Advanced cryptology (GBX9SY06)

Block ciphers

Pierre Karpman

2017-11

1 First definitions

A block cipher is a family of invertible mappings indexed by a key: $\mathscr{E} : \mathscr{K} \times \mathscr{M} \to \mathscr{C}$, where both $\mathscr{E}$ and its inverse $\mathscr{E}^{-1}$ are "efficiently computable"[1]. In practice, we only care about ciphers for which $\mathscr{M} = \mathscr{C}$, meaning that $\mathscr{E}$ defines a family of *permutations* over $\mathscr{M}$. In the vast majority of cases, we also have $\mathscr{M} = \{0,1\}^n$ for some integer n, often equal to 64 or 128 (though smaller, larger, and intermediate values are possible). In some corner cases, however, it may be that $\mathscr{M}$ has a different structure, for instance all integers smaller than a certain N, or the set of valid credit-card numbers. This is the concern of *format-preserving encryption*, which will not be addressed in this course. Note that as long as cryptography is implemented on digital circuits, there is no similar incentive to take $\mathscr{K}$ of another form than $\{0,1\}^\kappa$. To summarize, we will use the following definition.

Definition 1 (Block ciphers)**.** A block cipher is a family of mappings $\mathscr{E} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ such that for all $k \in \{0,1\}^\kappa$, $\mathscr{E}(k,\cdot)$ is a permutation. The quantities n and κ are positive integers, called the block size (or length) and key size, respectively.

Block ciphers are important *primitives* in symmetric-key cryptography. When used with a suitable *mode of operation*, they allow to ensure the confidentiality and authenticity of data. We will not address the issue of modes in this course, but it is important to remember that they are essential; without a proper mode, a block cipher is mostly useless.

There are many constructions that satisfy Definition 1, but not all of them are useful in a cryptographic context. For instance, $\mathscr{E}$ such that for all k the mapping $\mathscr{E}(k,\cdot)$ is the identity is not terribly good at providing confidentiality. We thus need to express one or several security properties that should hold for a "good" block cipher. Intuitively and informally, we typically require that:

1. Given $x_0, \ldots, x_m, \mathscr{E}(k,x_0), \ldots, \mathscr{E}(k,x_m)$, it should be "hard" to find k, even if the x_is span the entire message space $\{0,1\}^n$. An attack violating this property is called a *key-recovery attack*[2].

2. Given $x_0, \ldots, x_m, \mathscr{E}(k,x_0), \ldots, \mathscr{E}(k,x_m)$ and y_0 (resp. $\mathscr{E}(k,y_0)$), it should be "hard" to learn information about $\mathscr{E}(k,y_0)$ (resp. y_0) (except that it is distinct from the $\mathscr{E}(k,x_i)$s (resp. x_is)). An attack violating this property is usually called a *distinguishing* attack.

The first of these informal properties is maybe the most obvious one to think of, but it is not sufficient in itself to capture all the desired properties of a block cipher. For instance, one may

[1] As concrete block ciphers (usually) fix all their parameters, there is not much sense to argue about this efficiency in terms of asymptotic complexity, and we will not attempt to give a precise definition of what "efficient" means in this context.

[2] More generally, we may require that it should be hard to recover k' s.t. $\mathscr{E}(k,\cdot) = \mathscr{E}(k',\cdot)$ on most inputs.

consider a cipher $\mathscr{E}$ for which key-recovery attacks are hard, but such that for all k, $\mathscr{E}(k,m) = m$ whenever m starts with a zero bit. It is quite clear that such an $\mathscr{E}$ does not provide adequate confidentiality.

In order to be more precise about our requirements for good block ciphers, it is useful to first consider *ideal* block ciphers. This allows to set a standard for any security property we might be interested in, as these should be "maximally hard" for an ideal cipher. A concrete (non-ideal) cipher $\mathscr{E}$ is then considered to be secure if it is hard to decide if one is interacting with $\mathscr{E}$ or with an optimally-secure ideal cipher.

The definition of an ideal cipher is quite simple: it is simply a block cipher $\mathfrak{E}$ such that for all k, $\mathfrak{E}(k,\cdot)$ is a permutation uniformly drawn at random among all permutations over the same domain. Thus, all keys of $\mathfrak{E}$ define completely independent mappings that are all equally likely to be selected. Fixing the notation, we have the following definition.

Definition 2 (Ideal block cipher)**.** Let Π_n denote the set of all $2^n!$ permutations over $\{0,1\}^n$. For any finite set $\mathscr{S}$, we write $x \xleftarrow{\$} \mathscr{S}$ the uniform sampling of x over all elements of $\mathscr{S}$. Then, an ideal block cipher is a cipher $\mathfrak{E} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ such that for all $k \in \{0,1\}^\kappa$, $\mathfrak{E}(k,\cdot) \xleftarrow{\$} \Pi_n$.

Note that not all ideal ciphers are in fact secure. For instance, when randomly sampling an ideal cipher, there is a $1/(2^n! \times 2^\kappa)$ chance that $\mathfrak{E}$ be defined as the insecure identity-everywhere cipher. However, we claim that the odds of picking an "insecure" cipher are small enough for this eventuality to be ignored altogether.

As already said above, one may use the notion of an ideal cipher to define the security of a concrete one by how easy it is for an adversary to decide if he is interacting with an ideal cipher or not. This corresponds to the notion of *pseudorandom permutation* (PRP). Informally, this notion is expressed by having an adversary interacting with an *oracle* $\mathcal{O}$ which is randomly chosen to be either of $\mathscr{E}(K,\cdot)$ or $\mathfrak{E}(K,\cdot)$ for a random key K (as $\mathfrak{E}$ is ideal, the latter case is simplified by dropping the key and simply asking that $\mathfrak{E}$ be a uniformly chosen permutation). Then one considers what is the best *advantage* over a random choice that an adversary has of deciding how $\mathcal{O}$ was instantiated.

It seems clear that the advantage should in fact be a function of the *data complexity* (the number of queries to $\mathcal{O}$) and of the *time complexity* (where the unit is generally the time it takes to evaluate $\mathscr{E}$) of the adversary. For instance, this accounts for the fact that any cipher (even an ideal one) of key size κ can be broken by exhaustive key search in time 2^κ, and yet could still be secure if the adversary only has more limited resources. We then define the PRP security of $\mathscr{E}$ through the following PRP advantage function (see e.g. [BKR00, BR]).

Definition 3 (PRP advantage)**.** The *PRP advantage* of $\mathscr{E} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ is defined as:

$$\mathbf{Adv}^{\mathrm{PRP}}_{\mathscr{E}}(q,t) := \max_{A_{q,t}} \left| \Pr\left[A^{\mathcal{O}}_{q,t}(\cdot) = 1 : \mathcal{O} \xleftarrow{\$} \Pi_n \right] - \Pr\left[A^{\mathcal{O}}_{q,t}(\cdot) = 1 : \mathcal{O} = \mathscr{E}(K,\cdot), K \xleftarrow{\$} \{0,1\}^\kappa \right] \right|,$$

where $A^{\mathcal{O}}_{q,t}$ denotes an adversary A who makes at most q queries to its oracle $\mathcal{O}$, runs in time at most t, and returns a unique bit.

A similar notion of *strong* PRP (SPRP) can be defined, where the adversary is additionally granted access to the inverse of $\mathcal{O}$.

Now we can use Definition 3 to formulate some requirements about the security of a "good" block cipher $\mathscr{E}$. As, we would ideally want $\mathscr{E}$ to behave as close as possible to an ideal cipher $\mathfrak{E}$ (defined over the same domains), i.e. we wish that $\mathbf{Adv}^{\mathrm{PRP}}_{\mathscr{E}}(q,t) \approx \mathbf{Adv}^{\mathrm{PRP}}_{\mathfrak{E}}(q,t)$ for all q, t (for instance, we would like to have $\mathbf{Adv}^{\mathrm{PRP}}_{\mathscr{E}}(1,t) \approx t/2^\kappa$). By abuse of terminology, we usually say of such an hypothetical cipher that "it is a PRP".

Note that breaking PRP security does not even require to recover the key k when $\mathcal{O} = \mathscr{E}(k,\cdot)$, which was one of the informal goals stated at the beginning of this section. As recovering the key *does* however allow to break PRP security, focusing only on the latter does not weaken the

requirements on $\mathscr{E}$, and allows to capture the second informal goal of resisting distinguishing attacks.

The notion of PRP is useful to express desirable properties for concrete block ciphers. However, it is in itself useless to actually evaluate their security. That it is to say, the definition provides very little insight into how to compute e.g. $\mathbf{Adv}^{\text{PRP}}_{\text{AES-128}}$. It is in fact the goal of cryptanalysis in general, by finding explicit attacks, to lower-bound the advantage function for specific ciphers and complexities.

2 An ideal construction

We now present a generic construction of block ciphers due to Even and Mansour [EM91, EM97], which is very simple yet of considerable interest. Let $\mathscr{P} : \{0,1\}^n \to \{0,1\}^n$ be a "public" (i.e. not secret) permutation, then one simply defines the Even-Mansour cipher built from $\mathscr{P}$ as $\mathscr{E}(k_1||k_2, m) := \mathscr{P}(m \oplus k_1) \oplus k_2$, for all keys $(k_1||k_2)$ of $2n$ bits and all messages m. In fact, we can also go for an even simpler definition by taking $k_1 = k_2$, thence obtaining $\mathscr{E}(k, m) = \mathscr{P}(m \oplus k) \oplus k$. What is remarkable is that both of these constructions can be proven to be good PRPs, assuming that $\mathscr{P}$ is itself a "good" permutation. Intuitively, this implies for instance that $\mathscr{P}$ should not admit efficient distinguishers, in line with our second intuitive requirement for a block cipher. However, this is not an easy notion to formalize; the approach taken by Even and Mansour to prove their construction is then to only allow an adversary to make black-box oracle queries to $\mathscr{P}$. A possible interpretation is then to consider this number of queries itself to denote the "time complexity" of the attack.

2.1 Proof sketch

The security theorem proved by Even and Mansour does not use the notion of PRP, which was not formalised at the time. Instead, it proves that the success probability of an adversary for the *existential forgery problem* (EFP) is upper-bounded by $O(DT/2^n)$, where D is the number of queries to $\mathscr{E}$ or its inverse (with an unknown fixed random key) and T the number of black-box queries to $\mathscr{P}$ or its inverse. In other words, the scheme achieves security up to the birthday bound. Note however that the proof can be adapted to a PRP setting, but we will rather focus on the original formulation.

The EFP attack considered here does the following: the adversary is first given access to $\mathscr{E}(k,\cdot)$, $\mathscr{E}^{-1}(k,\cdot)$, $\mathscr{P}$ and $\mathscr{P}^{-1}$ to make any query he wants; then, he must provide a pair (p, c) such that $\mathscr{E}(k, p) = c$, and neither p nor c was queried for the $\mathscr{E}$ or $\mathscr{E}^{-1}$ oracle respectively. One can show without much trouble that this attack is not easier than decrypting a challenge ciphertext without knowing the key, or more generally recovering an unknown key.

We now briefly sketch the proof. The idea is to keep track of all queries made by the adversary and to count the number of possible *bad* keys. These are the keys whose consistency with $\mathscr{P}$ and $\mathscr{E}$ was "checked". After the last query is made, if the actual key is marked as bad, we assume that the adversary won. Otherwise, one can show that all *good* (i.e. not bad) keys have the same probability of being the actual key, and their number determines the success probability of the adversary. More precisely:

- A key k is *bad* if the adversary made a query relating x and y through $\mathscr{P}$ (i.e. queried $\mathscr{P}$ on x or $\mathscr{P}^{-1}$ on y) and a query relating x' and y' through $\mathscr{E}$ such that at least one of $x \oplus x'$ and $y \oplus y'$ is equal to k. Indeed, if k were the correct key used in $\mathscr{E}$, we would have $y' = \mathscr{E}(k, x') = \mathscr{P}(x' \oplus k) \oplus k$, so $x = x' \oplus k \Leftrightarrow y' = y \oplus k$, which is a property that can easily be checked from the queries.
- A key is *good* if it is not bad.

One can then see that the number of bad keys is at most approximately equal to the number of queries to $\mathscr{P}$ times the number of queries to $\mathscr{E}$, and the theorem mostly follows.

2.2 Generic attacks

An important and common question that arises in the context of security proofs such as the one above is the question of *tightness*. That is, we are interested in knowing if the actual security of a construction might be better than what the proof provides, or if there exists an attack with a complexity matching the provable bound. Ideally, we would like all proofs to be tight. If the status of a proof is not known, it is a natural research problem to either find an attack matching the bound of the proof or a better proof matching less efficient attacks.

In the case of the Even-Mansour construction, it was quickly found out by Daemen that the proof is indeed tight [Dae91]. In fact, it is quite straightforward to extract an "optimal" attack strategy from the proof sketch as given above. Several variants of what is generically called *slide attacks* exist (see e.g. [DKS12]), and we sketch two of them. The first one works as follows:

1. Pick N distinct random values x_i, query $\mathscr{E}(x_i)$ and $\mathscr{P}(x_i)$, and insert their sums $y_i := \mathscr{E}(x_i) \oplus \mathscr{P}(x_i)$ along with x_i in a table.
2. For all i, j such that $y_i = y_j$, return $x_i \oplus x_j$ as a key candidate.

First, let us show that if the unknown key k is equal to $x_i \oplus x_j$ for some (i, j), the test in step 2) will indeed return k as a candidate. This is straightforward, as by definition we have:

$$y_i = \mathscr{P}(x_i \oplus k) \oplus k \oplus \mathscr{P}(x_i) = \mathscr{P}(x_j) \oplus k \oplus \mathscr{P}(x_j \oplus k) = y_j.$$

Now what is the probability (in function of N) that this event happens? As the x_is are chosen randomly, their differences form $N \cdot (N-1)/2$ (non-necessarily distinct) random values, hence the probability that one of these is equal to k is roughly $N^2/2^n$. This already matches the bound of the security proof, however, we can go further and show that the attack can still be efficiently implemented if one additionally considers its time[3] and memory complexity (which are not accounted for in the theorem). First, by using a suitable data structure (e.g. a hash table), one only needs $O(N)$ memory to store the N (x_i, y_i) pairs, and every collision can be detected in constant time. Thus, we only need to look at the expected number of collisions and show that it is not more than N. If we make the reasonable assumption that $x \mapsto \mathscr{E}(x) \oplus \mathscr{P}(x)$ behaves as a random function, then the expected number of collisions is $O(N^2/2^n)$, which is much smaller than N for $N \ll 2^n$. We thus only expect a few "false positive" suggestions for k before finding the correct value and stopping the attack. Putting everything together, this attack has a time, memory, and data complexity of N, and a success probability upper-bounded by $O(N^2/2^n)$.

We only describe, and do not analyse the second attack. Unlike the first, it actually requires the keys k_1 and k_2 to be identical (the previous did not, even though it was presented in this simplified case), but it allows to trade queries to $\mathscr{E}$ for queries to $\mathscr{P}$. That is, it succeeds with probability upper-bounded by $DT/2^n$ with D and T free to take any value, whereas the previous attack required $D = T$. This attack works as follows:

1. Pick D distinct random values x_i, query $\mathscr{E}(x_i)$, and insert $y_i := \mathscr{E}(x_i) \oplus x_i$ along with x_i in a table.
2. Pick T distinct random values x'_i, query $\mathscr{P}(x'_i)$, and insert $y'_i := \mathscr{P}(x'_i) \oplus x'_i$ along with x'_i in a table.
3. For all i, j such that $y_i = y'_j$, return $x_i \oplus x'_j$ as a key candidate.

[3] We referred to the black-box queries to $\mathscr{P}$ as the "time complexity". While it is indeed reasonable to consider this as a *lower-bound* on the time complexity of an attack that does not exploit structural properties of $\mathscr{P}$, an actual attack algorithm such as the one above might imply additional processing beyond computations of $\mathscr{P}$.

2.3 Generalisations

There are several ways to generalise the Even-Mansour constructions, beyond the simple variant taking $k_1 = k_2$ that we already considered. One direction is to consider a block cipher $\mathscr{E}' : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ instead of a public permutation $\mathscr{P}$. This gives the "FX" construction $\mathscr{E}(k||k_1||k_2, m) := \mathscr{E}'(k, m \oplus k_1) \oplus k_2$, which was first suggested by Rivest and then analysed by Kilian and Rogaway [KR01]. Using a method similar to the one of Even and Mansour, they proved that the PRP advantage (this time exactly in the sense of Definition 3) of an adversary attacking $\mathscr{E}$ is upper-bounded by $D \cdot T/2^{\kappa+n-1}$, where D represents the number of queries to the unknown oracle $\mathscr{O}$ and T represents the "time complexity" of black-box accesses to the unkeyed $\mathscr{E}'$. One can see in particular that by taking $\kappa = 0$, i.e. fixing $\mathscr{E}'$ to be a public permutation, this upper-bound essentially matches the one of the Even-Mansour scheme.

Another natural way to generalise the scheme is to compose it with independent instantiations of itself, i.e. defining $\mathscr{E}(k_1||k_2||\ldots||k_{r+1}, m)$ as $\mathscr{P}_r(\ldots\mathscr{P}_1(m \oplus k_1) \oplus k_2)\ldots) \oplus k_{r+1}$. Chen and Steinberger showed that this achieved a PRP security up to $O(2^{\frac{rn}{r+1}})$ queries [CS14], while Lampe and Seurin showed an *indifferentiability* property for the twelve-round iterated scheme that uses distinct permutations $\mathscr{P}_1, \ldots, \mathscr{P}_{12}$ but equal keys $k_1 = \ldots = k_{13}$ [LS13]. However, interestingly, the simplest and most economical way to compose an Even-Mansour scheme, i.e. taking r identical permutations and $r+1$ equal keys does *not* work, in the sense that it does not provide much more security than the original non-iterated scheme. Let us show why.

Consider $\mathscr{E}(k, m) := \mathscr{P}(\ldots\mathscr{P}(m \oplus k)\ldots) \oplus k$ for an arbitrary large number of applications of $\mathscr{P}$. We first observe that if two values x and x' are related by $x' = x \oplus k$, then $\mathscr{E}(\mathscr{P}(x')) = \mathscr{P}(\mathscr{E}(x)) \oplus k$, and by symmetry $\mathscr{E}(\mathscr{P}(x)) = \mathscr{P}(\mathscr{E}(x')) \oplus k$. It follows that:

$$\mathscr{E}(\mathscr{P}(x)) \oplus \mathscr{P}(\mathscr{E}(x)) = \mathscr{P}(\mathscr{E}(x')) \oplus k \oplus \mathscr{P}(\mathscr{E}(x)) = \mathscr{E}(\mathscr{P}(x')) \oplus \mathscr{P}(\mathscr{E}(x')).$$

Thus, we can attack $\mathscr{E}$ with probability ≈ 1 by picking $2^{n/2}$ random values x and looking for collisions for the function $x \mapsto \mathscr{E}(\mathscr{P}(x)) \oplus \mathscr{P}(\mathscr{E}(x))$, each of them suggesting a candidate value for k. Note that unlike the attacks on the non-iterated Even-Mansour scheme, this one requires *chosen plaintexts* (but variants requiring only *known plaintexts* also exist, for instance by considering collisions between the lists $\{\mathscr{P}(\mathscr{E}(x)) \oplus x\}$ and $\{\mathscr{E}(x) \oplus \mathscr{P}^{-1}(x)\}$).

3 Related-key attacks

So far, we only focused on attacks from adversaries who could only access a single oracle, for instance corresponding to $\mathscr{E}(k, \cdot)$ for an unknown key k. A way to increase an adversary's power is then to allow access to a family of oracles, for instance corresponding to $\{\mathscr{E}(\varphi(k), \cdot),\ \varphi \in \Phi\}$, still for an unknown key k, and Φ a set of *related-key functions* $\{0,1\}^\kappa \to \{0,1\}^\kappa$. That is to say, the adversary is given access to $\mathscr{E}$ with several unknown keys that are related one to another through the elements of Φ. We call *related-key attack* an attack that corresponds to this model.

The security requirements associated with related-key attacks are the same as for the *single-key* case. For instance, we can informally state that it should be hard to recover the unknown k for a cipher to which we are given related-key oracle access. It is also possible to define a related-key (S)PRP notion that generalises the one of Definition 3, as was done by Bellare and Kohno [BK03]. We let again $\mathfrak{E}$ denote an ideal cipher, such that for all k, $\mathfrak{E}(k, \cdot) \xleftarrow{\$} \Pi_n$. We write $\mathfrak{E} \xleftarrow{\$} \Pi_n^\kappa$ the uniform sampling of such a cipher. We define a related-key oracle $\mathscr{E}_{\mathrm{RK}(\cdot,K)}(\cdot)$ for a block cipher $\mathscr{E} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ that takes a first input $\varphi : \{0,1\}^\kappa \to \{0,1\}^\kappa$ and a second input $m \in \{0,1\}^n$, and returns $\mathscr{E}(\varphi(K), m)$. We then have the following.

Definition 4 (Related-key PRP advantage restricted to Φ)**.** The *related-key PRP advantage* of $\mathscr{E}$ with respect to the related-key class Φ is defined as:

$$\mathbf{Adv}^{\text{PRP-RKA}}_{\Phi,\mathscr{E}}(q,t) := \max_{A_{q,t}} \left| \Pr\left[A^{\mathfrak{E}_{\mathrm{RK}(\cdot,K)}(\cdot)}_{q,t} = 1 : K \xleftarrow{\$} \{0,1\}^\kappa, \mathfrak{E} \xleftarrow{\$} \Pi_n^\kappa\right] - \Pr\left[A^{\mathscr{E}_{\mathrm{RK}(\cdot,K)}(\cdot)}_{q,t} = 1 : K \xleftarrow{\$} \{0,1\}^\kappa\right] \right|.$$

One can see that this definition is parameterized by the allowed related-key class Φ. In fact, this parameterization plays a major role in how secure a cipher can be according Definition 4: there are classes Φ with respect to which *no* cipher (even an ideal one) attains any meaningful level of security. For instance, assume that Φ includes a constant function $\varphi : x \mapsto c$; then an adversary may simply query its related-key oracle on φ and an arbitrary c' and compare the result with $\mathscr{E}(c, c')$; it succeeds with advantage close to one by answering 1 if the two values match and 0 otherwise.

There are less trivial related-key classes for which similar problems occur. One of the perhaps less intuitive examples is to consider $\Phi = \varphi^{\oplus} \cup \varphi^{+}$, where $\varphi^{\oplus} = \{x \mapsto x \oplus \Delta,\ \Delta \in \{0,1\}^{\kappa}\}$ and $\varphi^{+} = \{x \mapsto x + \Delta,\ \Delta \in \{0,1\}^{\kappa}\}$ (with $+$ denoting here the addition modulo 2^{κ}). Remarkably, non-trivial security *is* achievable w.r.t. $\varphi^{\oplus}$ or φ^{+} in isolation. We will not explore this formally, rather focusing on the intuition.

There are two key properties for Φ to allow non-trivial security. The first is that it be *collision-resistant*, and the second that it be *output-unpredictable*. Collision resistance means that it is hard to find two functions φ, φ' such that $\varphi(K) = \varphi'(K)$ for a uniform K. This is the property that does not hold for $\varphi^{\oplus} \cup \varphi+$ and that can be exploited to mount easy attacks using this class. However, it does hold trivially for, say $\varphi^{\oplus}$, as $\Delta \mapsto K \oplus \Delta$ is a permutation for any K. The output unpredictability property means that it is hard to guess $\varphi(K)$ for a uniform K; more generally, it should be hard to guess any of the elements of $\{\varphi(K), \varphi \in X \subseteq \Phi\}$ relatively to the size of X. It is easy to see that this property does not hold if Φ includes constant mappings, but that it does for $\Phi = \varphi^{\oplus}$.

We now illustrate the fact that if many related-key queries with different functions φ are allowed, the security of any cipher degrades significantly. Consider a simplified attack setting where the adversary interacts with a known cipher $\mathscr{E}$ with unknown key k and tries to recover k. To do so, he is allowed oracle access to $\mathscr{E}(\varphi(k), \cdot)$, $\varphi \in \varphi^{\oplus}$. A good (in fact basically optimal) attack that does not exploit any properties of $\mathscr{E}$ consists in querying $\mathscr{E}(k \oplus \Delta, 0)$ for $2^{\kappa/2}$ randomly selected Δ and putting the results in a list $\mathscr{L}$ along with the value Δ. Then the adversary tries $2^{\kappa/2}$ candidates k' for the key, each time computing $\mathscr{E}(k', 0)$. Any match with an element of $\mathscr{L}$ suggests a value for k, and one expects such a match with high probability.

3.1 Provably-secure constructions

Even when Φ is restricted to *meaningful* functions, the related-key model gives significantly more power to the adversary. The collision attack sketched above is an example; another one is to consider some very simple attacks on the Even-Mansour scheme. Let Φ include at least one function of the form $x \mapsto x \oplus \Delta$, then an adversary can distinguish an Even-Mansour scheme from an ideal cipher by querying $\mathscr{E}(k \oplus \Delta, \Delta)$, $\mathscr{E}(k, 0)$, and checking if the two values only differ by Δ.

This attack does not contradict the security proof of Even and Mansour, as it requires related-key queries that are not covered by the single-key proof. It does show however that a good cipher in the single-key setting might be terribly broken in a related-key setting, even w.r.t. "meaningful" related-key classes.

It is quite easy to see that the above attack also applies to iterated Even-Mansour schemes when independent keys are used at every round. A slightly more involved but also very efficient attack exists on the two-round scheme that uses identical keys. However, it was proven (in an ideal permutation model similar to the one of the single-key proof) that using three rounds or more with identical keys leads to a construction that is related-key secure w.r.t. $\varphi^{\oplus}$ [CS15, FP15].

3.2 Tweakable block ciphers

The fact that single-key and related-key security may be widely disconnected is part of the reason why not all concrete block ciphers are designed to be related-key secure. This comes

usually with the argument that resisting related-key attacks (for some a priori defined classes) would add some computational overhead and that the model is altogether unrealistic, as no properly designed protocol would allow an attacker to make related-key queries. We will not delve into this debate whose implications are rather complex, instead giving an example of a constructive use of related-key secure ciphers.

First, let us define *tweakable block ciphers*. These are simply block ciphers that take a second parameter called a *tweak*, such that all distinct pairs of keys and tweaks define (ideally) independent permutations. The difference between the key and the tweak is that the latter is public and may be freely chosen by the adversary. Using simplified expressions for the domains, we have the following.

Definition 5 (Tweakable block ciphers)**.** A tweakable block cipher is a family of mappings $\widetilde{\mathscr{E}} : \{0,1\}^\kappa \times \{0,1\}^\theta \times \{0,1\}^n \to \{0,1\}^n$ such that for all $k \in \{0,1\}^\kappa$, $t \in \{0,1\}^\theta$, $\mathscr{E}(k,t,\cdot)$ is a permutation.

Tweakable block ciphers are useful in that they allow to "diversify" a fixed-key instance without selecting a new, independent key. In other words, two communicating parties may first secretely share a secret key k with which to use $\widetilde{\mathscr{E}}$, and then publicly agree on a new tweak for each message to be exchanged. The concept of tweakable ciphers was formalised by Liskov et al. [LRW11], but already implicitly used for instance by Rogaway et al. to define an efficient mode of operation [RBBK01].

A simple way to build a tweakable block cipher $\widetilde{\mathscr{E}}$ from a "regular" cipher $\mathscr{E}$ is to define $\widetilde{\mathscr{E}}(k,t,\cdot)$ as $\mathscr{E}(k \oplus t,\cdot)$. The security of this construction fully reduces to the related-key security of $\mathscr{E}$ w.r.t. $\varphi^\oplus$, which may be non-trivial. For instance, this construction provably achieves a meaningful level of security in the ideal permutation model if $\mathscr{E}$ is a three-round iterated Even-Mansour scheme with identical keys. Note however that because of generic collision attacks, the security is limited w.r.t. the number of different tweaks for which $\widetilde{\mathscr{E}}$ is queried. The design of similar generic constructions that are secure beyond the birthday-bound is a rather active research topic.

References

[BK03] Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: Rka-prps, rka-prfs, and applications. In Eli Biham, editor, *Advances in Cryptology - EUROCRYPT 2003, International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4-8, 2003, Proceedings*, volume 2656 of *Lecture Notes in Computer Science*, pages 491–506. Springer, 2003.

[BKR00] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of the cipher block chaining message authentication code. *J. Comput. Syst. Sci.*, 61(3):362–399, 2000.

[BR] Mihir Bellare and Phillip Rogaway. Introduction to modern cryptography. Lecture notes available at `http://cseweb.ucsd.edu/~mihir/cse207/index.html`.

[CS14] Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 327–350. Springer, 2014.

[CS15] Benoit Cogliati and Yannick Seurin. On the provable security of the iterated even-mansour cipher against related-key and chosen-key attacks. In Elisabeth Oswald

and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 584–613. Springer, 2015.

[Dae91] Joan Daemen. Limitations of the even-mansour construction. In Imai et al. [IRM93], pages 495–498.

[DKS12] Orr Dunkelman, Nathan Keller, and Adi Shamir. Minimalism in cryptography: The even-mansour scheme revisited. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, volume 7237 of *Lecture Notes in Computer Science*, pages 336–354. Springer, 2012.

[EM91] Shimon Even and Yishay Mansour. A construction of a cioher from a single pseudorandom permutation. In Imai et al. [IRM93], pages 210–224.

[EM97] Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. *J. Cryptology*, 10(3):151–162, 1997.

[FP15] Pooya Farshim and Gordon Procter. The related-key security of iterated even-mansour ciphers. In Gregor Leander, editor, *Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers*, volume 9054 of *Lecture Notes in Computer Science*, pages 342–363. Springer, 2015.

[IRM93] Hideki Imai, Ronald L. Rivest, and Tsutomu Matsumoto, editors. *Advances in Cryptology - ASIACRYPT '91, International Conference on the Theory and Applications of Cryptology, Fujiyoshida, Japan, November 11-14, 1991, Proceedings*, volume 739 of *Lecture Notes in Computer Science*. Springer, 1993.

[KR01] Joe Kilian and Phillip Rogaway. How to protect DES against exhaustive key search (an analysis of DESX). *J. Cryptology*, 14(1):17–35, 2001.

[LRW11] Moses Liskov, Ronald L. Rivest, and David A. Wagner. Tweakable block ciphers. *J. Cryptology*, 24(3):588–613, 2011.

[LS13] Rodolphe Lampe and Yannick Seurin. How to construct an ideal cipher from a small set of public permutations. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 444–463. Springer, 2013.

[RBBK01] Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: a block-cipher mode of operation for efficient authenticated encryption. In Michael K. Reiter and Pierangela Samarati, editors, *CCS 2001, Proceedings of the 8th ACM Conference on Computer and Communications Security, Philadelphia, Pennsylvania, USA, November 6-8, 2001.*, pages 196–205. ACM, 2001.