rounds [14], and four rounds [4].

[3] The terminology is attributed to Mario Szegedy.

Gazi and Tessaro [18] considered a construction they named $\mathsf{2XOR}$, which is a variant of the DESX [23] and "Xor-Cascade" [17, 28] key-length extension methods. Given a block cipher E with message space $\{0,1\}^n$ and key space $\{0,1\}^\kappa$, the $\mathsf{2XOR}$ construction defines a new block cipher with message space $\{0,1\}^n$ and key space $\{0,1\}^{\kappa+n}$ as

$$\mathsf{2XOR}^E_{z,k}(x) = E_{z_2}(k \oplus E_{z_1}(k \oplus x)),$$

where (z_1, z_2) are pairwise distinct sub-keys derived from $z \in \{0,1\}^\kappa$. They showed that, when the underlying block cipher E is modeled as an ideal cipher, this construction is secure up to $\mathcal{O}(2^{\kappa+n/2})$ queries to E, even when the adversary can make all possible 2^n queries to the permutation oracle (which, in the indistinguishability experiment, is either $\mathsf{2XOR}^E_{z,k}$ or an independent random permutation). Considering a block cipher E with key-length $\kappa = 1$, one obtains a construction which is similar to the two-round Even-Mansour cipher of Figure 1, right, where the last key addition would be omitted.[4] Hence, the Gazi-Tessaro result says that this construction is secure for $q_e = 2^n$ and $q_p = \mathcal{O}(2^{n/2})$.[5] Our own results are incomparable with the one of [18]. First, the third key addition is omitted in the $\mathsf{2XOR}$ construction. Second, our bounds are more general: they hold for any value of q_e and q_p as long as $q_e < 2^{2n/3}$ and $q_p < 2^{2n/3}$. Though our bounds become meaningless for $q_e = 2^n$, they show that when $q_e < 2^{2n/3}$ (an interesting case in practice since an attacker will not always have access to the entire codebook), security is ensured up to $\widetilde{\mathcal{O}}(2^{2n/3})$ queries to the internal permutations (something that cannot be derived from the result of [18]).

Open Questions. Currently, our results only apply when the key derivation functions mapping the master key to the round keys are *linear* bijective functions of $\mathbb{F}_2^n$. This is due to the fact that the proof of our sum-capture theorem in Section 3 requires linear mappings. It is an open question whether this theorem can be extended to nonlinear (bijective) mappings as well. A second tantalizing yet challenging open problem is of course to generalize our results to larger numbers of rounds. Namely, for $r > 2$, can we find sufficient conditions on the key-schedule such that the r-round single-permutation Even-Mansour cipher ensures security up to $\widetilde{\mathcal{O}}(2^{\frac{rn}{r+1}})$ queries of the adversary? We stress that even the

[4] There is a slight subtlety here: in the $\mathsf{2XOR}$ construction used with a block cipher with key-length $\kappa = 1$, i.e., a pair of permutations (P_1, P_2), there is an additional key bit z (hidden to the distinguisher) which tells in which order the two permutations are called.

[5] This is in fact very closely related to the security result for the single-key one-round Even-Mansour cipher up to $\mathcal{O}(2^{n/2})$ queries to the inner and outer permutations [15]. In the Gazi-Tessaro case with $\kappa = 1$, the adversary is given an arbitrary permutation E, and must distinguish, given access to (P_1, P_2), whether P_1 and P_2 are independent, or whether $P_2(k \oplus P_1(k \oplus x)) = E(x)$ for some random key k. In the single-key one-round Even-Mansour case, the adversary must distinguish, given access to (P_1, P_2), whether P_1 and P_2 are independent, or whether $k \oplus P_1(k \oplus x) = P_2(x)$, i.e., $P_2^{-1}(k \oplus P_1(k \oplus x)) = x$. These are very similar problems, the latter being (up to changing P_2 into P_2^{-1}) a special case of the former with E the identity.

simpler case where permutations are independent and round keys are identical seems hard to tackle for $r > 2$: we currently have no idea of how to extend our sum-capture result in order to upper bound the probability of bad transcripts even in the case $r = 3$.

It would also be interesting to reduce the *time* complexity of attacks against the two-round Even-Mansour cipher (potentially down to $\mathcal{O}(2^{2n/3})$). Currently, the best known attack (for the case of independent permutations and identical round keys) has time complexity $\mathcal{O}(2^{n-\log_2 n})$ [15]. Since our focus in this paper is on query complexity, we have not investigated whether this attack applies to the single-permutation variant ($\star$) as well.

Organization. We start in Section 2 by setting the notation, giving the necessary background on the H-coefficient technique, and proving some helpful lemmas. In Section 3, which is self-contained, we prove our new sum-capture result, which might be of independent interest. Section 4 contains our main provable security result for the "minimized" variant of the single-permutation two-round Even-Mansour cipher (Figure 1, left). The case where the two permutations are independent and the three round keys are identical (Figure 1, right) is treated in the full version of the paper [10]. The permutation counting argument in that case serves as a good exercise before reading the corresponding one for the single-permutation case (Lemma 3). In the full version of the paper [10], we also detail slide attacks against the iterated Even-Mansour cipher.

2 Preliminaries

Notation. In all the following, we fix an integer $n \geq 1$, and we write $N = 2^n$. The set of all permutations on $\{0,1\}^n$ will be denoted $\mathcal{P}_n$. 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. Given $\mathcal{Q} = ((x_1, y_1), \ldots, (x_q, y_q))$, where the x_i's are pairwise distinct n-bit strings and the y_i's are pairwise distinct n-bit strings, and a permutation $P \in \mathcal{P}_n$, we say that P extends $\mathcal{Q}$, denoted $P \vdash \mathcal{Q}$, if $P(x_i) = y_i$ for $i = 1, \ldots, q$. When two sets A and B are disjoint, we denote $A \sqcup B$ their (disjoint) union. We denote $\mathbb{F}_2 \simeq \{0,1\}$ the field with two elements, and $\mathbb{F}_2^n$ the vector space of dimension n over $\mathbb{F}_2$. The general linear group of degree n over $\mathbb{F}_2$, i.e., the set of all automorphisms (linear bijective mappings) of $\mathbb{F}_2^n$, will be denoted $\mathsf{GL}(n)$.

The Generalized Even-Mansour Cipher. Fix integers $n, r, m, \ell \geq 1$. Let $\phi : \{1, \ldots, r\} \to \{1, \ldots, m\}$ be an arbitrary function, and $\boldsymbol{\gamma} = (\gamma_0, \ldots, \gamma_r)$ be a $(r+1)$-tuple of functions from $\{0,1\}^\ell$ to $\{0,1\}^n$. The r-round Generalized Even-Mansour construction $\mathsf{EM}[n, r, m, \ell, \phi, \boldsymbol{\gamma}]$ specifies, from any m-tuple $\boldsymbol{P} = (P_1, \ldots, P_m)$ of permutations on $\{0,1\}^n$, a block cipher with message space $\{0,1\}^n$ and key space $\{0,1\}^\ell$, simply denoted $\mathsf{EM}^{\boldsymbol{P}}$ in the following (parameters $[n, r, m, \ell, \phi, \boldsymbol{\gamma}]$ are implicit and will always be clear from the context), which maps a plaintext $x \in \{0,1\}^n$ and a key $K \in \{0,1\}^\ell$ to the ciphertext defined by (see Figure 2):

$$\mathsf{EM}^{\boldsymbol{P}}(K, x) = \gamma_r(K) \oplus P_{\phi(r)}(\gamma_{r-1}(K) \oplus P_{\phi(r-1)}(\cdots P_{\phi(1)}(\gamma_0(K) \oplus x) \cdots)).$$

We denote $\mathsf{EM}^{\boldsymbol{P}}_K : x \mapsto \mathsf{EM}^{\boldsymbol{P}}(K, x)$ the Even-Mansour cipher instantiated with key K (hence, syntactically, $\mathsf{EM}^{\boldsymbol{P}}_K$ is a permutation on $\{0, 1\}^n$).

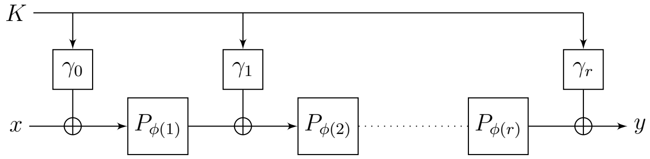

Fig. 2. The r-round Generalized Even-Mansour cipher.

All previous work about the indistinguishability of the Even-Mansour cipher [8, 26, 32, 9] considered the case where all permutations and all round keys are independent, namely $m = r$, ϕ is the identity function, $\ell = (r + 1)n$, and γ_i simply selects the i-th n-bit string of $K = (k_0, \ldots, k_r)$.

In the following, we will focus in particular on two special cases:

- the case where permutations are independent and the same n-bit key k is used at each round, namely $m = r$, ϕ is the identity function, $\ell = n$, and all γ_i's are the identity function, in which case we will simply denote $\mathsf{EMIP}[n, r]$ the resulting construction. Hence, for an r-tuple of permutations $\boldsymbol{P} = (P_1, \ldots, P_r)$, the block cipher $\mathsf{EMIP}^{\boldsymbol{P}}$ maps a plaintext $x \in \{0, 1\}^n$ and a key $k \in \{0, 1\}^n$ to the ciphertext defined by:
$$\mathsf{EMIP}^{\boldsymbol{P}}(k, x) = k \oplus P_r(k \oplus P_{r-1}(\cdots P_2(k \oplus P_1(k \oplus x)) \cdots)).$$
- the case where a single permutation P is used at each round, namely $m = 1$ and $\phi(i) = 1$ for $i = 1, \ldots, r$, in which case the resulting construction will simply be denoted $\mathsf{EMSP}[n, r, \ell, \boldsymbol{\gamma}]$. Hence, for a permutation P, the block cipher EMSP^P maps a plaintext $x \in \{0, 1\}^n$ and a key $K \in \{0, 1\}^\ell$ to the ciphertext defined by:
$$\mathsf{EMSP}^P(K, x) = \gamma_r(K) \oplus P(\gamma_{r-1}(K) \oplus P(\cdots P(\gamma_1(K) \oplus P(\gamma_0(K) \oplus x)) \cdots)).$$
When additionally $\ell = n$ (namely the master key length is equal to the block length), we overload the notation and simply denote $\mathsf{EMSP}[n, r, \boldsymbol{\gamma}]$ the resulting construction.

Security Definition. To study the indistinguishability of the Generalized Even-Mansour cipher (in the Random Permutation Model), we consider a distinguisher $\mathcal{D}$ which interacts with a set of $m + 1$ permutation oracles on n bits that we denote generically $(P_0, P_1 \ldots, P_m) = (P_0, \boldsymbol{P})$. The goal of $\mathcal{D}$ is to distinguish whether it is interacting with $(\mathsf{EM}^{\boldsymbol{P}}_K, \boldsymbol{P})$, where $\boldsymbol{P} = (P_1, \ldots, P_m)$ are random and independent permutations and K is randomly chosen from $\{0, 1\}^\ell$ (we will informally refer to this case as the "real" world), or with $(E, \boldsymbol{P})$, where

E is a random n-bit permutation independent from $\boldsymbol{P}$ (the "ideal" world). Note that in the latter case the distinguisher is simply interacting with $m+1$ independent random permutations. We sometimes refer to the first permutation P_0 as the *outer* permutation, and to permutations $P_1, \ldots, P_m$ as the *inner* permutations. The distinguisher is adaptive, and can make both forward and backward queries to each permutation oracle, which corresponds to the notion of adaptive chosen-plaintext and ciphertext security (CCA). We consider computationally unbounded distinguishers, and we assume *wlog* that the distinguisher is deterministic and never makes useless queries (which means that it never repeats a query, nor makes a query $P_i^{-1}(y)$ if it received y as the answer to a previous query $P_i(x)$, or vice-versa).

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

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

where the first probability is taken over the random choice of K and $\boldsymbol{P}$, and the second probability is taken over the random choice of E and $\boldsymbol{P}$. We recall that, even though this is not apparent from the notation, the distinguisher can make both forward and backward queries to each permutation oracle.

For q_e, q_p non-negative integers, we define the insecurity of the ideal[6] Generalized Even-Mansour cipher with parameters $(n, r, m, \ell, \phi, \boldsymbol{\gamma})$ as:

$$\mathbf{Adv}^{\mathrm{cca}}_{\mathsf{EM}[n,r,m,\ell,\phi,\boldsymbol{\gamma}]}(q_e, q_p) = \max_{\mathcal{D}} \mathbf{Adv}(\mathcal{D}),$$

where the maximum is taken over all distinguishers $\mathcal{D}$ making exactly q_e queries to the outer permutation and exactly q_p queries to each inner permutation. The notation is adapted naturally for the two special cases EMIP and EMSP defined above.

The H-Coefficient Technique. We give here all the necessary background on the H-coefficient technique [31, 9] that we will use throughout this paper. All the information gathered by the distinguisher when interacting with the system of $m+1$ permutations can be summarized in what we call the *transcript* of the interaction, which is the ordered list of queries and answers received from the system (i, b, z, z'), where $i \in \{0, \ldots, m\}$ names the permutation being queried, b is a bit indicating whether this is a forward or backward query, $z \in \{0,1\}^n$ is the actual value queried and z' the answer. We say that a transcript is *attainable* (with respect to some fixed distinguisher $\mathcal{D}$) if there exists a tuple of permutations $(P_0, \ldots, P_m) \in (\mathcal{P}_n)^{m+1}$ such that the interaction of $\mathcal{D}$ with $(P_0, \ldots, P_m)$ yields this transcript (said otherwise, the probability to obtain this transcript in the "ideal" world is non-zero). In fact, an attainable transcript can be represented in a more convenient way that we will use in all the following. Namely, from the transcript we can build $m+1$ lists of directionless queries

[6] By ideal, we mean that this insecurity measure is defined in the Random Permutation Model for $P_1, \ldots, P_m$.

$\mathcal{Q}_E = ((x_1, y_1), \ldots, (x_{q_e}, y_{q_e}))$, $\mathcal{Q}_{P_1} = ((u_{1,1}, v_{1,1}), \ldots, (u_{1,q_p}, v_{1,q_p}))$, $\ldots, \mathcal{Q}_{P_m} = ((u_{m,1}, v_{m,1}), \ldots, (u_{m,q_p}, v_{m,q_p}))$ as follows. For $j = 1, \ldots, q_e$, let $(0, b, z, z')$ be the j-th query to P_0 in the transcript: if this was a forward query then we set $x_j = z$ and $y_j = z'$, otherwise we set $x_j = z'$ and $y_j = z$. Similarly, for each $i = 1, \ldots, m$, and $j = 1, \ldots, q_p$, let (i, b, z, z') be the j-th query to P_i in the transcript: if this was a forward query then we set $u_{i,j} = z$ and $v_{i,j} = z'$, otherwise we set $u_{i,j} = z'$ and $v_{i,j} = z$. A moment of thinking should make it clear that for attainable transcripts there is a one-to-one mapping between these two representations. (Essentially this follows from the fact that the distinguisher is deterministic). Moreover, though we defined $\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_m}$ as ordered lists, the order is unimportant (our formalization keeps the natural order induced by the distinguisher).

For convenience, and following [9], we will be generous with the distinguisher by providing it, at the end of its interaction, with the actual key K when it is interacting with $(\mathsf{EM}_K^{\boldsymbol{P}}, \boldsymbol{P})$, or with a dummy key K selected uniformly at random when it is interacting with $(E, \boldsymbol{P})$. This is without loss of generality since the distinguisher is free to ignore this additional information. Hence, all in all a transcript τ is a tuple $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_m}, K)$. We refer to $(\mathcal{Q}_E, \mathcal{Q}_{P_1}, \ldots, \mathcal{Q}_{P_m})$ (without the key) as the *permutation transcript*, and we say that a transcript τ is attainable if the corresponding permutation transcript is attainable. We denote $\mathcal{T}$ the set of attainable transcripts. (Thus $\mathcal{T}$ depends on $\mathcal{D}$, as the notion of attainability depends on $\mathcal{D}$.) 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.

In order to upper bound the advantage of the distinguisher, we will repeatedly use the following strategy: we will partition the set of attainable transcripts $\mathcal{T}$ into a set of "good" transcripts $\mathcal{T}_1$ such that the probabilities to obtain some transcript $\tau \in \mathcal{T}_1$ are close in the real and in the ideal world, and a set $\mathcal{T}_2$ of "bad" transcripts such that the probability to obtain any $\tau \in \mathcal{T}_2$ is small in the ideal world. More precisely, we will use the following result, which is proved in the full version of the paper [10].

Lemma 1. *Fix a distinguisher $\mathcal{D}$. Let $\mathcal{T} = \mathcal{T}_1 \sqcup \mathcal{T}_2$ be a partition of the set of attainable transcripts. Assume that there exists ε_1 such that for any $\tau \in \mathcal{T}_1$, one has*[7]

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

and that there exists ε_2 such that $\Pr[T_{\mathrm{id}} \in \mathcal{T}_2] \leq \varepsilon_2$. *Then* $\mathbf{Adv}(\mathcal{D}) \leq \varepsilon_1 + \varepsilon_2$.

3 A Sum-Capture Theorem

In this section, we prove a variant of previous "sum-capture" results [3, 24, 33]. Informally, such results typically state that when choosing a random subset A

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

of $\mathbb{Z}_2^n$ (or more generally any abelian group) of size q, the value

$$\mu(A) = \max_{\substack{U,V \subseteq \mathbb{Z}_2^n \\ |U|=|V|=q}} |\{(a,u,v) \in A \times U \times V : a = u \oplus v\}|$$

is close to its expected value q^3/N (if A, U, V were chosen at random), except with negligible probability. Here, we prove a result of this type for the setting where A arises from the interaction of an adversary with a random permutation P, namely $A = \{x \oplus y : (x,y) \in \mathcal{Q}\}$, where $\mathcal{Q}$ is the transcript of the interaction between the adversary and P. In fact our result is even more general, the special case just mentioned corresponding to Γ being the identity in the theorem below.

Theorem 1. *Fix an automorphism $\Gamma \in \mathsf{GL}(n)$. Let P be a uniformly random permutation of $\{0,1\}^n$, and let $\mathcal{A}$ be some probabilistic algorithm making exactly q (two-sided) adaptive queries to P. Let $\mathcal{Q} = ((x_1, y_1), \ldots, (x_q, y_q))$ denote the transcript of the interaction of $\mathcal{A}$ with P. For any two subsets U and V of $\{0,1\}^n$, let*

$$\mu(\mathcal{Q}, U, V) = |\{((x,y),u,v) \in \mathcal{Q} \times U \times V : x \oplus u = \Gamma(y \oplus v)\}|.$$

Then, assuming $9n \leq q \leq N/2$, one has

$$\Pr_{P,\omega}\left[\exists U, V \subseteq \{0,1\}^n : \mu(\mathcal{Q}, U, V) \geq \frac{q|U||V|}{N} + \frac{2q^2\sqrt{|U||V|}}{N} + 3\sqrt{nq|U||V|}\right] \leq \frac{2}{N},$$

where the probability is taken over the random choice of P and the random coins ω of $\mathcal{A}$.

Proof. Deferred to the full version [10] for reasons of space. □

4 Security Proof for the Single Permutation Case

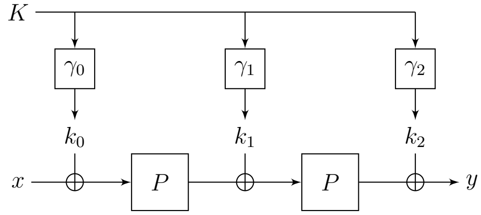

Fig. 3. The two-round Even-Mansour cipher with a single permutation and an arbitrary key-schedule.

In this section, we study the security of the two-round Even-Mansour construction where a single permutation P is used instead of two independent permutations, namely $\mathsf{EMSP}[n, r, \ell, \boldsymbol{\gamma}]$ (depicted on Figure 3). Because of the slide attack described in the full version of the paper [10], we know that we cannot simply use the same n-bit key k at each round if we aim at proving security beyond the birthday bound, so that some non-trivial key-schedule $\boldsymbol{\gamma} = (\gamma_0, \gamma_1, \gamma_2)$, with $\gamma_i : \{0,1\}^\ell \to \{0,1\}^n$, is needed (we remain as general as possible in a first phase, and will only specify the key-schedule later on). Given a key $K \in \{0,1\}^\ell$, we denote $k_0 = \gamma_0(K)$, $k_1 = \gamma_1(K)$, and $k_2 = \gamma_2(K)$, so that:

$$\mathsf{EMSP}_K^P(x) = P(P(x \oplus k_0) \oplus k_1) \oplus k_2.$$

Let $\tau = (\mathcal{Q}_E, \mathcal{Q}_P, K)$, with $|\mathcal{Q}_E| = q_e$, $|\mathcal{Q}_P| = q_p$, and $K \in \{0,1\}^\ell$ be an attainable transcript. As previously, we start by defining the set of bad transcripts. In all the following, we let

$$M = \frac{q_e}{N^{\frac{1}{3}}}.$$

Definition 1 (Bad transcript, single-permutation case). *We say that a transcript $\tau = (\mathcal{Q}_E, \mathcal{Q}_P, K) \in \mathcal{T}$ is* bad *if*

$$K \in \mathsf{BadK} = \bigcup_{1 \le i \le 10} \mathsf{BadK}_i$$

where

$$K \in \mathsf{BadK}_1 \Leftrightarrow \exists (x,y) \in \mathcal{Q}_E, \exists (u,v),(u',v') \in \mathcal{Q}_P : k_0 = x \oplus u \textit{ and } k_2 = v' \oplus y$$

$$K \in \mathsf{BadK}_2 \Leftrightarrow \exists (x,y) \in \mathcal{Q}_E, \exists (u,v),(u',v') \in \mathcal{Q}_P : k_0 = x \oplus u \textit{ and } k_1 = v \oplus u'$$

$$K \in \mathsf{BadK}_3 \Leftrightarrow \exists (x,y) \in \mathcal{Q}_E, \exists (u,v),(u',v') \in \mathcal{Q}_P : k_1 = v \oplus u' \textit{ and } k_2 = v' \oplus y$$

$$K \in \mathsf{BadK}_4 \Leftrightarrow \exists (x,y),(x',y') \in \mathcal{Q}_E, \exists (u,v) \in \mathcal{Q}_P : \\ k_0 = x \oplus u \textit{ and } k_0 \oplus k_1 = v \oplus x'$$

$$K \in \mathsf{BadK}_5 \Leftrightarrow \exists (x,y),(x',y') \in \mathcal{Q}_E, \exists (u,v) \in \mathcal{Q}_P : \\ k_1 \oplus k_2 = y' \oplus u \textit{ and } k_2 = v \oplus y$$

$$K \in \mathsf{BadK}_6 \Leftrightarrow |\{((x,y),(u,v)) \in \mathcal{Q}_E \times \mathcal{Q}_P : x \oplus u = k_0\}| > \frac{M}{3}$$

$$K \in \mathsf{BadK}_7 \Leftrightarrow |\{((x,y),(u,v)) \in \mathcal{Q}_E \times \mathcal{Q}_P : v \oplus y = k_2\}| > \frac{M}{3}$$

$$K \in \mathsf{BadK}_8 \Leftrightarrow |\{((x,y),(u,v)) \in \mathcal{Q}_E \times \mathcal{Q}_P : x \oplus v = k_0 \oplus k_1\}| > \frac{M}{3}$$

$$K \in \mathsf{BadK}_9 \Leftrightarrow |\{((x,y),(u,v)) \in \mathcal{Q}_E \times \mathcal{Q}_P : u \oplus y = k_1 \oplus k_2\}| > \frac{M}{3}$$

$$K \in \mathsf{BadK}_{10} \Leftrightarrow |\{((x,y),(x',y')) \in \mathcal{Q}_E \times \mathcal{Q}_E : x \oplus y' = k_0 \oplus k_1 \oplus k_2\}| > M.$$

Otherwise τ is said good*. We denote $\mathcal{T}_2$ the set of bad transcripts, and $\mathcal{T}_1 = \mathcal{T} \backslash \mathcal{T}_2$ the set of good transcripts.*

In this section, we focus on the case where $\ell = n$, namely the master key length is equal to the block length (and hence to the round keys length). We treat the (simpler) cases where the three round keys are independent, or derived from two independent n-bit keys, in the full version of the paper [10]. First, we specify conditions on the key-schedule that will allow us to upper bound the probability to obtain a bad transcript (in the ideal world).

Definition 2 (Good key-schedule). *We say that a key-schedule $\boldsymbol{\gamma} = (\gamma_0, \gamma_1, \gamma_2)$, where $\gamma_i : \{0,1\}^n \to \{0,1\}^n$, is* good *if it satisfies the following conditions:*

(i) $\gamma_0, \gamma_1, \gamma_2 \in \mathsf{GL}(n)$ *(i.e., each γ_i is a linear bijective map of $\mathbb{F}_2^n$);*
(ii) $\gamma_0 \oplus \gamma_1 \in \mathsf{GL}(n)$ *and* $\gamma_1 \oplus \gamma_2 \in \mathsf{GL}(n)$*;*
(iii) $\gamma_0 \oplus \gamma_1 \oplus \gamma_2$ *is a permutation over $\{0,1\}^n$ (non-necessarily linear over $\mathbb{F}_2^n$).*

A simple way to build a good key-schedule is to take for γ_0 and γ_2 the identity, and $\gamma_1 = \pi$, where π is a linear orthomorphism of $\mathbb{F}_2^n$ (recall that a permutation π of $\{0,1\}^n$ is an orthomorphism if $x \mapsto x \oplus \pi(x)$ is also a permutation), so that the sequence of round keys is $(k, \pi(k), k)$. We give two simple examples of linear orthomorphisms which are attractive from an implementation point of view:

- When n is even, and $k = (k_L, k_R)$ where k_L and k_R are respectively the left and right halves of k, then
$$\pi : (k_L, k_R) \mapsto (k_R, k_L \oplus k_R)$$
is a linear orthomorphism.
- Fix an irreducible polynomial p of degree n over $\mathbb{F}_2$ and identify $\mathbb{F}_2^n$ and the extension field $\mathbb{F}_{2^n}$ defined by p in the canonical way. Then, for any $c \in \mathbb{F}_{2^n} \backslash \{0,1\}$, $k \mapsto c \odot k$ (where $\odot$ denotes the extension field multiplication) is a linear orthomorphism.

Lemma 2. *Let $\boldsymbol{\gamma} = (\gamma_0, \gamma_1, \gamma_2)$ be a good key-schedule. Assume that $9n \le q_e, q_p \le N/2$. Then*

$$\Pr[T_{\mathrm{id}} \in \mathcal{T}_2] \le \frac{10}{N} + \frac{4q_e^2 q_p + 7q_e q_p^2 + 4q_p^2\sqrt{q_e q_p}}{N^2} + \frac{9q_p\sqrt{nq_e} + 6q_e\sqrt{nq_p}}{N} + \frac{q_e + 12q_p}{N^{\frac{2}{3}}}.$$

Proof. In the ideal world, sets BadK_i only depend on the random permutations E and P, and not on the key k, which is drawn uniformly at random at the end of the interaction of the distinguisher with (E, P). Moreover, the size of BadK_i for $i = 6$ to 10 can be upper bounded independently of E, P. Indeed, since γ_0, γ_2, $\gamma_0 \oplus \gamma_1$, $\gamma_1 \oplus \gamma_2$, and $\gamma_0 \oplus \gamma_1 \oplus \gamma_2$ are all permutations of $\{0,1\}^n$, one has, for any permutation transcript $(\mathcal{Q}_E, \mathcal{Q}_P)$,

$$|\mathsf{BadK}_6|, |\mathsf{BadK}_7|, |\mathsf{BadK}_8|, |\mathsf{BadK}_9| \le \frac{3q_e q_p}{M} \quad \text{and} \quad |\mathsf{BadK}_{10}| \le \frac{q_e^2}{M},$$

so that

$$\Pr\left[k \leftarrow_{\$} \{0,1\}^n : k \in \bigcup_{i=6}^{10} \mathsf{BadK}_i\right] \leq \frac{12q_eq_p}{NM} + \frac{q_e^2}{NM} \leq \frac{q_e + 12q_p}{N^{\frac{2}{3}}}.$$

On the other hand, in order to upper bound $|\mathsf{BadK}_i|$ for $i = 1$ to 5, we need to appeal to the sum-capture theorem of Section 3. For a permutation transcript $(\mathcal{Q}_E, \mathcal{Q}_P)$, let

$$\begin{aligned} X &= \{x \in \{0,1\}^n : (x,y) \in \mathcal{Q}_E\}, & Y &= \{y \in \{0,1\}^n : (x,y) \in \mathcal{Q}_E\}, \\ U &= \{u \in \{0,1\}^n : (u,v) \in \mathcal{Q}_P\}, & V &= \{v \in \{0,1\}^n : (u,v) \in \mathcal{Q}_P\} \end{aligned}$$

denote the domains and the ranges of $\mathcal{Q}_E$ and $\mathcal{Q}_P$, respectively. Then one has

$$\begin{aligned} |\mathsf{BadK}_1| &\leq \mu(\mathcal{Q}_E, U, V) \\ &\overset{\text{def}}{=} |\{((x,y),u,v) \in \mathcal{Q}_E \times U \times V : x \oplus u = \gamma_0 \circ \gamma_2^{-1}(y \oplus v)\}| \\ |\mathsf{BadK}_2| &\leq \mu(\mathcal{Q}_P, X, U) \\ &\overset{\text{def}}{=} |\{((u,v),x,u') \in \mathcal{Q}_P \times X \times U : x \oplus u = \gamma_0 \circ \gamma_1^{-1}(v \oplus u')\}| \\ |\mathsf{BadK}_3| &\leq \mu(\mathcal{Q}_P, V, Y) \\ &\overset{\text{def}}{=} |\{((u',v'),v,y) \in \mathcal{Q}_P \times V \times Y : v \oplus u' = \gamma_1 \circ \gamma_2^{-1}(v' \oplus y)\}| \\ |\mathsf{BadK}_4| &\leq \mu(\mathcal{Q}_P, X, X) \\ &\overset{\text{def}}{=} |\{((u,v),x,x') \in \mathcal{Q}_P \times X \times X : x \oplus u = \gamma_0 \circ (\gamma_0 \oplus \gamma_1)^{-1}(v \oplus x')\}| \\ |\mathsf{BadK}_5| &\leq \mu(\mathcal{Q}_P, Y, Y) \\ &\overset{\text{def}}{=} |\{((u,v),y,y') \in \mathcal{Q}_P \times Y \times Y : y' \oplus u = (\gamma_1 \oplus \gamma_2) \circ \gamma_2^{-1}(v \oplus y)\}|. \end{aligned}$$

By our assumption that the key-schedule is good, we have that $\gamma_0 \circ \gamma_2^{-1}$, $\gamma_0 \circ \gamma_1^{-1}$, $\gamma_1 \circ \gamma_2^{-1}$, $\gamma_0 \circ (\gamma_0 \oplus \gamma_1)^{-1}$, and $\gamma_0 \circ (\gamma_0 \oplus \gamma_1)^{-1}$ are all automorphisms of $\mathbb{F}_2^n$. Hence, we can apply Theorem 1 (note that in order to apply this theorem to upper bound, say, $|\mathsf{BadK}_1|$, we consider the combination of the distinguisher $\mathcal{D}$ and permutation P as a probabilistic adversary $\mathcal{A}$ interacting with permutation E, resulting in transcript $\mathcal{Q}_E$). Thus, if we set

$$\begin{aligned} C_1 &= \frac{q_eq_p^2}{N} + \frac{2q_e^2q_p}{N} + 3q_p\sqrt{nq_e} \\ C_2 = C_3 &= \frac{q_eq_p^2}{N} + \frac{2q_p^2\sqrt{q_eq_p}}{N} + 3q_p\sqrt{nq_e} \\ C_4 = C_5 &= \frac{q_e^2q_p}{N} + \frac{2q_eq_p^2}{N} + 3q_e\sqrt{nq_p}, \end{aligned}$$

one has $\Pr[E, P \leftarrow_{\$} \mathcal{P}_n : |\mathsf{BadK}_i| \geq C_i] \leq 2/N$ for each $i = 1$ to 5. Since

$$\Pr[T_{\text{id}} \in \mathcal{T}_2] \leq \sum_{i=1}^{5} \Pr[E, P \leftarrow_{\$} \mathcal{P}_n : |\mathsf{BadK}_i| \geq C_i] + \frac{\sum_{i=1}^5 C_i}{N} + \frac{q_e + 12q_p}{N^{\frac{2}{3}}},$$

we get the final result. □

In the second stage of the proof, it remains to show that for any good transcript τ, the ratio between the probabilities to obtain τ in the ideal world and the real world is close to 1. We have the following lemma, proved in the full version of the paper [10].

Lemma 3. *Assume that $N \geq 7^3$ and $4q_e + 2q_p \leq N$. Let $\tau = (\mathcal{Q}_E, \mathcal{Q}_P, K) \in \mathcal{T}_1$ be a good transcript. Then*

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

where

$$\varepsilon_1 = \frac{4q_e(q_e + q_p)^2}{N^2} + \frac{2q_e^2}{N^{\frac{4}{3}}} + \frac{20q_e}{N^{\frac{2}{3}}}.$$

Combining Lemmas 1, 2, and 3, we obtain the main theorem of this paper.

Theorem 2 (Single permutation and non-independent round keys). *Consider the single-permutation two-round Even-Mansour cipher $\mathsf{EMSP}[n, 2, \boldsymbol{\gamma}]$ with a good key-schedule $\boldsymbol{\gamma}$ (see Definition 2). Assume that $N \geq 7^3$, $9n \leq q_e, q_p \leq N/2$, and $4q_e + 2q_p \leq N$. Then*

$$\begin{aligned}\mathbf{Adv}^{\mathrm{cca}}_{\mathsf{EMSP}[n,2,\boldsymbol{\gamma}]}(q_e, q_p) \leq \frac{10}{N} &+ \frac{4q_e^3 + 12q_e^2 q_p + 11q_e q_p^2 + 4q_p^2\sqrt{q_e q_p}}{N^2} + \frac{2q_e^2}{N^{\frac{4}{3}}} \\ &+ \frac{9q_p\sqrt{nq_e} + 6q_e\sqrt{nq_p}}{N} + \frac{21q_e + 12q_p}{N^{\frac{2}{3}}}.\end{aligned}$$

Letting $q = \max(q_e, q_p)$, and assuming $q \leq N^{\frac{2}{3}}$, the upper bound of Theorem 2 simplifies into

$$\frac{10}{N} + \frac{31q^3}{N^2} + \frac{2q^2}{N^{\frac{4}{3}}} + \frac{15\sqrt{n}q^{\frac{3}{2}}}{N} + \frac{33q}{N^{\frac{2}{3}}} \leq \frac{10}{N} + \frac{81\sqrt{n}q}{N^{\frac{2}{3}}} = \frac{10}{2^n} + \frac{81q}{2^{\frac{2n}{3} - \frac{1}{2}\log_2 n}}.$$

Hence, security is ensured up to $\mathcal{O}(2^{\frac{2n}{3} - \frac{1}{2}\log_2 n}) = \widetilde{\mathcal{O}}(2^{\frac{2n}{3}})$ queries of the adversary.

Acknowledgments

Shan and John were supported by National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003, and by the China Ministry of Education grant number 20121088050. Rodolphe was supported by the French Direction Générale de l'Armement and the French National Agency of Research through the PRINCE project (contract ANR-10-SEGI-015). Jooyoung was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2013R1A1A2007488). Yannick was partially supported by the French National Agency of Research through the BLOC project (contract ANR-11-INS-011).

References

[1] N. Alon, T. Kaufman, M. Krivelevich, and D. Ron. Testing Triangle-Freeness in General Graphs. *SIAM J. Discrete Math.*, 22(2):786–819, 2008.

[2] E. Andreeva, A. Bogdanov, Y. Dodis, B. Mennink, and J. P. Steinberger. On the Indifferentiability of Key-Alternating Ciphers. In *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.

[3] L. Babai. The Fourier Transform and Equations over Finite Abelian Groups: An introduction to the method of trigonometric sums. Lecture notes, December 1989. Available at http://people.cs.uchicago.edu/~laci/reu02/fourier.pdf.

[4] E. Biham, Y. Carmeli, I. Dinur, O. Dunkelman, N. Keller, and A. Shamir. Cryptanalysis of Iterated Even-Mansour Schemes with Two Keys. IACR Cryptology ePrint Archive, Report 2013/674, 2013. Available at http://eprint.iacr.org/2013/674.

[5] A. Biryukov and D. Wagner. Slide Attacks. In *Fast Software Encryption - FSE '99*, volume 1636 of *LNCS*, pages 245–259. Springer, 1999.

[6] A. Biryukov and D. Wagner. Advanced Slide Attacks. In *Advances in Cryptology - EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 589–606. Springer, 2000.

[7] A. Bogdanov, L. R. Knudsen, G. Leander, C. Paar, A. Poschmann, M. J. Robshaw, Y. Seurin, and C. Vikkelsoe. PRESENT: An Ultra-Lightweight Block Cipher. In *Cryptographic Hardware and Embedded Systems - CHES 2007*, volume 4727 of *LNCS*, pages 450–466. Springer, 2007.

[8] A. Bogdanov, L. R. Knudsen, G. Leander, F.-X. Standaert, J. P. Steinberger, and E. Tischhauser. Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations - (Extended Abstract). In *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62. Springer, 2012.

[9] S. Chen and J. Steinberger. Tight Security Bounds for Key-Alternating Ciphers. In *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, 2014.

[10] S. Chen, R. Lampe, J. Lee, Y. Seurin, and J. Steinberger. Minimizing the Two-Round Even-Mansour Cipher. Full version of this paper. Available at http://eprint.iacr.org/2014/443.

[11] J. Daemen. Limitations of the Even-Mansour Construction. In *Advances in Cryptology - ASIACRYPT '91*, volume 739 of *LNCS*, pages 495–498. Springer, 1991.

[12] J. Daemen and V. Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer, 2002.

[13] J. Daemen and V. Rijmen. Probability Distributions of Correlations and Differentials in Block Ciphers. ePrint Archive, Report 2005/212, 2005. Available at http://eprint.iacr.org/2005/212.pdf.

[14] I. Dinur, O. Dunkelman, N. Keller, and A. Shamir. Key Recovery Attacks on 3-round Even-Mansour, 8-step LED-128, and Full AES^2. In *Advances in Cryptology - ASIACRYPT 2013 (Proceedings, Part I)*, volume 8269 of *LNCS*, pages 337–356. Springer, 2013. Full version available at http://eprint.iacr.org/2013/391.

[15] O. Dunkelman, N. Keller, and A. Shamir. Minimalism in Cryptography: The Even-Mansour Scheme Revisited. In *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 336–354. Springer, 2012.

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

[17] P. Gazi. Plain versus Randomized Cascading-Based Key-Length Extension for Block Ciphers. In *Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I)*, volume 8042 of *LNCS*, pages 551–570. Springer, 2013.

[18] P. Gazi and S. Tessaro. Efficient and Optimally Secure Key-Length Extension for Block Ciphers via Randomized Cascading. In *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 63–80. Springer, 2012.

[19] S. W. Golomb, G. Gong, and L. Mittenthal. Constructions of Orthomorphisms of $\mathbb{Z}_n^2$. In *Proceedings of The Fifth International Conference on Finite Fields and Applications*, pages 178–195. Springer, 1999.

[20] J. Guo, T. Peyrin, A. Poschmann, and M. J. B. Robshaw. The LED Block Cipher. In *Cryptographic Hardware and Embedded Systems - CHES 2011*, volume 6917 of *LNCS*, pages 326–341. Springer, 2011.

[21] T. P. Hayes. A Large-Deviation Inequality for Vector-Valued Martingales. Manuscript, 2005. Available at http://www.cs.unm.edu/~hayes/papers/VectorAzuma.

[22] P. Junod and S. Vaudenay. FOX : A New Family of Block Ciphers. In *Selected Areas in Cryptography - SAC 2004*, volume 3357 of *LNCS*, pages 114–129. Springer, 2004.

[23] J. Kilian and P. Rogaway. How to Protect DES Against Exhaustive Key Search (an Analysis of DESX). *Journal of Cryptology*, 14(1):17–35, 2001.

[24] E. Kiltz, K. Pietrzak, and M. Szegedy. Digital Signatures with Minimal Overhead from Indifferentiable Random Invertible Functions. In *Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I)*, volume 8042 of *LNCS*, pages 571–588. Springer, 2013.

[25] X. Lai and J. L. Massey. A Proposal for a New Block Encryption Standard. In *Advances in Cryptology - EUROCRYPT '90*, volume 473 of *LNCS*, pages 389–404. Springer, 1990.

[26] R. Lampe, J. Patarin, and Y. Seurin. An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher. In *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 278–295. Springer, 2012.

[27] R. Lampe and Y. Seurin. How to Construct an Ideal Cipher from a Small Set of Public Permutations. In *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.

[28] J. Lee. Towards Key-Length Extension with Optimal Security: Cascade Encryption and Xor-cascade Encryption. In *Advances in Cryptology - EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 405–425. Springer, 2013.

[29] L. Mittenthal. Block Substitutions Using Orthomorphic Mappings. *Advances in Applied Mathematics*, 16(1):59–71, 1995.

[30] I. Nikolica, L. Wang, and S. Wu. Cryptanalysis of Round-Reduced LED. In *Fast Software Encryption - FSE 2013*, 2013. To appear.

[31] J. Patarin. The "Coefficients H" Technique. In *Selected Areas in Cryptography - SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, 2008.

[32] J. 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.

[33] J. Steinberger. Counting solutions to additive equations in random sets. arXiv Report 1309.5582, 2013. Available at http://arxiv.org/abs/1309.5582.

[34] S. Vaudenay. On the Lai-Massey Scheme. In *Advances in Cryptology - ASIACRYPT '99*, volume 1716 of *LNCS*, pages 8–19. Springer, 1999.