Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations

Andrey Bogdanov[1], Lars R. Knudsen[2], Gregor Leander[2], Francois-Xavier Standaert[3], John Steinberger[4], and Elmar Tischhauser[1]

[1] KU Leuven and IBBT {Andrey.Bogdanov,Elmar.Tischhauser}@esat.kuleuven.be
[2] Technical University of Denmark
{G.Leander,Knudsen}@mat.dtu.dk
[3] Université catholique de Louvain, UCL Crypto Group
fstandae@uclouvain.be
[4] Tsinghua University
jpsteinb@gmail.com

Abstract. This paper considers—for the first time—the concept of key-alternating ciphers in a provable security setting. Key-alternating ciphers can be seen as a generalization of a construction proposed by Even and Mansour in 1991. This construction builds a block cipher PX from an n-bit permutation P and two n-bit keys k_0 and k_1, setting $PX_{k_0,k_1}(x) = k_1 \oplus P(x \oplus k_0)$. Here we consider a (natural) extension of the Even-Mansour construction with t permutations $P_1, \ldots, P_t$ and $t+1$ keys, $k_0, \ldots, k_t$. We demonstrate in a formal model that such a cipher is secure in the sense that an attacker needs to make at least $2^{2n/3}$ queries to the underlying permutations to be able to distinguish the construction from random. We argue further that the bound is tight for $t = 2$ but there is a gap in the bounds for $t > 2$, which is left as an open and interesting problem. Additionally, in terms of statistical attacks, we show that the distribution of Fourier coefficients for the cipher over all keys is close to ideal. Lastly, we define a practical instance of the construction with $t = 2$ using AES referred to as AES^2. Any attack on AES^2 with complexity below 2^{85} will have to make use of AES with a fixed known key in a non-black box manner. However, we conjecture its security is 2^{128}.

Keywords: Block ciphers, provable security, Even-Mansour construction, AES

1 Introduction

Block ciphers are one of the fundamental primitives in symmetric cryptography. Often called the work horses of cryptography, they form the backbone of today's secure communication. Therefore, their design has been an important research focus over the last 20 years, giving rise to different well-established strategies to prevent large classes of attacks. As typical examples, one can mention the practical security approach against linear and differential cryptanalysis [22], and

the wide-trail strategy [14] that lead to the design of the AES Rijndael [13]. Another line of research is the so-called provable security approach against statistical attacks, that served as foundation for the block cipher MISTY [26, 27]. One can also mention the decorrelation theory [32] and the design of the ciphers C [1] and KFC [2]. At a high level, the three main design paradigms for block ciphers are Feistel structures such as DES, Lai-Massey ciphers such as IDEA [23], and key-alternating ciphers [11, 13, 14] for which the AES Rijndael is a prominent representative. State-of-the-art block ciphers are quite well understood and provide security against all known attacks. Though there has recently been remarkable progress in the cryptanalysis of AES [7], these results are far from being any threat for the use of AES in practice. Thus, from a practical point of view, block ciphers in general and key-alternating ciphers in particular can be seen as a success story.

Given the degree of confidence in properly designed key-alternating ciphers on the practical side (e.g. with AES approved for the encryption of secret and top secret data in the USA), it is even more surprising that there has been no provable setting developed so far for the design of key-alternating ciphers on the theoretical side. Nobody seems to have even formulated the problem of whether the key-alternating cipher makes sense from this point of view. Clearly, given the state of the art, proving AES secure in any strict sense is out of reach. However, by modeling the round functions as fixed public randomly chosen permutations, we are able to precisely formulate and—as we shall see—prove the soundness of the key-alternating cipher design. The cipher we are dealing with is depicted in Figure 2 and detailed in Section 2.

We note the difference of our setting to that of an idealized Feistel cipher, often called the Luby-Rackoff construction [25], or to that of similar results obtained for the Lai-Massey schemes [33]. In these former works, for each key it is assumed that the function used in the Feistel (resp. Lai-Massey) construction is chosen at random. Directly adopting this model to the case of a key-alternating cipher immediately results in an ideal cipher (even for one round). At the same time, in most key-alternating ciphers including AES, the key is the only part of the design to define the cipher permutation and all round permutations are fixed for the entire cipher, not varying from key to key. In other words, working along the lines of [25] does not elucidate how to mix the key into the state. It is exactly this point we deal with in the present paper, both at a high-level, i.e. in a provable setting, as well as at lower-levels, i.e. considering statistical attacks and as a guideline for actually designing ciphers.

Interestingly, another look at the construction and its properties arises from the question of how to design the key schedule of a block cipher. This has been an open problem in symmetric-key cryptography for decades. While some ciphers are based upon simple linear or nearly linear key schedules [8, 17], a number of others opt for heavier and often highly nonlinear key schedules, sometimes as complex as the round functions [3] or the cipher itself [30]. In the prominent case of AES, for instance, the key schedule is iterative, mainly linear, and provides

relatively slow diffusion in the backward direction. It is precisely these properties that facilitated the related-key cryptanalysis of the full AES-192 and AES-256, e.g. [5,6] as well as the recent biclique cryptanalysis of all three full AES versions in the classical single-key model [7]. In general, these examples emphasize a relatively weak understanding of key scheduling algorithms, compared to the design of block cipher rounds. In this context, the results of this paper can be seen as a case for simple key schedules (or even no key scheduling at all). Hence, they provide new insights into the design of block ciphers.

1.1 Related Work

An exception from the above-mentioned lack of theoretical studies of key-alternating block ciphers is the Even-Mansour construction [15] depicted in Figure 1. This

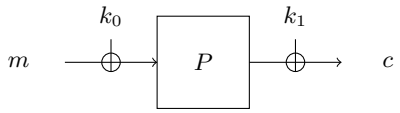

Fig. 1. The Even-Mansour construction

construction can be seen as a one-round variant of a key-alternating cipher. Informally, Even and Mansour proved that in order to have a reasonable success probability in decrypting an (unqueried) message, an attacker has to make roughly $2^{n/2}$ queries to the permutation P. In this setting, the attacker is given oracle access to P, its inverse, and to an encryption and decryption oracle. Later, Daemen [10] showed that this bound is actually tight. He presented a differential attack on the Even-Mansour scheme that allows to successfully recover the key with a good probability, after $2^{n/2}$ evaluations of both the permutation P and the encryption oracle.

1.2 Our Contribution

Our contributions in this paper are twofold.

On the theoretical side (cf. Section 3), we provide the first treatment of the concept of key-alternating ciphers in a provable security setting. We prove below that, for any t-round version of the cipher with randomly drawn and fixed underlying permutations, $t \geq 2$, depicted in Figure 2, an attacker needs to make at least $2^{2n/3}$ queries before being able to distinguish the encryption oracle from a random permutation. Here n is the block size of the cipher. Furthermore, we provide a simple attack that shows that an attacker, by making $2^{\frac{t}{t+1}n}$ queries, is able to recover the secret key used in the decryption oracle. We do conjecture that this lower bound — being tight only for $t = 2$ — is the actual bound. We

leave proving this as an important open question (see also Section 7). Note that in this setup, we necessarily only consider the query complexity of an attacker, ignoring the computational complexity. It seems unlikely that an attack with a comparable computational complexity exists. Such an attack would in particular imply an attack on e.g. AES-256 with a complexity of around 2^{120} operations.

On the practical side, we propose to actually use the construction of Figure 2. Given our theoretical results, the merit of this approach is the following: Any attack on a key-alternating cipher with complexity below $2^{2n/3}$ will have to make use of the round functions in a non-black box manner.

However, and we feel that it is important to make this point explicit even though it might be obvious, the theoretical result does not carry over to any efficient instance, as one must consider the round functions as black-boxes—i.e. objects which the adversary must query to evaluate—in order to meaningfully discuss the distinguishability of the cipher from a random permutation by an information-theoretic adversary.

This fact and the fact that, as mentioned above, the theoretical bounds are likely to be lower than the computational complexity of any attack, motivates us to study the security of our proposal with respect to such statistical attacks as linear cryptanalysis (see Section 5).

To capture the difference between the single-round Even-Mansour cipher and the multiple-round key-alternating construction with respect to linear cryptanalysis, we study the Fourier spectrum of the ciphers. We prove that once the fixed underlying permutations are close to average (which is the case for randomly drawn permutations with high probability), the distribution of Fourier coefficients for the key-alternating cipher over all keys for $t \geq 2$ gets close to that over all permutations — the natural reference point for any block cipher. At the same time, we demonstrate that this is not the case for the original Even-Mansour construction with $t = 1$ where the Fourier coefficients almost do not change from key to key. It seems therefore unlikely that linear attacks are able to break the multiple-round key-alternating cipher with $t \geq 2$.

Finally, as the crypto community likes targets and we anticipate that having a concrete proposal is a valuable stimulation for further research, we propose an actual cipher called AES^2 following the 2-round version of the general construction (see Section 6). Here we replace the random permutations by two instantiations of AES-128 with fixed known keys. Given the new AES instructions on recent Intel processors, AES^2 performs very competitively on those platforms, with as few as 2.65 cycles per byte required in the counter mode.

We conclude with a section dedicated to open questions and further work (Section 7), discussing how to possibly improve and extend the research we consider in the paper.

2 The Construction

The cipher we consider is an idealized model of a key-alternating cipher — the notion introduced under this name in [13, 14] in connection with the design of

AES and used without being explicitly named even before that [11] in similar contexts. Such a cipher consists of round functions interleaved with xoring round keys to the current state. In our idealized model, the round functions are the public, randomly chosen permutations P_i and the key consists of $t+1$ independent round-keys are k_i. More precisely, let $P_1, \ldots, P_t$ be permutations from $\{0,1\}^n$ to $\{0,1\}^n$, $t \geq 1$. Let $k_0, \ldots, k_t \in \{0,1\}^n$ be keys. The block cipher $E = E_{k_0,\ldots,k_t} : \{0,1\}^n \to \{0,1\}^n$ we consider is defined by

$$E(x) = E_{k_0 \cdots k_t}(x) = P_t(\ldots P_2(P_1(x \oplus k_0) \oplus k_1)\ldots) \oplus k_t \tag{1}$$

for $x \in \{0,1\}^n$. The cipher is shown in Figure 2.

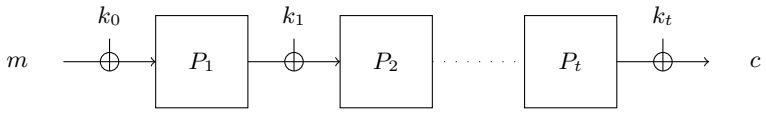

Fig. 2. A key-alternating cipher

3 Indistinguishability Analysis

Putting $N = 2^n$, we define the PRP security of E against an adversary A expecting a $(t+1)$-tuple of oracles as

$$\mathbf{Adv}^{\mathrm{PRP}}_{E,N,t}(A) = \Pr[k_0 \cdots k_t \leftarrow \{0,1\}^n; A^{E_{k_0 \cdots k_t}, P_1, \ldots, P_t} = 1] - \Pr[A^{Q, P_1, \ldots, P_t} = 1]$$

where in each experiment $Q, P_1, \ldots, P_t$ are independent and uniformly sampled random permutations. Here A can make inverse queries to each of its oracles. Thus, an attacker has to tell apart two worlds, depicted below.

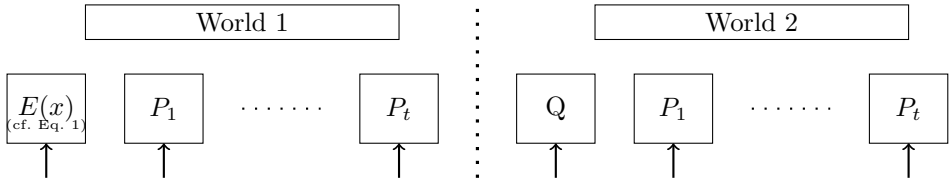

We note that one *must* consider the permutations $P_1, \ldots, P_t$ as random (or pseudorandom) black-boxes—i.e. objects which the adversary must query to evaluate—in order to meaningfully discuss the distinguishability of $E_{k_0,\ldots,k_t}$ from a random permutation by an information-theoretic adversary.

We define

$$\mathbf{Adv}^{\mathrm{PRP}}_{E,N,t}(q) = \max_A \mathbf{Adv}^{\mathrm{PRP}}_E(A)$$

where the maximum is taken over all adversaries A making at most q queries. (We note the parameters n and t are elided from both of the notations $\mathbf{Adv}^{\mathrm{PRP}}_{E}(A)$ and $\mathbf{Adv}^{\mathrm{PRP}}_{E}(q)$; but it should be understood that $\mathbf{Adv}^{\mathrm{PRP}}_{E}(q)$ is a function n and t as well as of q.)

Our main security result is the following:

Theorem 1. *Let $N = 2^n$ and let $q = N^{\frac{t}{t+1}}/Z$ for some $Z \geq 1$. Then, for any $t \geq 1$, and assuming $q < N/100$, we have*

$$\boldsymbol{Adv}^{\mathrm{PRP}}_{E,N,t}(q) \leq \frac{4.3q^3t}{N^2} + \frac{t+1}{Z^t}.$$

For $t \geq 2$ the limiting term in the above bound is $4q^3t/N^2$, which caps q at around $N^{2/3}$. The following corollary is more telling.

Corollary 1. *Assume $t \geq 2$. Let $q = N^{\frac{2}{3}}/\lambda\sqrt[3]{t}$ for some $\lambda \geq 1$. Then, assuming $q < N/100$,*

$$\boldsymbol{Adv}^{\mathrm{PRP}}_{E,N,t}(q) \leq \frac{4.3}{\lambda^3} + \frac{t+1}{(\sqrt[3]{t}\lambda)^t}.$$

We also note that $q < N/100$ as long as $n \geq 20$; this condition is therefore compatible with practical parameters. We note that Corollary 1's security of $q \approx N^{\frac{2}{3}}$ is optimal for $t = 2$ (cf. Section 3.1) and suboptimal for $t > 2$, in which case we conjecture a security of $q \approx N^{\frac{t}{t+1}}$. Closing this gap might be obtained by a tightening of Proposition 2 below.

Theorem 1 is proved by a hybrid argument involving an intermediate game. In order to outline this hybrid argument we start by developing some new notation.

Note firstly that if E is defined as in (1) then, putting $P_0 = E^{-1}$, we have

$$P_0(P_t(\cdots P_1(\cdot \oplus k_0)\cdots) \oplus k_t) = id.$$

Applying P_0^{-1} to both sides and then substituting $P_0(\cdot)$ for the input, we find

$$P_t(\cdots P_2(P_1(P_0(\cdot) \oplus k_0) \oplus k_1)\cdots) \oplus k_t = id. \tag{2}$$

It is easy to see that, for fixed $k_0, \ldots, k_t$, randomly sampling $P_1, \ldots, P_t$, defining E as in (1) and giving an adversary access to the tuple of oracles $(E, P_1, \ldots, P_t)$ (and their inverses) is equivalent to sampling $P_0, \ldots, P_t$ uniformly at random from all $(t+1)$-tuples of permutations satisfying (2) and giving the adversary access to $(P_0^{-1}, P_1, \ldots, P_t)$ (and their inverses). Moreover, it is just a notational change to give the adversary access to $(P_0, P_1, \ldots, P_t)$, since the adversary is allowed inverse queries anyway (of course, the adversary is alerted to the fact that its first oracle is now P_0 and not P_0^{-1}).

We now formally implement the interface $(P_0, \ldots, P_t)$ via an oracle $O(N, t)$ taking $k_0, \ldots, k_t$ as implicit parameters. Rather than sampling $P_0, \ldots, P_t$ uniformly at random from those sequences satisfying (2) at the start of the experiment, $O(N, t)$ implements the permutations $P_0, \ldots, P_t$ by lazy sampling. More

precisely, $P_0, \ldots, P_t$ are initially set to be undefined everywhere. When the adversary makes a query $P_i(x)$ or $P_i^{-1}(y)$, the adversary defines P_i at the relevant point using the following procedure, illustrated for the case of a forward query $P_i(x)$ (the case of a backward query is analogous):

- Let $\mathcal{P} = \mathcal{P}(P_0, \ldots, P_t)$ be the set of all $(t+1)$-tuples of permutations $(\overline{P}_0, \ldots, \overline{P}_t)$ such that $\overline{P}_i$ extends the currently defined portion of P_i, and such that

$$\overline{P}_t(\cdots \overline{P}_2(\overline{P}_1(\overline{P}_0(\cdot) \oplus k_0) \oplus k_1) \cdots \oplus k_{t-1}) \oplus k_t = id. \tag{3}$$

  Then $O(N,t)$ samples uniformly at random an element $(\overline{P}_0, \ldots, \overline{P}_t)$ from $\mathcal{P}$. The adversary sets $P_i(x) = \overline{P}_i(x)$ and returns this value.

After the above, the adversary "forgets" about $\overline{P}_0, \ldots, \overline{P}_t$, and samples these afresh at the next query. It is clear that this lazy sampling process gives the same distribution as sampling the tuple $(P_0, \ldots, P_t)$ at the start of the game. Thus, giving the adversary oracle access to $O(N,t)$ is equivalent to giving the adversary oracle access to $(E, P_1, \ldots, P_t)$, up to the cosmetic change that E is replaced by E^{-1}. We therefore have:

Proposition 1. *With $O(N,t)$ defined as above, we have:*

$$\boldsymbol{Adv}^{\mathrm{PRP}}_{E,N,t}(A) = \Pr[k_0 \cdots k_t \leftarrow \{0,1\}^n; A^{O(N,t)} = 1] - \Pr[A^{Q_0, Q_1, \ldots, Q_t} = 1]$$

where $Q_0, \ldots, Q_t$ are independent random permutations.

(We emphasize that $k_0, \ldots, k_t$ are implicit arguments to $O(N,t)$.)

Our hybrid will be an oracle $\tilde{O}(N,t)$ (also taking $k_0, \ldots, k_t$ as implicit inputs) that uses a slightly different lazy sampling procedure to define the permutations $P_0, \ldots, P_t$. Say that a sequence of partially defined permutations is *consistent* if $\mathcal{P}(P_0, \ldots, P_t) \neq \emptyset$, with $\mathcal{P}(\cdot)$ defined as in the description of $O(N,t)$ above. Initially, $\tilde{O}(N,t)$ also sets the permutations $P_0, \ldots, P_t$ to be undefined everywhere. Upon receiving (say) a forward query $P_i(x)$, $\tilde{O}(N,t)$ uses the following lazy sampling procedure to answer:

- Let $U \subseteq \{0,1\}^n$ be the set of values y such that defining $P_i(x) = y$ maintains the consistency of $P_0, \ldots, P_t$, besides maintaining the fact that P_i is a permutation. Then $\tilde{O}(N,t)$ samples a value y uniformly from U, sets $P_i(x) = y$, and returns y.

Inverse queries are lazy sampled the same way. While not immediately apparent, the above lazy sampling procedure produces a slightly *different* distribution of outputs than the first lazy sampling procedure.

Theorem 1 is an direct consequence of Proposition 1 and of the following two propositions.

Proposition 2. *Let $q < N/100$. With $O(N,t)$ and $\tilde{O}(N,t)$ defined as above,*

$$\Pr[k_0, \ldots, k_t \leftarrow \{0,1\}^n; A^{O(N,t)} = 1] - \Pr[k_0, \ldots, k_t \leftarrow \{0,1\}^n; A^{\tilde{O}(N,t)} = 1] \leq \frac{4.3q^3t}{N^2}$$

for every adversary A making at most q queries.

Proposition 3. *Let $q = N^{\frac{t}{t+1}}/Z$ for some $Z \geq 1$ be such that $q < N/3$. With $\tilde{O}(N,t)$ defined as above,*

$$\Pr[k_0, \ldots, k_t \leftarrow \{0,1\}^n; A^{\tilde{O}(N,t)} = 1] - \Pr[A^{Q_0,\ldots,Q_t} = 1] \leq \frac{t+1}{Z^{t+1}}.$$

for every adversary A making at most q queries, where $Q_0, \ldots, Q_t$ are independent random permutations.

Proposition 2 is the main technical hurdle in our proof. Its proof, however, is entirely combinatorial, given that we actually show this bound holds even when A sees the keys $k_0, \ldots, k_t$. The presence of keys is therefore actually irrelevant for this proposition[1]. We refer to Appendix A for more details.

The proof of Proposition 3, on the other hand, is fairly accessible, and also contains those ingredients that have the most "cryptographic interest".

Proof (of Proposition 3.). We make the standard assumption that the adversary never makes a redundant query (querying $P_i^{\pm 1}(x)$ twice or querying, e.g., $P_i(x)$ after obtaining x as an answer to a query $P_i^{-1}(y)$).

We modify $\tilde{O}(N,t)$ to use a slightly different lazy sampling method, equivalent to $\tilde{O}(N,t)$'s original sampling method. In this new method, we also maintain a flag bad which is originally set to false.

$\tilde{O}(N,t)$'s new sampling method is as follows: when faced with a query $P_i(x)$, $\tilde{O}(N,t)$ samples a value y uniformly at random from the remaining range of $P_i(x)$, that is, uniformly at random from

$$\{0,1\}^n \backslash \{P_i(x') : x' \in \{0,1\}^n, P_i(x') \text{ is defined}\}.$$

$\tilde{O}(N,t)$ then checks if setting $P_i(x) = y$ would make $P_0, \ldots, P_t$ inconsistent; if so, it sets $\mathsf{bad} = \mathsf{true}$, and resumes its original sampling method for the rest of the game (including to answer the last query); otherwise, it sets $P_i(x) = y$, and returns y. Inverse queries are treated the same.

We can also define a value for the bad flag when the adversary has oracle access to the random permutations $(Q_0, Q_1, \ldots, Q_t)$. Originally, set $\mathsf{bad} = \mathsf{false}$ and select random values $k_0, \ldots, k_t$. Set $Q_0, \ldots, Q_t$ to be undefined at all points, and use lazy sampling to define them by simulating the lazy sampling process for $P_0, \ldots, P_t$ up until $\mathsf{bad} = \mathsf{true}$; after $\mathsf{bad} = \mathsf{true}$, simply keep lazy sampling each permutation Q_i while ignoring bad as well as $k_0, \ldots, k_t$.

Obviously, the probability bad is set to true is equal in both worlds, and the two worlds behave identically up until $\mathsf{bad} = \mathsf{true}$. Thus (a standard argument shows that) the adversary's advantage is upper bounded by the probability that bad is set to true.

For simplicity, we upper bound the probability that bad becomes true when the adversary has oracle access to $Q_0, \ldots, Q_t$. In this case, note that it is equivalent to set the bad flag by sampling the values $k_0, \ldots, k_t$ randomly at the end

[1] We note that the bound of Proposition 2 is the bottleneck of Theorem 1. A potential improvement of Proposition 2 might exploit the fact that $k_0, \ldots, k_t$ aren't known to the adversary.

of the game, and then checking whether these values are inconsistent with the partially defined permutations $Q_0, \ldots, Q_t$. (To recall, $k_0, \ldots, k_t$ are inconsistent with $Q_0, \ldots, Q_t$ if there exist no permutations $\overline{Q}_0, \ldots, \overline{Q}_t$ such that

$$\overline{Q}_t(\cdots \overline{Q}_2(\overline{Q}_1(\overline{Q}_0(\cdot) \oplus k_0) \oplus k_1) \cdots \oplus k_{t-1}) \oplus k_t = id.)$$

Given the partially defined permutations $Q_0, \ldots, Q_t$ and values $k_0, \ldots, k_t$ a *contradictory path* is a sequence of values $(x_0, y_0), \ldots, (x_t, y_t)$ such that (i) $Q_i(x_i) = y_i$ for all i and (ii) $|\{i : y_i \oplus x_{i+1} = k_i, 0 \leq i \leq t\}| = t$, where we put $x_{t+1} = x_0$. Because $q < N/3$, Lemma 3 of Section A implies[2] that $Q_0, \ldots, Q_t$ is consistent with $k_0, \ldots, k_t$ if and only if there exists no contradictory path. Since each Q_i contains at most q defined input-output pairs (x_i, y_i) at the end of the game, there are at most q^{t+1} possible different sequences $((x_0, y_0), \ldots, (x_t, y_t))$ such that $Q(x_i) = y_i$ for $0 \leq i \leq t$. For each of these sequences, the probability that the random selection of $k_0, \ldots, k_t$ creates a contradictory path is upper bounded by $(t+1)N^{-t}$, since the condition $k_i = y_i \oplus x_{i+1}$ must be satisfied for all but one value of i, $0 \leq i \leq t$, and we can union bound over this value of i. Hence, by a union bound over the (at most) q^{t+1} possible different sequences, the probability that bad is set to true is at most $\frac{(t+1)q^{t+1}}{N^t} = \frac{t+1}{Z^t}$ as desired.

3.1 An upper bound

For any number of rounds t, there is an (non-adaptive) attack with a query complexity of roughly $t2^{\frac{t}{t+1}n}$, thus meeting the bound on the query complexity for $t = 2$. Note that this is not an attack in the practical sense, as the computational cost is higher than brute force. The idea of this attack is to actually construct (with high probability) a contradictory path for each possible key.

1. Make $2^{\frac{t}{t+1}n}$ queries to E and each of the oracles P_1 to P_t. Denote the set of queries to P_i by $\mathcal{P}_i$ and queries to E_k by $\mathcal{M}$.
2. For each key candidate $(k_0, k_1, \ldots, k_t)$ do:
   (a) Find all sequences of values $(x_1, \ldots, x_{t-1})$ such that $x_1 \in \mathcal{M}$ and $x_i \oplus k_{i-1} \in \mathcal{P}_i$, $\forall 1 \leq i \leq t$ and $P_i(x_i \oplus k_{i-1}) = x_{i+1}$, $\forall 1 \leq i \leq t-1$.
   (b) Check if $P_t(x_t \oplus k_{t-1}) \oplus k_t = E(x_1)$ for all these sequences.
   (c) If so, assume $(k_0, k_1, \ldots, k_t)$ is the correct value of the key;
   (d) otherwise, it is certainly the wrong value of the key.

To get a better reduction on key-candidates, a bit more than $t2^{\frac{t}{t+1}n}$ queries are sufficient.

[2] More precisely, Lemma 3 is applied by setting the edges of the matching M_i to be all pairs $(x_i, y_i \oplus k_i)$ such that $Q_i(x_i)$ is defined; that is M_i encodes the permutation $Q_i(\cdot) \oplus k_i$.

4 Attacks

The bounds proved earlier are information-theoretic bounds which take into account only the number of queries of the random permutations made by an adversary. Of equal interest are attacks which take the computational complexity into account. In this section we consider only attacks in the single key-model. Note that, in the case where all round-keys are independent, related-key attacks exist trivially. However, the situation might be very different in the case where all round-keys are identical, see Section 7 for further discussion on this point.

4.1 Daemen's attack for $t = 1$

For the original Even-Mansour construction (in our setting, this corresponds to $t = 1$), a differential attack has been published by Daemen [10] meeting the lower bound of $2^{n/2}$ evaluations of P proven by Even and Mansour. It can be described as follows:

1. Choose s plaintext pairs (m_i, m_i^*), $1 \leq i \leq s$, with $m_i \oplus m_i^* = \Delta$ for any nonzero constant Δ.
2. Get the encryptions (c_i, c_i^*) of the s pairs.
3. For $2^n/s$ values v:
   (a) Compute $w' := P(v) \oplus P(v \oplus \Delta)$.
   (b) If $w' = c_i \oplus c_i^*$ for some i: Output $k_0 := v \oplus m_1$ and $k_1 := c_1 \oplus P(m_1 \oplus k_0)$ and stop.

For a random permutation P, only very few values of v are expected to satisfy $P(v) + P(v + \Delta) = c_i \oplus c_i^*$. The wrong candidates can be easily filtered in step (3b) by testing them on a few additional encryptions. After encrypting s plaintext pairs, one has to perform about $2 \cdot 2^n/s$ evaluations of P. The expression $2(s + 2^n/s)$ is minimal for $s = 2^{n/2}$. In this case, the time complexity is $2^{n/2}$ with a storage requirement of $2^{n/2}$ plaintext pairs.

4.2 A meet in the middle attack

There is a meet in the middle attack on the t-permutation construction which finds the keys in time and space $2^{tn/2}$ for $t > 1$. This is a straight-forward attack given here for the case $t = 2$:

1. From a pair of messages (m_1, m_2), compute and save in a sorted table, T, the values $P(m_1 \oplus k) \oplus P(m_2 \oplus k)$ for all possible 2^n values of k.
2. Get the encryptions c_1 and c_2 of m_1 respectively m_2.
3. For all 2^n possible values of k' compute $Q^{-1}(c_1 \oplus k') \oplus Q^{-1}(c_2 \oplus k')$ and look for a match in T.
4. Each match gives candidate values for the three keys, which are tested against additional encryptions.

5 Statistical Properties

A fundamental cryptographic property of a block cipher is its Fourier spectrum that completely defines the cipher via the Fourier transform and whose distribution is closely related to the resistance against linear cryptanalysis [9].

To support security claims, block cipher designs usually come with arguments why these Fourier coefficients cannot take values exploitable by an attacker. In most cases, however, formal proofs of these properties appear technically infeasible and designers limit themselves to demonstrating upper bounds on trail probabilities, that can be seen as summands to obtain the actual Fourier coefficients. This solution is usually denoted as the practical security approach for statistical cryptanalysis. Such an approach does not allow an accurate estimation of the data complexity of statistical attacks, that typically depends on numerous trails [24, 28].

As opposed to that, we analyze the construction of key alternating cipher following a provable security approach, by directly investigating its Fourier coefficients. In addition, we provide a more informative analysis than for standard block ciphers, as we study the distribution of the Fourier coefficients for the cipher over all keys, rather than bounding the mean value of this distribution. This is made possible by the use of fixed public permutations in our construction. More precisely, in a key-alternating cipher using $t \geq 2$ fixed public permutations, we study the distribution of the Fourier coefficients over all cipher keys. If these permutations are close to the average over all permutations, we show that this distribution turns out to be very close to that over all permutations, suggesting that the t-round key-alternating construction is theoretically sound from this perspective. This implies that it behaves well with respect to linear cryptanalysis.

On the contrary, the distribution of Fourier coefficients for a fixed point in the Fourier spectrum is nearly degenerated for the key-alternating cipher with $t = 1$ (the Even-Mansour cipher). This emphasizes the constructive effect of having 2 and more rounds in the key-alternating cipher.

5.1 Fourier coefficients over all permutations

Here we recall the definitions of Fourier coefficients and Fourier spectrum as well as the distribution of Fourier coefficients over all permutations. We also introduce some notations we will be using throughout the section.

Notations. The canonical scalar product of two vectors $a, b \in \{0,1\}^n$ is denoted by $a^T b$. We denote the normal distribution with mean μ and variance σ^2 as $\mathcal{N}(\mu, \sigma^2)$. By $X \sim_v \mathcal{D}$, we denote a random variable X following a distribution $\mathcal{D}$ taken over all values of v. The expectation of X with respect to v is denoted by $\mathbf{E}_v[X]$, its variance (with respect to v) by $\mathbf{Var}_v[X]$.

Fourier coefficients and Fourier spectrum. For a permutation $P : \{0,1\}^n \to \{0,1\}^n$, its *Fourier coefficient* at point (α, β) is defined as

$$W^P_{\alpha,\beta} \stackrel{\text{def}}{=} \sum_{x \in \{0,1\}^n} (-1)^{\alpha^T x + \beta^T P(x)}.$$

The collection of Fourier coefficients at all points $(\alpha, \beta) \in \{0,1\}^n \times \{0,1\}^n$ is called the *Fourier spectrum* of P. For a block cipher F, we denote the Fourier coefficient at point (α, β) as $W^F_{\alpha,\beta}[K]$ to emphasize its dependency on key K. If F is the t-round key-alternating cipher, this is denoted by $W^{P_1,\ldots,P_t}_{\alpha,\beta}[K]$.

The following characterisation for the distribution of Fourier coefficients in a Boolean permutation has been proven.

Fact 1 ([12, Corollary 4.3, Lemma 4.6]). *When $n \geq 5$, the distribution of the Fourier coefficient $W^P_{\alpha_0,\beta_0}$ with $\alpha_0, \beta_0 \neq 0$ over all n-bit permutations can be approximated by the following distribution up to continuity correction:*

$$W^P_{\alpha_0,\beta_0} \sim_P \mathcal{N}(0, 2^n). \tag{4}$$

The distribution of Fact 1 is the reference point throughout the section: A block cipher cannot have a better distribution of Fourier coefficients than that close to Fact 1.

5.2 Fourier coefficients in the single-round Even-Mansour cipher

Let F be the basic single-round Even-Mansour cipher, that is, a fixed public permutation P surrounded by two additions with keys k_0 and k_1, respectively (see Figure 1). If $W^P_{\beta_0,\beta_1}$ is the Fourier coefficient for the underlying permutation P at point (β_0, β_1), then the Fourier coefficient for the cipher at this point is

$$W^F_{\beta_0,\beta_1} = (-1)^{\beta_0^T k_0 \oplus \beta_1^T k_1} W^P_{\beta_0,\beta_1}.$$

Now consider the distribution of $W^F_{\beta_0,\beta_1}$ with $\beta_0 \neq 0$, $\beta_1 \neq 0$ taken over all keys (k_0, k_1). Its support contains exactly two points: $W^P_{\beta_0,\beta_1}$ and $-W^P_{\beta_0,\beta_1}$. Thus, the value of $W^F_{\beta_0,\beta_1}$ almost does not vary from key to key. This is crucially different from the reference point – the distribution over all permutations of Fact 1.

5.3 Fourier coefficients in the t-round key-alternating cipher

Now we state the main result of this section. The proof is given in Appendix B.

Theorem 2. *Fix a point (β_0, β_t) with $\beta_0, \beta_t \neq 0$ in the Fourier spectrum of the t-round key-alternating n-bit block cipher with round permutations $P_1, \ldots, P_t$ for $t \geq 2$ and sufficiently high n. Then the distribution of the Fourier coefficient $W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}$ at this point over all keys K is approximated by:*

$$W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}[K] \sim_K \mathcal{N}\left(0, (1+\varepsilon)\left(\frac{2^n-1}{2^n}\right)^{t-1} 2^n\right), \tag{5}$$

assuming that the distributions over points of the Fourier spectra of the permutations P_i, $1 \leq i \leq t$, have variances satisfying

$$\underset{(\beta_{i-1},\beta_i)}{\mathbf{Var}}\left[W^{P_i}_{\beta_{i-1},\beta_i}\right] \geq 2^{n/2}, \tag{6}$$

and that for any given key K, the signs of the Fourier coefficients behave independently for different points. The deviation of the permutations P_i from the mean over all permutations Q_i is quantified by factor $(1+\varepsilon)$:

$$\begin{aligned} &\textstyle\sum_{(\beta_1,\ldots,\beta_{t-1})}\left(W^{P_1}_{\beta_0,\beta_1}\cdots W^{P_t}_{\beta_{t-1},\beta_t}\right)^2 \\ &= (1+\varepsilon)\cdot \mathbf{E}_{Q_1,\ldots,Q_t}\left[\textstyle\sum_{(\beta_1,\ldots,\beta_{t-1})}\left(W^{Q_1}_{\beta_0,\beta_1}\cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right)^2\right]. \end{aligned} \tag{7}$$

Interestingly, the latter deviation ε from the mean in (7) is small for most choices of the P_i. For instance, in case $t = 2$, it can be shown that over all permutations, mean and variance of each summand in (7) are 2^{2n} and 2^{4n+2}, respectively. The whole sum then approximately follows a normal distribution $\mathcal{N}(2^{3n} - 2^{2n}, 2^{5n+2} - 2^{4n+2})$. This means that for *randomly drawn permutations* P_1, P_2, the sum $\sum_{\beta_1}\left(W^{P_1}_{\beta_0,\beta_1}W^{P_2}_{\beta_1,\beta_2}\right)^2$ will be within d standard deviations from its mean with probability $\operatorname{erf}\left(d/\sqrt{2}\right)$. Notably, this implies $\Pr(|\varepsilon| \leq 2^{-n/2+3}) \approx 0.9999$, i.e. $|\varepsilon|$ only very rarely exceeds $2^{-n/2+3}$.

Theorem 2 gives the distribution over all keys of the Fourier coefficient $W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}$ individually for each nontrivial point (β_0, β_t). Appropriate choices for the P_i should have distributions close to $\mathcal{N}(0, 2^n)$ for each nontrivial point, not only for some of them. Conversely, the distribution of the Fourier coefficient at the (trivial) point $(\beta_0, 0)$ differs from (5) for any choice of the P_i, since it is constant over the keys.

Note also that the result of Theorem 2 does not require the underlying permutations to be different. Moreover, it does not require the permutations P_i to be randomly drawn from the set of all permutations, but holds for any fixed choice of permutations satisfying (6). To obtain a distribution close to ideal, however, the set of underlying permutations has to ensure a small deviation ε in (7). As argued above, drawing the underlying permutations at random from the set of all permutations is highly likely to result in a very small deviation ε from the average.

Summarising, the results of Theorem 2 suggest that once the small number of $t \geq 2$ underlying permutations are carefully chosen and fixed, the t-round key-alternating cipher for each secret key is likely to be statistically sound which rules out some crucial cryptanalytic distinguishers. More precisely, the distributions of the Fourier coefficients for the t-round key-alternating cipher over all keys become close to those over all permutations.

Note that, in contrast to the reference point, it is possible to identify large but efficiently representable subsets of keys where the distribution is again degenerated, as in the case for $t = 1$. Examples of such sets are sets of keys where

one fixes all keys k_1 up to k_{t-1}. For any point (β_0, β_1) the value of $W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}$ takes on only two possible values - over all possible sub-keys k_0, k_t. However, it seems unlikely that this can be used in an attack.

6 Practical constructions

In this section, we discuss possible practical realisations of the t-round key-alternating cipher.

A natural approach to building a practical cipher following the t-permutation construction is to base the t fixed permutations on a block cipher by fixing some keys. With $t = 1$, this corresponds to the original Even-Mansour construction, so the security level is limited to $2^{n/2}$ operations with n denoting the cipher's block length. With a 128-bit block cipher such as the AES, we therefore only obtain a security level of 2^{64} in terms of computational complexity, so it is advisable to choose $t > 1$.

In the following we describe a sample construction with $t = 2$, that is, we consider the 2-round key alternating construction with permutations P_1 and P_2 and the keys k_0, k_1, k_2.

6.1 AES^2: a block cipher proposal based on AES

The construction is defined by fixing two randomly chosen 128-bit AES-128 keys, which specifies the permutations P_1 and P_2. The key is comprised by three independently chosen 128-bit secret keys k_0, k_1, k_2.

Let AES$[k]$ denote the (10-round) AES-128 algorithm with the 128-bit key k and the 128-bit quantities π_1, π_2 be defined based on the first 256 bits of the binary digit expansion of $\pi = 3.1415\ldots$:

$$\pi_1 := \texttt{0x243f6a8885a308d313198a2e03707344} \quad \text{and}$$
$$\pi_2 := \texttt{0xa4093822299f31d0082efa98ec4e6c89}.$$

Then we denote the resulting 2-permutation construction by AES$^2[k_0, k_1, k_2]$. Its action on the 128-bit plaintext m is defined as:

$$\text{AES}^2[k_0, k_1, k_2](m) := \text{AES}[\pi_2](\text{AES}[\pi_1](m \oplus k_0) \oplus k_1) \oplus k_2. \tag{8}$$

Security. Any attack on AES^2 in the single secret-key model with complexity below 2^{85} will have to make use of AES with a fixed known key in a non-black box manner. On the other hand, we are aware of no attack with a computational complexity of less than 2^{128}. Moreover, if the distribution of Fourier coefficients for AES$[\pi_1]$ and AES$[\pi_2]$ meets the assumption of average behaviour, Theorem 2 suggests that the Fourier coefficients for AES^2 are distributed close to ideal which implies resistance against basic linear cryptanalysis and some of its variants. Intuitively, this construction can be seen to arguably transfer the security properties for AES with a single randomly fixed key to the entire cipher as a set of permutations. For AES^2, we explicitly do not claim any related-, known- or chosen-key security.

Performance. AES^2 can be implemented very efficiently in software on general-purpose processors. The two AES keys π_1 and π_2 are fixed and, therefore, the round keys for the two AES transformations can be precomputed, so there is no need to implement the key scheduling algorithm of AES. This ensures high key agility of AES^2.

On the Westmere architecture generation of Intel general-purpose processors, AES^2 can be implemented using the AES-NI instruction set [18]. As the AES round instructions are pipelined, we fully utilise the pipeline by processing four independent plaintext blocks in parallel implementing the basic electronic codebook mode (ECB) and counter mode (CTR). The performance of these implementations on recent processors is demonstrated and compared to two conventional implementations of AES-128 (i.e. without AES-NI instructions) – the bitsliced implementation of [20] and the OpenSSL 1.0.0e implementation based on lookup tables. All numbers are given in cycles per byte (cpb).

	Intel Xeon X5670 2.93 GHz, 12 MB L3 cache	Intel Core i7 640M 2.8 GHz, 4 MB L3 cache
AES^2, AES-NI, ECB	2.54 cpb	2.69 cpb
AES^2, AES-NI, CTR	2.65 cpb	2.76 cpb
AES-128, AES-NI, ECB	1.18 cpb	1.25 cpb
AES-128, AES-NI, CTR	1.32 cpb	1.36 cpb
AES-128, bitsliced, CTR	7.08 cpb	7.84 cpb
AES-128, OpenSSL, CTR	15.73 cpb	16.76 cpb

It turns out that on both platforms, the performance of AES^2 is almost equal to half that of AES, indicating that the overhead is very low. Compared to the best implementations of the AES which are in widespread use now on standard platforms, AES^2 provides a performance improvement of almost factor three and higher with the AES-NI instruction set.

7 Conclusion, Open Problems and Future Work

In this paper we gave the first formal treatment of the key-alternating cipher in a provable setting. For two or more rounds an attacker needs to query the oracles at least $2^{2n/3}$ times for having a reasonable success probability. Furthermore, we studied the security of the construction with respect to statistical attacks, arguing that even for $t = 2$ linear attacks do not seem to be applicable. Finally we gave a concrete proposal mimicking the construction for $t = 2$. There are several lines of future work and open problems we like to mention.

On the theoretical side, it seems unlikely that the bounds given here are tight. Thus, improving them is an important open problem. We actually conjecture that the correct bound on the query complexity is roughly $2^{t/(t+1)n}$. As a first step, deriving bounds that increase with the number of rounds is a goal worth aiming for. Secondly, for now, we have to assume that all round keys are

independent. For aesthetical reasons, but also from a practical point of view (see below) it would be nice to prove bounds for the case that all round keys are identical.

On the practical side, mainly for efficiency reasons but also due to resistance against related-key attacks, several variants for $t = 2$ are worth studying. First of all, since the security level is at most 2^n, due to the meet in the middle attack, one could be tempted to derive three n-bit keys k_0, k_1, and k_2 from one n-bit word. The simplest case here is to have all three keys identical. Taking P and Q different, we are not aware of any attack with computational complexity below 2^n. Furthermore, it seems reasonable to assume that such a construction provides some security against certain types of related-key attacks as well. The best attacks we are aware of in such a setting has birthday complexity $2^{n/2}$. See Appendix C for the details.

Eventually, it is an interesting open problem to determine whether the results in this work can be used as directions for alternative block cipher designs, e.g. with minimum key scheduling algorithms. As a typical example, one could consider the possibility to generate public permutations from a variant of the AES, where the round keys would be replaced with simple constants. In general, such an approach could lead to efficient lightweight designs. Interestingly, it is also the direction taken, to a certain extent, by the recently proposed block cipher LED [19]. In its 64-bit version, this cipher just iterates blocks made of 4 rounds and the addition of the master key.

Another tempting way, in order to increase efficiency, is to choose $Q = P$. Similarly, it may be advantageous to have $Q = P^{-1}$, which has the further advantage that the decryption and encryption operations are similar, except for using the keys in reverse order. However, with $Q = P^{-1}$ there is an attack which finds the value of $k_0 \oplus k_2$ using $2^{n/2}$ queries and similar time. After $k_0 \oplus k_2$ is known the cipher is easily distinguishable from a random permutation. Also, with $Q = P$ but now assuming that $k_0 \oplus k_2$ is known, one finds the secret keys using $2^{n/2}$ queries and similar time.

Acknowledgements. Andrey Bogdanov is a postdoctoral fellow of the Fund for Scientific Research - Flanders (FWO). Francois-Xavier Standaert is associate researcher of the Belgian fund for scientific research (FNRS-F.R.S.). This work has been funded in parts by the ERC project 280141 (acronym CRASH). John Steinberger is supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61061130540, 61073174, and by NSF grant 0994380.Elmar Tischhauser is a doctoral fellow of the Fund for Scientific Research - Flanders (FWO). This work is supported in part by the IAP Programme P6/26 BCRYPT of the Belgian State, by the European Commission under contract number ICT-2007-216676 ECRYPT NoE phase II, by KU Leuven-BOF (OT/08/027), and by the Research Council KU Leuven (GOA TENSE).

References

1. Thomas Baignères and Matthieu Finiasz. Dial C for Cipher. Selected Areas in Cryptography, LNCS 4356, pp. 76–95, Springer-Verlag, 2006.
2. Thomas Baignères and Matthieu Finiasz. KFC - The Krazy Feistel Cipher. ASIACRYPT 2006, LNCS 4284, pp. 380–395, Springer-Verlag, 2006.
3. Paulo S.L.M. Barreto and Vincent Rijmen. The KHAZAD Legacy-Level Block Cipher. First open NESSIE Workshop, 15 pages, Leuven, Belgium, November 2000.
4. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. Keccak sponge function family main document. Submission to NIST (Round 2), 2009.
5. Alex Biryukov, Orr Dunkelman, Nathan Keller, Dmitry Khovratovich and Adi Shamir. Key Recovery Attacks of Practical Complexity on AES-256 Variants with up to 10 Rounds. EUROCRYPT 2010, LNCS 6110, pp. 299–319, Springer-Verlag, 2010.
6. Alex Biryukov and Dmitry Khovratovich. Related-Key Cryptanalysis of the Full AES-192 and AES-256. ASIACRYPT 2009, LNCS 5912, pp. 1–18, Springer-Verlag, 2009.
7. Andrey Bogdanov, Dmitry Khovratovich and Christian Rechberger. Biclique Cryptanalysis of the Full AES. ASIACRYPT 2011, LNCS 7073, pp. 344–371, Springer-Verlag, 2011.
8. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin and C. Vikkelsoe: PRESENT: An Ultra-Lightweight Block Cipher. CHES 2007, LNCS 4727, pp. 450–466, Springer-Verlag, 2007.
9. Florent Chabaud and Serge Vaudenay. Links between differential and linear cryptanalysis. EUROCRYPT 94, LNCS 950, pp. 356–365, Springer-Verlag, 1995.
10. Joan Daemen. Limitations of the Even-Mansour Construction. ASIACRYPT 1991, LNCS 739, pp. 495–498, Springer-Verlag, 1991.
11. Joan Daemen, Rene Govaerts and Joos Vandewalle. Correlation matrices. FSE 1994, LNCS 1008, pp. 275–285, Springer-Verlag, 1995.
12. Joan Daemen and Vincent Rijmen. Probability distributions of correlations and differentials in block ciphers. Journal on Mathematical Cryptology 1(3), pp. 221–242, 2007.
13. Joan Daemen and Vincent Rijmen. The Design of Rijndael. Springer-Verlag, 2002.
14. Joan Daemen and Vincent Rijmen. The Wide Trail Design Strategy. IMA Int. Conf., LNCS 2260, pp. 222–238, Springer-Verlag, 2001.
15. Shimon Even and Yishay Mansour. A Construction of a Cipher from a Single Pseudorandom Permutation. J. Cryptology, vol. 10, num. 3, pp. 151–162, 1997.
16. Shimon Even and Yishay Mansour. A Construction of a Cipher From a Single Pseudorandom Permutation. ASIACRYPT 1991, LNCS 739, pp. 210–224, Springer-Verlag, 1993.
17. FIPS PUB 46-3: DATA ENCRYPTION STANDARD (DES). 1999.
18. Shay Gueron. Intel Mobility Group, Israel Development Center, Israel: Intel Advanced Encryption Standard (AES) Instructions Set, 2010. Available at http://software.intel.com/file/24917.
19. Jian Guo, Thomas Peyrin, Axel Poschmann and Matt Robshaw: The LED Block Cipher. CHES 2011, LNCS 6917, pp. 326–341, Spinger-Verlag, 2011.
20. Emilia Käsper and Peter Schwabe. Faster and Timing-Attack Resistant AES-GCM. CHES 2009, LNCS 5747, pp. 1–17, Springer-Verlag, 2009.

21. Liam Keliher, Henk Meijer and Stafford E. Tavares. Improving the Upper Bound on the Maximum Average Linear Hull Probability for Rijndael. Selected Areas in Cryptography, LNCS 2259, pp. 112–128, Springer-Verlag, 2001.
22. Lars R. Knudsen. Practically Secure Feistel Ciphers. FSE 1993, LNCS 809, pp. 211–221, Springer-Verlag, 1991.
23. Xuejia Lai and James L. Massey. A Proposal for a New Block Encryption Standard. EUROCRYPT 1990, LNCS 473, pp. 389–404, Springer-Verlag, 1990.
24. Xuejia Lai and James L. Massey. Markov Ciphers and Differentail Cryptanalysis. EUROCRYPT 1991, LNCS 547, pp. 17–38, Springer-Verlage, 1991.
25. Michael Luby and Charles Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. SIAM J. Comput., vol. 17, num. 2, pp. 373–386, 1988.
26. Mitsuru Matsui. New Block Encryption Algorithm MISTY. FSE 1997, LNCS 1267, pp. 54–68, Springer-Verlag, 1997.
27. Mitsuru Matsui. New Structure of Block Ciphers with Provable Security against Differential and Linear Cryptanalysis. FSE 1996, LNCS 1039, pp. 205–218, Springer-Verlag, 1996.
28. Kaisa Nyberg. Linear Approximation of Block Ciphers. EUROCRYPT 1994, LNCS 950, pp. 439–444, Springer-Verlag, 1994.
29. Luke O'Connor. Properties of Linear Approximation Tables. FSE 1994, LNCS 1008, pp. 131–136, Springer-Verlag, 1995.
30. Vincent Rijmen, Joan Daemen, Bart Preneel, Antoon Bosselaers and Erik De Win. The Cipher SHARK. FSE 1996, LNCS 1039, pp. 99–111, Springer-Verlag, 1996.
31. Aris Spanos. Probability Theory and Statistical Inference: Econometric Modeling with Observational Data. Cambridge University Press, 1999.
32. Serge Vaudenay. Decorrelation: A Theory for Block Cipher Security. J. Cryptology, vol. 16, num. 14, pp. 249–286, 2003.
33. Serge Vaudenay. On the Lai-Massey Scheme. ASIACRYPT 1999, LNCS 1716, pp. 8–19, Springer-Verlag, 1999.

A Proof of Proposition 2

In this section we provide a proof of Proposition 2, which constitutes the most technical part of our paper. The argument is structured as follows: Firstly, we allow the adversary A to see the values $k_0, \ldots, k_t$ (in fact, we even allow A to choose these values); obviously, such an adversary can only perform better than an adversary without knowledge of $k_0, \ldots, k_t$. Secondly, we argue that the k_i's can be set to 0^n without any loss in advantage. The problem then reduces to upper bounding the adversary's advantage at distinguishing two different methods of lazy sampling permutations $P_0, \ldots, P_t$ such that

$$P_t(\cdots P_2(P_1(\cdot))\cdots) = id.$$

We then directly argue (by "low-level" combinatorics) that, for any given query, the statistical distance between the two types of sampling is small. In order to facilitate the application of this statistical distance bound on a single query to the general q-query setting, we introduce another abstraction (of potential interest on its own) that we call *sample distinguishability*. As it applies to our

setting, the sample distinguishability game lets the adversary "set up" each of its queries as it wants; namely, it can partially define the permutations $P_0, \ldots, P_t$ on at most q points each (subject to the consistency constraint), and then request its oracle—which is either $O(N,t)$ or $\tilde{O}(N,t)$—to answer a query of its choice, for that choice of $P_0, \ldots, P_t$; at its next query, the adversay can set up $P_0, \ldots, P_t$ again from scratch, and so on. Clearly such an adversary with "set-up power" has advantage at least that of a standard adversary. We then combine the statistical distance bound for a single query (Lemma 2 below) with a simple lemma relating sampling distinguishability to single-sample statistical distance (Lemma 1 below) to obtain the final result (restated as Lemma 5 below, that is equivalent to Proposition 2).

Define

$$\mathbf{Adv}^{O\tilde{O}}_{N,t}(A) = \Pr[A \to k_0, \ldots, k_t; A^{O(N,t)} = 1] - \Pr[A \to k_0, \ldots, k_t; A^{\tilde{O}(N,t)} = 1]$$

and

$$\mathbf{Adv}^{O\tilde{O};0^n}_{N,t}(A) = \Pr[k_0 = \cdots = k_t = 0^n; A^{O(N,t)} = 1] - \Pr[k_0 = \cdots = k_t = 0^n; A^{\tilde{O}(N,t)} = 1].$$

Obviously, it suffices to show that $\mathbf{Adv}^{O\tilde{O}}_{N,t}(A) \leq 4.3q^3t/N^2$ for all A making at most $q < N/100$ queries in order to prove Proposition 2. (Indeed A is free to choose $k_0, \ldots, k_t$ randomly and then forget about these values.) Our first proposition shows that, in fact, it is sufficient to upper bound $\mathbf{Adv}^{O\tilde{O};0^n}_{N,t}(\cdot)$.

Proposition 4. *For every q-query adversary A there exists a q-query adversary A' such that*

$$\boldsymbol{Adv}^{O\tilde{O};0^n}_{N,t}(A') = \boldsymbol{Adv}^{O\tilde{O}}_{N,t}(A).$$

Proof. A' simulates A; let $k_0, \ldots, k_t$ be the keys chosen by A. When A makes a query $P_i(x)$, A' queries $P_i(x)$ and returns $P_i(x) \oplus k_i$ to A; when A queries $P_i^{-1}(x)$, A' queries $P_i^{-1}(x \oplus k_i)$. It is easy to check that when A''s oracle is $O(N,t)$ (resp. $\tilde{O}(N,t)$), then A' provides A with a perfect simulation of $O(N,t)$ (resp. $\tilde{O}(N,t)$) on keys $k_0, \ldots, k_t$. It follows that A''s advantage is exactly A's.

The rest of our effort focuses on upper bounding $\mathbf{Adv}^{O\tilde{O};0^n}_{N,t}(A)$ for a q-query adversary A; namely, by Proposition 4, it is sufficient to show that $\mathbf{Adv}^{O\tilde{O};0^n}_{N,t}(A) < 4.3q^3t/N^2$ when $q < N/100$. The latter upper bound is finally established in Corollary 2 below.

We now abstract the problem of distinguishing the oracles $O(N,t)$, $\tilde{O}(N,t)$ into a more general type of game (that is also more generous to the adversary). This game is the "sample distinguishability" game referred to above.

Let $(X_\alpha)_{\alpha \in B}$ be a sequence of random variables indexed by some finite set B, where each X_α takes values in some finite set S. We write $A^{(X_\alpha)_{\alpha \in B}}$ to mean an adversary A with oracle access to a sequence of random variables indexed by the elements of B. More precisely, the adversary's query sequence has the form $(\alpha_1, \alpha_2, \ldots, \alpha_q)$, where each $\alpha_i \in B$; such a query sequence is answered by sampling $X_{\alpha_1}, \ldots, X_{\alpha_q}$, and returning these values to the adversary. Every sample is

taken independently of previous samples; in particular, if the adversary queries the same X_α twice, then X_α is sampled twice independently. The adversary is adaptive, and can query its oracles in any order it wishes.

We next define the notion of (adaptive) sample distinguishability.

Definition 1. *Let $(X_\alpha, Y_\alpha)_{\alpha\in B}$ denote a family of pairs of random variables indexed by a finite set B, where each random variable takes values in the same finite range S. We define*

$$\boldsymbol{Adv}^{\mathsf{samp-dist}}_{(X_\alpha,Y_\alpha)_{\alpha\in B}}(A) = \Pr[A^{(X_\alpha)_{\alpha\in B}} = 1] - \Pr[A^{(Y_\alpha)_{\alpha\in B}} = 1]$$

with the probabilities being taken over the randomness of the distributions and over the adversary's coins, if any. We also define

$$\boldsymbol{Adv}^{\mathsf{samp-dist}}_{(X_\alpha,Y_\alpha)_{\alpha\in B}}(q) = \max_A \boldsymbol{Adv}^{\mathsf{samp-dist}}_{(X_\alpha,Y_\alpha)_{\alpha\in B}}(A)$$

where the maximum is taken over all adversaries A making at most q queries.

We note that non-adaptive sample distinguishability—in which case the adversary must announce its sequence of queries $(\alpha_1, \ldots, \alpha_q)$ before receiving any answers—reduces to upper bounding the maximum statistical distance of the form

$$\Delta((X_{\alpha_i})_{i=1}^q, (Y_{\alpha_i})_{i=1}^q)$$

where $(X_{\alpha_i})_{i=1}^q$ is the product distribution[3] $(X_{\alpha_1}, \ldots, X_{\alpha_q})$ and likewise for $(Y_{\alpha_i})_{i=1}^q$, and with this maximum being taken over all possible sequences $(\alpha_1, \ldots, \alpha_q)$. Since all samples are taken independently (in the adaptive as well as in the non-adaptive game), it might intuitively seem that adapativity doesn't help, but, surprisingly, it does. (An example appears at the end of this section.)

Lemma 1. *Let $(X_\alpha, Y_\alpha)_{\alpha\in B}$ be a set of pairs of random variables indexed by the finite set B. Then*

$$\boldsymbol{Adv}^{\mathsf{samp-dist}}_{(X_\alpha,Y_\alpha)_{\alpha\in B}}(q) \le q \cdot \max_{\alpha\in B} \Delta(X_\alpha, Y_\alpha).$$

Proof. We use a coupling argument. For each $\alpha \in B$, let $(\tilde{X}_\alpha, \tilde{Y}_\alpha)$ be a maximal coupling of X_α and Y_α; this means $\tilde{X}_\alpha$ and $\tilde{Y}_\alpha$ are defined on the same probability space, such that

$$\Pr[\tilde{X}_\alpha \ne \tilde{Y}_\alpha] = \Delta(X_\alpha, Y_\alpha)$$

and such that $\tilde{X}_\alpha$ is equidistributed to X_α and $\tilde{Y}_\alpha$ is equidistributed to Y_α. (That such distributions $\tilde{X}_\alpha$, $\tilde{Y}_\alpha$ exist is a standard fact.) When we sample $\tilde{X}_\alpha$ we thus "automatically" sample $\tilde{Y}_\alpha$, and vice-versa. When an adversary A interacts with oracle $(\tilde{X}_\alpha)_{\alpha\in B}$, let bad be the event that, for one of the queries α asked by A, $\tilde{X}_\alpha \ne \tilde{Y}_\alpha$ (note that only $\tilde{X}_\alpha$ is returned to A). We likewise define bad when the

[3] An r.v. X of the form $X = (X_1, \ldots, X_q)$ is a *product distribution* when the X_i's are independent.

adversary interacts with oracle $(\tilde{Y}_\alpha)_{\alpha\in B}$. Note that as long as $\mathsf{bad} = \mathsf{false}$, the two oracles are equidistributed. (Indeed, the entity performing the queries is "doing the same thing" in either world: namely, sampling the required pair $(\tilde{X}_\alpha, \tilde{Y}_\alpha)$ and returning the common value of this pair.) The adversary's advantage is at most the probability of setting $\mathsf{bad} = \mathsf{true}$ which, by a union over the q queries, is at most

$$q \cdot \max_\alpha \Pr[\tilde{X}_\alpha \neq \tilde{Y}_\alpha] = q \cdot \max_\alpha \Delta(X_\alpha, Y_\alpha).$$

Note: We believe Lemma 1 is far from tight. Moreover, it currently constitutes the main "bottleneck" in our security bound. One could improve Lemma 1 by showing, for example, that the advantage of an adaptive sample distinguishability adversary is upper bounded by a constant (e.g., 2) times the advantage of a non-adaptive adversary, but we do not know of such a bound. See also the example relating adaptivity to non-adaptivity at the end of this section.

To state our main technical result we need to define the family of pairs $(X_\alpha, Y_\alpha)_{a\in B}$ that we are interested in applying Lemma 1 to. Parameters for this family will be N, q and t. Recall that upper bounding $\mathbf{Adv}^{O\tilde{O};0^n}_{N,t}(A)$ for a q-query adversary A means upper bounding A's distinguishing advantage between two different lazy sampling methods for a sequence of permutations $P_0, \ldots, P_t$ such that

$$P_t(\cdots P_1(P_0(\cdot))\cdots) = id. \tag{9}$$

In the following, we model the (partially defined) permutations by their associated (partial) matchings. That is, a partially defined permutation on $\{0,1\}^n$ is defined by a matching with left vertex set $\{0,1\}^n$ and right vertex set $\{0,1\}^n$, in the natural way. Composing sevaral permutations corresponds to gluing the associated matchings side by side.

Let $V_0, \ldots, V_t, V_{t+1}$ be sets of vertices with $|V_i| = N$, and where we identify V_{t+1} with V_0 (i.e., V_{t+1} and V_0 are two different names for the same set).

A sequence of matchings $\overline{M} = (\overline{M}_0, \ldots, \overline{M}_{t+1})$ where $\overline{M}_i$ is a perfect matching between V_i and V_{i+1} is called *circular* if every path starting at a vertex $v \in V_0$ following the edges in $\overline{M}_0, \ldots, \overline{M}_t$ ends at the same vertex $v \in V_{t+1} = V_0$. Thus, circularity is the matching equivalent of (9).

Given a sequence $M = (M_0, \ldots, M_t)$ where each M_i is a partial matching between V_i and V_{i+1}, we let

$$\mathcal{M}(M)$$

be the set of all circular sequences $\overline{M}$ extending M, i.e. the set of all sequences $\overline{M} = (\overline{M}_0, \ldots, \overline{M}_t)$ such that $\overline{M}_i$ extends M_i for each i and such that $\overline{M}$ is circular. We say M is *consistent* if $\mathcal{M}(M) \neq \emptyset$. (This fits our previous definition of consistency, restricted to the case $k_0 = \ldots = k_t = 0^n$.)

Let a *q-configuration* be a pair (v_0, M) such that (i) $M = (M_0, \ldots, M_t)$ is a consistent sequence of partial matchings such that $|M_i| \leq q$ for all i, (ii) $v_0 \in V_0$ is nonadjacent to M_0. Our index set B for the family of pairs $(X_\alpha, Y_\alpha)_{\alpha\in B}$ will be exactly the set of all q-configurations. That is,

$$B = \{(v_0, M) : (v_0, M) \text{ is a } q\text{-configuration}\}.$$

We now describe, for a given $\alpha = (v_0, M) \in B$, the distributions X_α and Y_α.

Let $\alpha = (v_0, M)$ be a q-configuration, $M = (M_0, \ldots, M_t)$. For any vertex $u \in V_1$ nonadjacent to M_0, we write $M \cup \{(v_0, u)\}$ for the sequence of partial matchings $(M_0 \cup \{(v_0, u)\}, M_1, \ldots, M_t)$. Let $U \subseteq V_1$ be the set of vertices u such that $M \cup \{(v_0, u)\}$ is consistent. We define

$$\Pr[X_\alpha = u] := \frac{\mathcal{M}(M \cup \{(v_0, u)\})}{\mathcal{M}(M)}.$$

We note that X_α is a probability distribution on U, and that X_α is equidistributed to $O(N, t)$ queried at $P_0(v_0)$ with keys $k_0 = \ldots = k_t = 0^n$ and with $P_0, \ldots, P_t$ defined such that $P_i(x) = y \iff (x, y) \in M_i$. As for Y_α, it is simply the uniform distribution on U. Thus Y_α is equidistributed to $\tilde{O}(N, t)$ under the same correspondence. (Note the restriction to queries of the form $P_0(\cdot)$ is without loss of generality, since the adversary can "set up" the matchings as its wants in the sample distinguishability game.)

The crux of our proof is the following lemma:

Lemma 2. *Let $q < N/100$ and let $t \geq 1$. Then for any q-configuration $\alpha = (v_0, M)$, we have*

$$\Delta(X_\alpha, Y_\alpha) \leq \frac{2q\rho}{N - 2q}.$$

where $\rho = 2.05qt/N$, with X_α and Y_α defined as above.

We need two more small results before giving the proof of Lemma 2. For a sequence of partial matchings $M = (M_0, \ldots, M_t)$, a path of length $t + 1$ using edges from the partial matchings $M_0, \ldots, M_t$ (possibly "wrapping around" through $V_0 = V_{t+1}$) is *contradictory* if it contains $t + 2$ vertices (i.e., if it is not a cycle—this is also a restatement of our previous definition of a contradictory path, restricted to the case where the k_i's are 0^n). Obviously, if a partial matching contains a contradictory path it cannot be consistent. The next lemma gives a partial converse.

Lemma 3. *Let $q \leq N/3$. Then a partial sequence of matchings $M = (M_0, \ldots, M_t)$ where each M_i has at most q edges each is consistent if and only if it contains no contradictory path.*

Proof. We show that M can be extended to a circular sequence of perfect matchings $\overline{M}$. The extension follows three steps: (i) for each edge in M_0, if such an edge is not already in a path of length $t + 1$, then we complete a non-contradictory path of length $t+1$ containing that edge; (ii) we arbitrarily extend the matchings $M_1, \ldots, M_t$ to perfect matchings $\overline{M}_1, \ldots, \overline{M}_t$; (iii) we complete the matching M_0 to a perfect matching $\overline{M}_0$ in the unique way that will make $\overline{M} = (\overline{M}_0, \ldots, \overline{M}_t)$ circular.

Steps (ii) and (iii) can obviously carried out if step (i) succeeds, so it remains to prove that step (i) is possible.

In the process of carrying out step (i), let (v_0, v_1) be an edge in M_0 that is not yet in a cycle, $v_0 \in V_0$, $v_1 \in V_1$. Say that a node is "free" if it is adjacent to

no edges (note that to start with, there are at least $N/3$ free nodes in each V_i). Let

$$(v_\ell, v_{\ell+1}, \ldots, v_t, v_{t+1} = v_0, v_1, \ldots, v_k)$$

be the maximal path containing (v_0, v_1), where $v_i \in V_i$, where $\ell \leq t+1$ and $k \geq 1$. By assumption that there are no contradictory paths, $k < \ell$. If $k = \ell - 1$ then we can simply connect v_k and v_ℓ by an edge. Otherwise, as long as there are free nodes left in each of the layers $V_{k+1}, \ldots, V_{\ell-1}$, we can use these to connect v_k to v_ℓ by a path. However, we start with at least $N/3$ free nodes in each layer, and we have only at most $N/3$ paths to create (one for each of M_0). Hence such free nodes will always exist.

The following is an elementary observation that trusting readers can take for granted.

Lemma 4. *Let a set U be the disjoint union of sets R, T, and let $\rho \in [0, \frac{1}{2}]$. Let Y be the uniform distribution over U and let X be a random variable such that $\Pr[X = u_1] = \Pr[X = u_2]$ for all $u_1, u_2 \in R$ and such that*

$$\Pr[X = u_1] \in [(1-\rho)\Pr[X = u_2], (1-\rho)^{-1}\Pr[X = u_2]] \tag{10}$$

for all $u_1, u_2 \in U$. Then

$$\Delta(X, Y) \leq \frac{2\rho|T|}{|U|}.$$

Proof. We start by noting that since there must exist some $u_1 \in U$ such that $\Pr[X = u_1] \leq 1/|U|$, and also some $u_2 \in U$ such that $\Pr[X = u_2] \geq 1/|U|$, the second condition implies that

$$\Pr[X = s] \in [(1-\rho)/|U|, (1-\rho)^{-1}/|U|]$$

for all $s \in U$. We also note that $\rho \in [0, \frac{1}{2}]$ implies $(1-\rho)^{-1} \leq 1 + 2\rho$. Since $|\Pr[X = s] - \Pr[Y = s]| \leq 2\rho/|U|$ for all s, the lemma obviously holds when $|T| = |U|$. We can therefore assume $R \neq \emptyset$.

Let p be the probability $\Pr[X = u]$ for some $u \in R$ (where by assumption this probability does not depend on the choice of $u \in R$). We consider two cases according to whether $p \geq 1/|U|$ or $p \leq 1/|U|$. Assume first that $p \leq 1/|U|$. Then $\Pr[X = s] \leq \Pr[Y = s]$ for all $s \in R$, so

$$\begin{aligned}
\Delta(X, Y) &= \max_{S \subseteq U} \sum_{s \in S} \Pr[X = s] - \Pr[Y = s] \\
&= \max_{S \subseteq T} \Pr[X = s] - \Pr[Y = s] \\
&\leq \sum_{s \in T} (1-\rho)^{-1}\Pr[Y = s] - \Pr[Y = s] \\
&\leq |T| 2\rho/|U|
\end{aligned}$$

as desired. If $p \geq 1/|U|$ then $\Pr[X = s] \geq \Pr[Y = s]$ for all $s \in R$, so

$$\begin{aligned}\Delta(X, Y) &= \max_{S \subseteq U} \sum_{s \in S} \Pr[Y = s] - \Pr[X = s] \\ &= \max_{S \subseteq T} \Pr[Y = s] - \Pr[X = s] \\ &\leq \sum_{s \in T} \Pr[Y = s] - (1 - \rho) \Pr[Y = s] \\ &\leq |T|\rho/|U|.\end{aligned}$$

Proof (Proof of Lemma 2). Assume first there is a path in $M_1, \ldots, M_t$ ending at $v_0 \in V_{t+1} = V_0$. Then, obviously, $|U| = 1$ and $\Delta(X_\alpha, Y_\alpha) = 0$. Thus, we can assume there is no such path.

In view of applying Lemma 4, let $R \subseteq U$ be the set of free nodes in V_1 (as defined in the proof of Lemma 3), and let $T = V_1 \backslash R \leq q$. Because $q < N/100 < N/3$, Lemma 3 implies that R in fact consists of all free nodes in V_1. Thus $|U| \geq |R| \geq N - 2q$, and

$$\frac{2\rho|T|}{|U|} \leq \frac{2q\rho}{N - 2q}. \tag{11}$$

Put $X = X_\alpha$. It is easy to check that $\Pr[X = u_1] = \Pr[X = u_2]$ for all $u_1, u_2 \in R$. Indeed, an easy path-switching argument shows that when $u_1, u_2 \in R$ there is a bijection between $\mathcal{M}(M \cup \{(v_0, u_1)\})$ and $\mathcal{M}(M \cup \{(v_0, u_2)\})$. In order to apply Lemma 4 and conclude the proof it thus only remains to show that

$$\frac{\Pr[X = u_1]}{\Pr[X = u_2]} \geq 1 - \rho$$

for all $u_1, u_2 \in U$. (Note this indeed implies (10).) By definition of $\Pr[X = u]$, this is equivalent to showing

$$\frac{\mathcal{M}(M \cup \{(v_0, u_1)\})}{\mathcal{M}(M \cup \{(v_0, u_2)\})} \geq 1 - \rho \tag{12}$$

for all $u_1, u_2 \in U$.

For every circular matching sequence $\overline{M} \in \mathcal{M}(M)$, let $C(\overline{M})$ be the consistent sequence of partial matchings obtained by restricting $\overline{M}$ to edges that are either in M or else in a path that contains an edge in M_0. Note that each partial matching in $C(\overline{M})$ has size at most $2q$, and that the matching from V_0 to V_1 in $C(\overline{M})$ coincides with M_0. Moreover, let

$$\mathcal{C}(M) = \{C(\overline{M}) : \overline{M} \in \mathcal{M}(M)\}$$

be the set of all such sequences of partial matchings. We note that every element of $\mathcal{M}(M)$ extends some (in fact, exactly one) element of $\mathcal{C}(M)$. (Though several elements of $\mathcal{M}(M)$ may extend the same element of $\mathcal{C}(M)$.)

Note that

$$|\mathcal{M}(M \cup \{(v_0, u_1)\})| = \sum_{K \in \mathcal{C}(M)} \mathcal{M}(K \cup \{(v_0, u_1)\})$$
$$|\mathcal{M}(M \cup \{(v_0, u_2)\})| = \sum_{K \in \mathcal{C}(M)} \mathcal{M}(K \cup \{(v_0, u_2)\})$$

Also note that neither v_0 nor u_1 nor u_2 are endpoints of an edge in the first matching of any $K \in \mathcal{C}(M)$, since the first matching of K is M_0.

We will show (12) by showing, more strongly, that

$$\frac{|\mathcal{M}(K \cup \{(v_0, u_1)\})|}{|\mathcal{M}(K \cup \{(v_0, u_2)\})|} \geq 1 - \rho \tag{13}$$

for any $K \in \mathcal{C}(M)$. The fact that $K \cup \{(v_0, u_1)\}$ and $K \cup \{(v_0, u_2)\}$ are consistent follows from the fact that $K \cup \{(v_0, u_1)\}$, $K \cup \{(v_0, u_2)\}$ contain no contradictory path (completing cycles cannot add a contradictory path) and that K has at most $2q < N/3$ edges per matching.

Fix therefore $K \in \mathcal{C}(M)$ and let $L_1 := K \cup \{(v_0, u_1)\}$, $L_2 := K \cup \{(v_0, u_2)\}$. Note that

$$L_1 = (M_0 \cup \{(v_0, u_1)\}, K_1, \ldots, K_t)$$
$$L_2 = (M_0 \cup \{(v_0, u_2)\}, K_1, \ldots, K_t)$$

since $K = (K_0 = M_0, K_1, \ldots, K_t)$. Note there is a bijection between elements of $\mathcal{M}(L_j)$ and tuples $(\overline{K}_1, \ldots, \overline{K}_t)$ such that $\overline{K}_i$ is a complete matching extending K_i and such that $u_j \in V_1$ is connected to $v_0 \in V_{t+1}$ by a path of edges from $\overline{K}_1, \ldots, \overline{K}_t$. (This uses the fact that K is picked from $\mathcal{C}(M)$.) Letting $\mathcal{K}_j$ be the set of such sequences $(\overline{K}_1, \ldots, \overline{K}_t)$ for $j = 1, 2$, it therefore suffices to show that

$$|\mathcal{K}_1|/|\mathcal{K}_2| \geq 1 - \rho. \tag{14}$$

Note that any element of $\mathcal{K}_j$ can be "built" the following way: first we extend each K_i, $i \geq 1$, to a partial matching K'_i by adding at most one edge to K_i, such that u_j is connected by a path of edges in $K'_1, \ldots, K'_t$ to $v_0 \in V_{t+1}$, and such that each edge in $K'_i \backslash K_i$ (if any) is an edge on this path; second, we complete each K'_i to a complete matching $\overline{K}_i$, arbitrarily for each i. Furthermore, we can construct the partial matchings $K'_1, \ldots, K'_t$ by the following process. We choose a path from $u_j \in V_1$ to $v_0 \in V_{t+1}$ that is compatible with the matchings $K_1, \ldots, K_t$, and augment these matchings by the edges on that path. More specifically, let $w_1 = u_j$. Let $t' \geq 1$ be the smallest value such that there exists a path from $v_0 \in V_{t+1}$ to a vertex in $V_{t'}$ by edges in $K_t, K_{t-1}, \ldots, K_{t'}$, and let $w_{t'} \in V_{t'}$ be the endpoint of this path. (Possibly, $t' = t + 1$ and $w_{t'} = v_0$.) In fact, $t' \geq 2$, as follows from the fact that M_1 and M_2 are both consistent. For $1 \leq i \leq t' - 1$, we construct $w_{i+1} \in V_{i+1}$ from $w_i \in V_i$ as follows: if w_i is incident to an edge of the matching K_i, let w_{i+1} be the other endpoint of this edge; otherwise, let w_{i+1} be any vertex in V_{i+1} that does not lie on a path of edges in $K_{i+1}, \ldots, K_{t'-1}$

whose endpoint in $V_{t'}$ is not $w_{t'}$ (i.e., either a path of length $t'-i-1$ starting at w_{i+1} does not exist in $K_{i+1},\ldots,K_{t'-1}$, or else the endpoint of this path is $w_{t'}$). It is easy to see that such a w_{i+1} always exists by the consistency of M_j. Furthermore, we note for future use that w_{i+1} can always be chosen to be any free vertex in V_{i+1}, if such a vertex exists, as long as $i+1<t'$. Once $w_1,\ldots,w_{t'}$ are defined, we add to K_i the edge (w_i,w_{i+1}) (if this edge is not already present) for $i=1,\ldots,t'-1$, and we leave $K_{t'},\ldots,K_t$ untouched, resulting in the sequence of partial matchings $(K'_1,\ldots,K'_t)$. There is obviously, by construction, a path from $u_j\in V_1$ to $v_0\in V_{t+1}$ using edges in $K'_1,\ldots,K'_t$, and K'_i differs from K_i only, if at all, by an edge in this path. Furthermore, any sequence of partial matchings $(K'_1,\ldots,K'_t)$ can be obtained by this process.

We have described a two-stage construction of an element of $\mathcal{K}_j$, whereby the matchings $K'_1,\ldots,K'_t$ are first constructed (i.e., a path from u_j to v_0 is first constructed, using the process described above), followed by an arbitrary extension of these matchings to full matchings $\overline{K}_1,\ldots,\overline{K}_t$. We now make a cosmetic change to this process which will help us count the size of $\mathcal{K}_j$. We will first construct $\overline{K}_1$, then $\overline{K}_2$, etc. Let t' and $w_{t'}$ be as above; also let $w_1=u_j$ as above. For $i=1$ to $t'-1$ we do the following: (i) choose w_{i+1} as described above, and add the edge (w_i,w_{i+1}) to K_{i-1} to form K'_i; (ii) extend K'_i arbitrarily to a full matching $\overline{K}_i$. Finally, for $i=t'$ to t, let $K'_i=K_i$ and extend K'_i to an arbitrary full matching $\overline{K}_i$.

The above sequence of choices determining $\overline{K}_1,\ldots,\overline{K}_t$ can be viewed as a tree of depth t, whereby the i-th level of the tree corresponds to the construction of $\overline{K}_t$. The number of leaves in this tree is $|\mathcal{K}_j|$. To upper and lower bound $|\mathcal{K}_j|$ we will upper and lower bound the degree of each non-leaf node.

Let $e_i=|K_i|$ be the number of edges in K_i for $1\le i\le t$. Consider a node r at level i of the tree (where the root has level 1). Say, first, that $i<t'-1$. This node r specifies (among others) a choice of $w_1,\ldots,w_i$, since the first $i-1$ levels of the tree determine $\overline{K}_1,\ldots,\overline{K}_{i-1}$. We distinguish two cases: when w_i is incident to an edge in K_i, and when it is not. If w_i is incident to an edge in K_i then there is a single choice for w_{i+1} and exactly $(N-e_i)!$ ways completing the matching $\overline{K}_i$, since the number of ways to complete the matching K_i is the number of permutations on $N-e_i$ points. In this case, therefore, r has degree $(N-e_i)!$. In the second case, when w_i is not incident to an edge in K_i, then there are at least $N-e_i-e_{i+1}$ choices for w_{i+1}, by the observation made above that w_{i+1} can be any free node in V_{i+1}. Once w_{i+1} is chosen, determining K'_i, there are $(N-e_i-1)!$ ways to extend K'_i to $\overline{K}_i$. Thus altogether, r has degree at least $(N-e_i-e_{i+1})(N-e_i-1)!$ and at most $(N-e_i)!$, in this case. Nextly, when $i=t'-1$, we note that by construction of t' and $w_{t'}$, w_i cannot be adjacent to an edge of K_i; in this case, therefore, r has degree $(N-e_i-1)!$ (since there is a unique choice for $w_{i+1}=w_{t'}$). Finally, when $i\ge t'$, r has degree $(N-e_i)!$ since we just need to extend $K'_i=K_i$ to $\overline{K}_i$.

Altogether, therefore, a lower bound for the number of leaves in the tree (i.e. a lower bound for $|\mathcal{K}_j|$) is

$$(N - e_{t'-1} - 1)! \cdot \prod_{i=1}^{t'-2} (N - e_i - e_{i+1})(N - e_i - 1)! \cdot \prod_{i=t'}^{t} (N - e_i)!$$

and an upper bound for the number of leaves is

$$(N - e_{t'-1} - 1)! \cdot \prod_{i=1}^{t'-2} (N - e_i)! \cdot \prod_{i=t'}^{t} (N - e_i)!.$$

Since $t' \leq t + 1$, dividing the lower bound by the upper bound gives

$$\prod_{i=1}^{t'-2} \frac{N - e_i - e_{i+1}}{N - e_i} = \prod_{i=1}^{t'-2} \big(1 - \frac{e_{i+1}}{N - e_i}\big) \geq \big(1 - \frac{2q}{N - 2q}\big)^{t-1} \geq 1 - \frac{2qt}{N - 2q}.$$

Therefore,

$$\frac{|\mathcal{K}_1|}{|\mathcal{K}_2|} \geq 1 - \frac{2qt}{N - 2q}. \tag{15}$$

Since $q < N/100$, $N - 2q > \frac{49}{50}N$, and therefore $(2qt)/(N - 2q) < 2.05qt/N = \rho$. Thus (15) implies (14), which concludes the proof.

Lemmas 1 and 2 immediately imply:

Lemma 5. *Let $(X_\alpha, Y_\alpha)_{\alpha \in B}$ be the family of random variable pairs described before the statement of Lemma 2 (parameters for which are N, t and q), with $q < N/100$. Then*

$$\boldsymbol{Adv}^{\mathsf{samp-dist}}_{(X_\alpha, Y_\alpha)_{\alpha \in B}}(q) \leq \frac{2q^2\rho}{N - 2q} \leq \frac{2.05q^2\rho}{N} \leq \frac{4.3q^3t}{N^2}$$

(where $\rho = 2.05qt/N$).

A sampling distinguishing adversary for $(X_\alpha, Y_\alpha)_{\alpha \in B}$ can obviously simulate a "standard" adversary for the $O(N, t)$-$\tilde{O}(N, t)$ distinguishing with keys $k_0, \ldots, k_t = 0^n$, with equal advantage (see the remarks before Lemma 2). Thus, we obtain the following corollary, that completes the proof of Proposition 2.

Corollary 2. *For $q < N/100$, we have $\boldsymbol{Adv}^{O\tilde{O};0^n}_{N,t}(A) \leq \frac{4.3q^3t}{N^2}$.*

An example where adaptivity helps for sample distinguishability. We conclude by showing, for general interest, an example for which adaptivity helps in the sample distinguishability game.

We use only two pairs of random variables (X_1, Y_1), (X_2, Y_2) taking values in a range $S = \{a, b, c\}$. Let $\varepsilon, \varepsilon', \delta > 0$ with $\varepsilon' < \varepsilon$. Define:

$$\begin{array}{lll} \Pr[X_1 = a] = 1 - \delta & \Pr[X_1 = b] = 0 & \Pr[X_1 = c] = \delta \\ \Pr[Y_1 = a] = 1 - \delta - \varepsilon' & \Pr[Y_1 = b] = \varepsilon' & \Pr[Y_1 = c] = \delta \end{array}$$

and

$$\Pr[X_2 = a] = \tfrac{1}{2} + \varepsilon \ \Pr[X_2 = b] = \tfrac{1}{2} - \varepsilon \ \Pr[X_2 = c] = 0$$
$$\Pr[Y_2 = a] = \tfrac{1}{2} - \varepsilon \ \Pr[Y_2 = b] = \tfrac{1}{2} + \varepsilon \ \Pr[Y_2 = c] = 0.$$

We put ε small (so that ε^2 is negligible) and put $\varepsilon' = 1.99\varepsilon$. We also put $\delta = 0.1\varepsilon^2$. For ε sufficiently small, we have that $\varepsilon' + \delta = \Delta(X_1, Y_1) < \Delta(X_2, Y_2) = 2\varepsilon$.

We give the adversary two queries. The best non-adaptive strategy is then for the adversary to query (X_1, Y_1) twice, even though $\Delta(X_1, Y_1) < \Delta(X_2, Y_2)$. Indeed, $\Delta(X_1^2, Y_1^2) \approx 4\varepsilon'$ whereas $\Delta(X_2^2, Y_2^2) \approx 4\varepsilon < 4\varepsilon'$ and $\Delta(X_1X_2, Y_1Y_2) \approx 6\varepsilon < 4\varepsilon'$.

On the other hand, choosing (X_1, Y_1) twice can be improved upon with an adaptive strategy, since if the adversary sees c after its first query to (X_1, Y_1) it is better for the adversary to query (X_2, Y_2), given that $\Delta(X_1, Y_1) < \Delta(X_2, Y_2)$ and that $\Pr[X_1 = c] = \Pr[Y_1 = c]$.

B Proof of Theorem 2

Consider a fixed point (β_0, β_t), $\beta_0, \beta_t \neq 0$, in the Fourier spectrum for the t-round key-alternating cipher with keys $K := (k_0, \ldots, k_t)$. Denote by β_i, $1 \leq i < t$, the intermediate selection pattern at the addition of k_i, and set $\beta := (\beta_1, \ldots, \beta_{t-1})$ and $\Gamma := (\beta_0, \ldots, \beta_t)$. By the theorem of trail composition (Theorem 7.8.1 in [13]), we have

$$W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}[K] = 2^{n(1-t)} \sum_{\beta} W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t} \cdot (-1)^{\Gamma^T K}, \tag{16}$$

with $W^{P_i}_{\beta_{i-1},\beta_i}$ denoting the Fourier coefficient of P_i at point (β_{i-1}, β_i). For each $\beta \neq 0$, define the random variable X_β as

$$X_\beta := W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t} \cdot (-1)^{\Gamma^T K}, \tag{17}$$

so that

$$W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}[K] = \sum_{\beta} X_\beta. \tag{18}$$

If, for any given key K, the quantities $\Gamma^T K$ behave independently over different β, as assumed in the claim of the theorem, we have that

$$X_\beta \sim_K W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t} \cdot (-1)^r, \tag{19}$$

with $r \sim \mathrm{Bern}(\frac{1}{2})$, where the distribution is taken over the keys, and $\mathrm{Bern}(p)$ denotes the Bernoulli distribution with success probability p.

Note that $\mathbf{E}[X_\beta] = \frac{1}{2}(W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t} - W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t}) = 0$. The variance of X_β is given by

$$\begin{aligned}\mathbf{Var}[X_\beta] &= \frac{1}{2}\left(W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t}\right)^2 + \frac{1}{2}\left(-W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t}\right)^2 \\ &= \left(W^{P_1}_{\beta_0,\beta_1} \cdots W^{P_t}_{\beta_{t-1},\beta_t}\right)^2.\end{aligned}$$

Furthermore, with $b := 2^{tn} + 1$, we have

$$\lim_{m\to\infty} \Pr(|X_m| < b) = 1, \tag{20}$$

as each of the t multiplicands $W^{P_i}_{\beta_{i-1},\beta_i}$ of X_m are bounded by 2^n. On the other hand, the variance of all partial sums is unbounded by assumption (6) that $\mathbf{Var}_{\beta_{i-1},\beta_i}\left[W^{P_i}_{\beta_{i-1},\beta_i}\right] \geq 2^{n/2}$ and a standard comparison test:

$$\lim_{m\to\infty} \sum_{i=1}^{m} 2^{n/2} = \infty \Longrightarrow \lim_{m\to\infty} \mathbf{Var}\left[\sum_{i=1}^{m} X_i\right] = \lim_{m\to\infty} \sum_{i=1}^{m} \mathbf{Var}\left[X_i\right] = \infty. \tag{21}$$

A sequence of independent (one can consider the X_β as independent since the signs are independent) random variables fulfilling (20) and (21) obeys the Lindeberg formulation of the central limit theorem [31, p. 488] (note that though we operate with finite numbers of summands, the conditions at infinity have to be checked for any application of the central limit theorem). Therefore, we have the following approximation, since the number of summands is high (it is exponential in n and in all interesting cases $n \geq 32$):

$$\sum_{\beta} X_\beta \sim_K \mathcal{N}(0, s^2) \tag{22}$$

with $s^2 := \sum_\beta \mathbf{Var}[X_\beta]$. The mean of s^2 over all permutations $Q_1, \ldots, Q_t$ can now be determined as $\mathbf{E}_{Q_1,\ldots,Q_t}[s^2] = \mathbf{E}_{Q_1,\ldots,Q_t}\left[\sum_\beta \left(W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right)^2\right] = \sum_\beta \mathbf{E}_{Q_1,\ldots,Q_t}\left[\left(W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right)^2\right] = \sum_\beta \mathbf{Var}_{Q_1,\ldots,Q_t}\left[W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right] + \left(\mathbf{E}_{Q_1,\ldots,Q_t}\left[W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right]\right)^2$ by linearity of expectation and definition of variance. By Fact 1, $W^{Q_i}_{\beta_{i-1},\beta_i} \sim_{Q_i} \mathcal{N}(0, 2^n) = 2^{n/2}\mathcal{N}(0,1)$ for each i, so $W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t} \sim 2^{t(n/2)}\mathcal{N}(0,1)\cdots\mathcal{N}(0,1)$, where the product is over t standard normal distributions. The mean of this distribution is zero, and the variance of the product of two independent standard normal distributions $Z := \mathcal{N}(0,1)\mathcal{N}(0,1)$ can be calculated via its moment-generating function

$$M_Z(y) = \frac{1}{2\pi}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-\frac{1}{2}x_1^2 - \frac{1}{2}x_2^2} e^{y x_1 x_2}\, dx_1\, dx_2 = \frac{1}{\sqrt{1-y^2}}.$$

Expanding the logarithm of $M_Z(y)$ in a power series in y, we find

$$\ln(M_Z(y)) = \sum_{n=0}^{\infty} m_n \frac{y^n}{n!} = \sum_{k=1}^{\infty} \frac{1}{2k} y^{2k} = \frac{1}{2}y^2 + \frac{1}{4}y^4 + \cdots,$$

and therefore $\mathbf{Var}[Z] = 1$. The same applies to $t > 2$. Consequently, $\mathbf{Var}_{Q_1,\ldots,Q_t}\left[W^{Q_1}_{\beta_0,\beta_1} \cdots W^{Q_t}_{\beta_{t-1},\beta_t}\right] = \left(2^{t(n/2)}\right)^2 \cdot 1 = 2^{tn}$ for each β. Note that we have $(2^n - 1)^{t-1}$ values of β with no $\beta_i = 0$, so $\mathbf{E}[s^2] = (2^n - 1)^{t-1}2^{nt}$.

Recall from (7) that for the t-round cipher with the permutations $P_1, \ldots, P_t$, we have that $s^2 = \sum_\beta \mathbf{Var}[X_\beta] = (1+\varepsilon)\,\mathbf{E}_{Q_1,\ldots,Q_t}[s^2]$. The distribution of $W^{P_1,\ldots,P_t}_{\beta_0,\beta_t}$ over all keys is therefore given by

$$\begin{aligned} W^{P_1,\ldots,P_t}_{\beta_0,\beta_t} &\sim_K 2^{n(1-t)}\mathcal{N}(0,(1+\varepsilon)(2^n-1)^{t-1}2^{nt}) \\ &= \mathcal{N}(0,(1+\varepsilon)\left(\frac{2^n-1}{2^n}\right)^{t-1}2^n), \end{aligned}$$

as claimed. □

We require condition (6) essentially to ensures that we sum over sufficiently many possible selection patterns for β such that we can invoke the central limit theorem. This in particular excludes the trivial case where all P_i are linear, in which their variances would be zero, and the sum in (16) would only have one summand.

C Attacks on the variants of the double construction

C.1 Attack on variant with $Q = P$ and $k_0 \oplus k_2 = \alpha$ known

This variant succumbs to a (variant of the) slide attack. The assumptions of the attack are that $P = Q$ and that $k_0 \oplus k_2 = \alpha$ is known.

Slide attacks consider *slid pairs*. A slid pair is a pair of encryptions such that an intermediate value in one encryption equals the plaintext value of the other encryption.

In our case a slid pair is two encryptions (m, c) and $(\tilde{m}, \tilde{c})$ such that

$$P(m \oplus k_0) = \tilde{m} \oplus k_0 \oplus k_1 \tag{23}$$
$$P(c \oplus k_1 \oplus k_2) = \tilde{c} \oplus k_2. \tag{24}$$

Since P is bijective this is the same as

$$P(m \oplus k_0) = \tilde{m} \oplus k_0 \oplus k_1 \tag{25}$$
$$P^{-1}(\tilde{c} \oplus k_2) = c \oplus k_1 \oplus k_2. \tag{26}$$

This implies that for a slid pair it holds that

$$P(m \oplus k_0) \oplus P^{-1}(\tilde{c} \oplus k_2) = c \oplus \tilde{m} \oplus \alpha \tag{27}$$

In an attack one tries to identify a slid pair, which gives candidate values for the secret key. The attack proceeds as follows.

1. Compute a sorted table T consisting of the elements $b_i = P(a_i) \oplus P^{-1}(a_i)$ for $i = 1, \ldots, 2^{n/2}$, where a_i are randomly chosen values.
2. Get the encryptions c_i for $2^{n/2}$ arbitrary messages m_i for $i = 1, \ldots, 2^{n/2}$.
3. Get the decryptions $\tilde{m}_i$ for ciphertexts $\tilde{c}_i$, where $\tilde{c}_i = m_i \oplus \alpha$ for $i = 1, \ldots, 2^{n/2}$.

4. Find pairs (i, j) such that $c_i \oplus \tilde{m}_i \oplus \alpha = b_j$.
5. For each match:
   (a) Set $k'_2 = \tilde{c}_j \oplus a_i$;
   (b) Set $k'_0 = m_j \oplus a_i$;
   (c) From one encryption (m', c'), compute k'_1 from (k'_0, k'_2), i.e., $k'_1 = P(m' \oplus k'_0) \oplus P^{-1}(c' \oplus k'_2)$;
   (d) Test the computed values (k'_0, k'_1, k'_2) one additional encryptions.

We expect to get one slid pair in the above collection of known and chosen texts. There may be other matches but they are easily discarded in a test on additional encryptions.

This is a chosen ciphertext attack of complexity roughly $2^{n/2}$. (There is a similar attack which uses chosen plaintexts instead of chosen ciphertexts.)

C.2 Attack on variant with $Q = P^{-1}$

Note that this variant has a key size of $3n$. A meet in the middle attack has complexity 2^n.

Here is an attack which finds n bits of the key using $2^{n/2}$ encryptions. After that, one can easily distinguish the cipher from random.

Set $k_0 \oplus k_2 = \alpha$. Let (m, c) and (m', c') be two arbitrary encryptions, where $m \neq m'$. It follows that if $m \oplus c' = \alpha$, then this implies that $m' \oplus c = \alpha$. In a chosen plaintext-ciphertext attack, one can find α using $2^{n/2}$ queries.

1. Choose $2^{n/2}$ messages, $m_i = (i \mid m_0)$, where $i = 1, \ldots, 2^{n/2}$ and m_0 is an $(n/2)$-bit constant. Get the corresponding encryptions c_i.
2. Choose $2^{n/2}$ ciphertexts, $c'_j = (c_0 \mid i)$, where $i = 1, \ldots, 2^{n/2}$ and c_0 is an $(n/2)$-bit constant. Get the corresponding messages m'_j.
3. Find a match (i, j) such that $m_i \oplus c_i = m'_j \oplus c'_j$. For each match compute a candidate value of $k_0 \oplus k_2 = m_i \oplus c'_j$.
4. Note that α will appear as one of the candidate values in the previous step. Repeat the attack, until only one candidate value, namely α, remains.

When the value of α is found, the cipher is easily distinguished from the ideal cipher. Let m be a message and c the corresponding ciphertext. Then the message $c \oplus \alpha$ will be encrypted to $m \oplus \alpha$.

C.3 Related-key attacks

In certain scenarios one considers also related-key attacks where the adversary is allowed to get encryptions under several related keys. In the case where all round-keys are independent, related-key attacks exist trivially. Thus, we here focus on the case of identical round-keys. Furthermore, we restrict to the case of $t = 2$, as this is the case which is most relevant for practical purposes. The following attack requires that an attacker can get encryptions under a key $k = (k_0, k_0, k_0)$ and under a key $\tilde{k} = (k_0 \oplus \alpha, k_0 \oplus \alpha, k_0 \oplus \alpha)$ for a known value of α.

1. Assume that attacker can get encryptions under k and under $\tilde{k}$ for a known and fixed value of α.
2. Compute a sorted table T with entries $P(x) \oplus P(x \oplus \alpha) \oplus \alpha$ for 2^a distinct, randomly chosen values of x.
3. Choose 2^b messages m_i and get the corresponding encryptions c_i under k.
4. Choose 2^b messages $\tilde{m}_i = m_i \oplus \alpha_0$ and get corresponding encryptions $\tilde{c}_i$ under $\tilde{k}$.
5. Find a match between the values $c_i \oplus \tilde{c}_i$ and the values in T.
6. For each match, find potential values of the key and test these values on further encryptions.

Following the birthday bound, using roughly $a = b = n/2$ one gets a probability of success of about one half.

D Implementation of AES^2 with AES-NI

On the Westmere architecture generation of Intel general-purpose processors, AES^2 can be implemented using the AES-NI instruction set [18]. Since the key schedule for the 22 AES round keys can be precomputed, the cipher basically only consists of 18 `aesenc` and 2 `aesenclast` instructions, bracketed by 5 XORs with the three keys k_0, k_1, k_2 and the two (constant) first AES subkeys. The AES round instructions are pipelined, with a documented latency of 6 cycles and throughtput 2. Practical measurements using recent Westmere processors indicate an actual latency of 4. Therefore, we can fully utilise the pipeline by processing multiple independent plaintext blocks in parallel in the basic electronic codebook mode (ECB) and counter mode (CTR).

All performance figures were obtained by using one core, with hyperthreading and Turbo Boost disabled to ensure fair comparison.