

A Domain Extender for the Ideal Cipher

Jean-Sébastien Coron², Yevgeniy Dodis¹, Avradip Mandal²,
and Yannick Seurin³

¹ New York University

² University of Luxembourg

³ Orange Labs

Abstract. We describe the first domain extender for ideal ciphers, *i.e.* we show a construction that is indistinguishable from a $2n$ -bit ideal cipher, given a n -bit ideal cipher. Our construction is based on a 3-round Feistel, and is more efficient than first building a n -bit random oracle from a n -bit ideal cipher (as in [9]) and then a $2n$ -bit ideal cipher from a n -bit random oracle (as in [10], using a 6-round Feistel). We also show that 2 rounds are not enough for indistinguishability by exhibiting a simple attack. We also consider our construction in the standard model: we show that 2 rounds are enough to get a $2n$ -bit tweakable block-cipher from a n -bit tweakable block-cipher and we show that with 3 rounds we can get beyond the birthday security bound.

Keywords: Ideal cipher model, indistinguishability, tweakable block-cipher.

1 Introduction

A block cipher is a primitive that encrypts a n -bit string using a k -bit key. The standard security notion for block-ciphers is to be indistinguishable from a random permutation, for a polynomially bounded adversary, when the key is generated at random in $\{0, 1\}^k$. A block-cipher is said to be a strong pseudo-random permutation (or chosen-ciphertext secure) when computational indistinguishability holds even when the adversary has access to the inverse permutation.

When dealing with block-ciphers, it is sometimes useful to work in an idealized model of computation, in which a concrete block-cipher is replaced by a publicly accessible random block-cipher (or ideal cipher); this is a block cipher with a k -bit key and a n -bit input/output, that is chosen uniformly at random among all block ciphers of this form; this is equivalent to having a family of 2^k independent random permutations. All parties including the adversary can make both encryption and decryption queries to the ideal block cipher, for any given key; this is called the Ideal Cipher Model (ICM). Many schemes have been proven secure in the ICM [5, 11, 13, 15, 19, 20, 27]; however, it is possible to construct artificial schemes that are secure in the ICM but insecure for any concrete block cipher (see [4]). Still, a proof in the ideal cipher model seems useful because it

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-3-642-11799-2_36](https://doi.org/10.1007/978-3-642-11799-2_36)

shows that a scheme is secure against generic attacks, that do not exploit specific weaknesses of the underlying block cipher.

It was shown in [9,10] that the Ideal Cipher Model and the Random Oracle Model are equivalent; the random oracle model is similar to the ICM in that a concrete hash function is replaced by a publicly accessible random function (the random oracle). The authors of [9] proved that a random oracle (taking arbitrary long inputs) can be replaced by a block cipher-based construction, and the resulting scheme will remain secure in the ideal cipher model. Conversely, it was shown in [10] that an ideal cipher can be replaced by a 6-round Feistel construction, and the resulting scheme will remain secure in the random oracle model. Both directions were obtained using an extension of the classical notion of indistinguishability, called *indifferentiability*, introduced by Maurer *et al.* in [24].

Since a block cipher can only encrypt a string of fixed length, one must consider the encryption of longer strings. A *mode of operation* of a block-cipher is a method used to extend the domain of applicability from fixed length strings to variable length strings. Many modes of operations have been defined that provide both privacy and authenticity (such as OCB [28]). A mode of operation can also be a permutation; in this case, one obtains an extended block cipher that must satisfy the same property as the underlying block-cipher, *i.e.* it must be a (strong) pseudo-random permutation. Many constructions of domain extender for block-ciphers have been defined that satisfy this security notion, for example PEP [6], XCB [14], HCTR [30], HCH [7] and TET [18].

However, it is easy to see that none of those constructions provide the indifferentiability property that enables to get a $2n$ -bit ideal cipher from a n -bit ideal cipher. This is because these constructions were proposed with privacy concerns in mind (mainly for disk encryption purposes) and proven secure only in the classical pseudo-random permutation model. Therefore, these constructions cannot be used when security must hold under the random permutation model (or ideal cipher model). Consider for example the public-key encryption scheme described by Phan and Pointcheval in [27]. The scheme requires a public random permutation with the same size as the RSA modulus, say 1024 bits. In order to replace a 1024-bit random permutation by a construction based on a smaller primitive (for example a 128-bit block cipher), indifferentiability with respect to a 1024-bit random permutation is required. Given a 128-bit block-cipher, none of the previous constructions can provide such property; therefore if one of these constructions is plugged into the Phan and Pointcheval scheme, nothing can be said about the security of the resulting scheme.

In this paper we construct the first domain extender for the ideal cipher; that is we provide a construction of an ideal cipher with $2n$ -bit input from an ideal cipher with n -bit input. Given an ideal cipher with n -bit input/output, one could in principle use the construction in [9] to get a random oracle with n -bit output, and then use the 6-round Feistel in [10] to obtain an ideal cipher with $2n$ -bit input/output, but that would be too inefficient. Moreover the security bound in [10] is rather loose, which implies that the construction only works for large

values of n .¹ In this paper we describe a more efficient construction, based on a 3-round Feistel only, and with a better security bound; we view this as the main result of the paper. More precisely, we show that the 3-round construction in Figure 1 (left) is enough to get a $2n$ -bit random permutation from a n -bit ideal cipher, and that its variant in Figure 1 (right) provides a $2n$ -bit ideal cipher. We also show that 2 rounds are not enough by providing a simple attack. Interestingly, in the so called honest-but-curious model of indistinguishability [12], we show that 2 rounds are sufficient.

Our construction is similar to that of Luby-Rackoff [23]. However we stress that the “indifferentiable construction” security notion is very different from the classical indistinguishability notion. The well known Luby-Rackoff result that 4 rounds are enough to obtain a strong pseudo-random permutation from pseudo-random functions [23], is proven under the classical indistinguishability notion. Under this notion, the adversary has only access to the input/output of the Luby-Rackoff construction, and tries to distinguish it from a random permutation; in particular it does not have access to the input/output of the inner pseudo-random functions. On the contrary, in our setting, the distinguisher can make oracle calls to the inner block-ciphers E_i 's (see Fig. 1); the indistinguishability notion enables to accommodate these additional oracle calls in a coherent definition.

The indistinguishability security notion still requires a (small) ideal component. We stress that it is unknown how to instantiate such ideal component (be it a random oracle or an ideal cipher, as opposed to a PRF or a PRP) and that the security guarantee does not hold anymore once that component is instantiated. Moreover the recent related-key attacks on AES [23] show that AES-192 and AES-256 do not behave as ideal ciphers; as of 2009 it is unclear if we have a candidate block-cipher with key-size larger than block-size that behaves like an ideal cipher.

Finally, we also analyze our construction in the standard model. In this case, we use a *tweakable* block-cipher as the underlying primitive. Tweakable block-ciphers were introduced by Liskov, Rivest and Wagner in [22] and provide an additional input - the tweak - that enables to get a *family* of independent block-ciphers; efficient constructions of tweakable block-ciphers were described in [22], given ordinary block-ciphers. In this paper we show that our construction with 2 rounds enables to get a $2n$ -bit tweakable block-cipher from a n -bit tweakable block-cipher. Moreover we show that with 3 rounds we achieve a security guarantee beyond the birthday paradox.

1.1 Related Work

At FSE 2009, Minematsu [25] provided two constructions of a $2n$ -bit block-cipher from an n -bit tweakable block-cipher :

¹ The security bound in [10] for the 6-round Feistel random oracle based construction is $q^{16}/2^n$, where q is the number of distinguisher's queries. This implies that for $q = 2^{64}$, one must take at least $n = 1024$, which corresponds to a 2048-bit permutation.

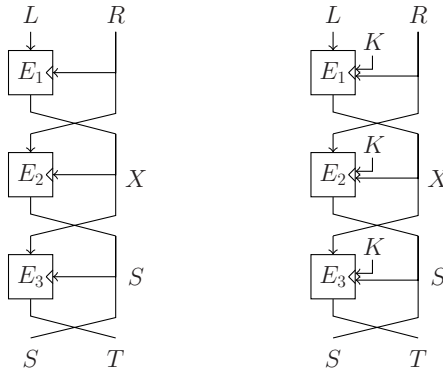


Fig. 1. Construction of a $2n$ -bit permutation given a n -bit ideal cipher with n -bit key (left). Construction of a $2n$ -bit ideal cipher with k -bit key, given a n -bit ideal cipher with $(n + k)$ -bit key (right).

1. A 3-round Feistel construction with universal hashing in the 1st round and tweakable block ciphers in the 2nd and the 3rd rounds. This construction is a secure pseudo-random permutation beyond the birthday bound.
2. A 4-round Feistel with universal hashing in the 1st and the 4th rounds and tweakable block ciphers in the 2nd and the 3rd rounds. This construction is a secure strong pseudo-random permutation beyond the birthday bound.

On the other hand, our construction in this paper is a 3-round Feistel, with tweakable block ciphers in every round, and it gives a secure (tweakable) strong pseudo-random permutation beyond the birthday bound. Therefore, the construction in [25] is more efficient as only 2 calls are required to the underlying tweakable block-cipher, instead of 3 calls in our construction (this is assuming very fast universal hashing, e.g. [21]). However, we stress that the constructions in [25] are secure only in the symmetric-key setting; it is easy to see that none of the two constructions from [25] can achieve the indistinguishability property (the attack is similar to the attack against 2-round Feistel described in Section 3).

2 Definitions

We first recall the notion of indistinguishability of random systems, introduced by Maurer *et al.* in [24]. This is an extension of the classical notion of indistinguishability, where one or more oracles are publicly available, such as random oracles or ideal ciphers.

As in [24], we define an *ideal primitive* as an algorithmic entity which receives inputs from one of the parties and delivers its output immediately to the querying party. In this paper, we consider ideal primitives such as random oracle, random permutation and ideal cipher. A *random oracle* [1] is an ideal primitive which provides a random output for each new query; identical input queries are given

the same answer. A *random permutation* is an ideal primitive that provides oracle access to a random permutation $P : \{0,1\}^n \rightarrow \{0,1\}^n$ and to P^{-1} . An *ideal cipher* is a generalization of a random permutation that models a random block cipher $E : \{0,1\}^k \times \{0,1\}^n \rightarrow \{0,1\}^n$. Each key $k \in \{0,1\}^k$ defines an independent random permutation $E_k = E(k, \cdot)$ on $\{0,1\}^n$. The ideal primitive also provides oracle access to E and E^{-1} ; that is, on query $(0, k, m)$, the primitive answers $c = E_k(m)$, and on query $(1, k, c)$, the primitive answers m such that $c = E_k(m)$. We stress that in the ideal cipher model, the adversary has oracle access to a publicly available ideal cipher and must send both the key and the plaintext in order to obtain the ciphertext; this is different from the standard model in which the key is privately generated by the system.

The notion of indistinguishability [24] enables to show that an ideal primitive \mathcal{P} (for example, a random permutation) can be replaced by a construction C that is based on some other ideal primitive E ; for example, C can be the Feistel construction illustrated in Fig. 1 (left).

Definition 1 ([24]). A Turing machine C with oracle access to an ideal primitive E is said to be $(t_D, t_S, q, \varepsilon)$ -indistinguishable from an ideal primitive \mathcal{P} if there exists a simulator S with oracle access to \mathcal{P} and running in time at most t_S , such that for any distinguisher D running in time at most t_D and making at most q queries, it holds that:

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

C^E is simply said to be indistinguishable from \mathcal{P} if ε is a negligible function of the security parameter n , for polynomially bounded q , t_D and t_S .

The previous definition is illustrated in Figure 2, where C is our 3-round construction of Figure 1 (left), E is an ideal cipher, \mathcal{P} is a random permutation and S is the simulator. In this paper, for a 3-round construction, we denote these ideal ciphers by E_1, E_2, E_3 (see Fig. 1). Equivalently, one can consider a single ideal cipher E and encode in the first 2 key bits which round ideal cipher E_1, E_2 , or E_3 is actually called. The distinguisher has either access to the system




Fig. 2. The indistinguishability notion

formed by the construction C and the ideal cipher E , or to the system formed by the random permutation P and a simulator \mathcal{S} . In the first system (left), the construction C computes its output by making calls to the ideal cipher E (equivalently the 3 ideal ciphers E_1, E_2 and E_3); the distinguisher can also make calls to E directly. In the second system (right), the distinguisher can either query the random permutation P , or the simulator that can make queries to P . If the distinguisher first makes a call to the construction C , and then makes the corresponding calls to ideal cipher E , he will get the same answer. This must remain true when the distinguisher interacts with permutation P and simulator \mathcal{S} . The role of simulator \mathcal{S} is then to simulate the ideal ciphers E_i 's so that 1) the output of \mathcal{S} should be indistinguishable from that of ideal ciphers E_i 's and 2) the output of \mathcal{S} should look “consistent” with what the distinguisher can obtain independently from P . We note that in this model the simulator does not see the distinguisher’s queries to P ; however, it can call P directly when needed for the simulation.

It is shown in [24] that the indifferentiability notion is the “right” notion for substituting one ideal primitive with a construction based on another ideal primitive. That is, if C^E is indifferentiable from an ideal primitive \mathcal{P} , then C^E can replace \mathcal{P} in any cryptosystem, and the resulting cryptosystem is at least as secure in the E model as in the \mathcal{P} model; see [24] or [9] for a proof.

3 An Attack against 2 Rounds

In this section we show that 2 rounds are not enough when the inner ideal ciphers are publicly accessible, that is we exhibit a property for 2 rounds that does not exist for a random permutation.

Formally, the 2 round construction is defined as follows (see Fig. 3). Let $E_1 : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a block cipher, where $c = E_1(K, m)$ is the n -bit ciphertext corresponding to n -bit key K and n -bit input message m ; let E_2 be defined similarly. We define the permutation $\Psi_2 : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ as:

$$\Psi_2(L, R) := (E_1(R, L), E_2(E_1(R, L), R))$$




Fig. 3. The 2-round Feistel construction $\Psi_2(L, R)$

It is easy to see that this defines an invertible permutation over $\{0, 1\}^{2n}$. Namely, given a ciphertext (S, T) the value R is recovered by “decrypting” T with block-cipher E_2 and key S , and the value L is recovered by “decrypting” S with block-cipher E_1 and key R .

The attack against permutation Ψ_2 is straightforward; it is based on the fact that the attacker can arbitrarily choose both R and S . More precisely, the attacker selects $R = 0^n$ and $S = 0^n$ and queries $L = E_1^{-1}(R, S)$ and $T = E_2(S, R)$. This gives $\Psi_2(L, R) = (S, T)$ as required. However, it is easy to see that with a random permutation P and a polynomially bounded number of queries, it is impossible to find L, R, S, T such that $P(L\|R) = S\|T$ with both $R = 0^n$ and $S = 0^n$, except with negligible probability. Therefore, the 2-round construction cannot replace a random permutation.

Theorem 1. *The 2-round Feistel construction Ψ_2 is not indifferntiable from a random permutation.*

In the full version of the paper [8] we also analyse existing constructions of domain extender for block ciphers and show that they are not indifferntiable from an ideal cipher; more precisely, we show that the CMC [16] and EME [17] constructions are not indifferntiable from an ideal cipher. We stress that our observations do not imply anything concerning their security in the standard pseudo-random permutation model.

4 Indifferntiability of 3-Round Feistel Construction

We now prove our first main result: the 3-round Feistel construction is indifferntiable from a random permutation. To get an ideal cipher, it suffices to prepend a key K to the 3 ideal ciphers E_1, E_2 and E_3 ; one then gets a family of independent random permutation, parametrised by K , i.e. an ideal cipher (see Fig. 1 for an illustration).

Formally, the 3 round permutation $\Psi_3 : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ is defined as follows, given block ciphers E_1, E_2 and E_3 with n -bit key (first variable) and n -bit input/output (second variable):

$$\begin{aligned} X &= E_1(R, L) \\ S &= E_2(X, R) \\ T &= E_3(S, X) \\ \Psi_3(L, R) &:= (S, T) \end{aligned}$$

The 3 round block cipher $\Psi'_3 : \{0, 1\}^k \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ is defined as follows, given block ciphers E_1, E_2 and E_3 with $(k + n)$ -bit key and n -bit input/output:

$$\begin{aligned} X &= E_1(K\|R, L) \\ S &= E_2(K\|X, R) \\ T &= E_3(K\|S, X) \\ \Psi'_3(K, (L, R)) &:= (S, T) \end{aligned}$$

Theorem 2. *The 3-round Feistel construction Ψ_3 is $(t_D, t_S, q, \varepsilon)$ -indifferentiable from a random permutation, with $t_S = \mathcal{O}(qn)$ and $\varepsilon = 5q^2/2^n$. The 3-round block-cipher construction Ψ'_3 is $(t_D, t_S, q, \varepsilon)$ -indifferentiable from an ideal cipher, with $t_S = \mathcal{O}(qn)$ and $\varepsilon = 5q^2/2^n$.*

Proof. We only consider the 3-round permutation Ψ_3 ; the extension to block-cipher Ψ'_3 is straightforward. We must construct a simulator \mathcal{S} such that the two systems formed by (Ψ_3, E) and (P, \mathcal{S}) are indistinguishable (see Fig. 2).

Our simulator maintains an history of already answered queries for E_1 , E_2 and E_3 . Formally, when the simulator answers X for a $E_1(R, L)$ query, it stores $(1, R, L, X)$ in history; the simulator proceeds similarly for E_2 and E_3 queries. We write that the simulator “simulates” $E_1(R, L) \leftarrow X$ when it first generates a random $X \in \{0, 1\}^n \setminus \mathcal{B}$, where \mathcal{B} is the set of already defined values for $E_1(R, \cdot)$, and then stores $(1, R, L, X)$ in history, meaning that $E_1(R, L) = X$; we use similar notations for E_2 and E_3 . The distinguisher’s queries are answered as follows by the simulator:

- | | |
|--|--|
| $E_1(R, L)$ query: <ol style="list-style-type: none"> 1. Simulate $E_1(R, L) \leftarrow X$ 2. $(S, T) \leftarrow \text{Adapt}(L, R, X)$ 3. Return X | $E_1^{-1}(R, X)$ query <ol style="list-style-type: none"> 1. Simulate $E_1^{-1}(R, X) \leftarrow L$ 2. $(S, T) \leftarrow \text{Adapt}(L, R, X)$ 3. Return L |
| $E_2(X, R)$ query: <ol style="list-style-type: none"> 1. Simulate $E_1^{-1}(R, X) \leftarrow L$ 2. $(S, T) \leftarrow \text{Adapt}(L, R, X)$ 3. Return S | $\text{Adapt}(L, R, X)$: <ol style="list-style-type: none"> 1. $S \ T \leftarrow P(L \ R)$ 2. Store $E_2(X, R) = S$ in history 3. Store $E_3(S, X) = T$ in history. 4. Return (S, T). |

The procedure for answering the other queries is essentially symmetric; we provide it for completeness:

- | | |
|--|--|
| $E_3^{-1}(S, T)$ query: <ol style="list-style-type: none"> 1. Simulate $E_3^{-1}(S, T) \leftarrow X$ 2. $(L, R) \leftarrow \text{Adapt}^{-1}(S, T, X)$ 3. Return X | $E_3(S, X)$ query <ol style="list-style-type: none"> 1. Simulate $E_3(S, X) \leftarrow T$ 2. $(L, R) \leftarrow \text{Adapt}^{-1}(S, T, X)$ 3. Return T |
| $E_2^{-1}(X, S)$ query: <ol style="list-style-type: none"> 1. Simulate $E_3(S, X) \leftarrow T$ 2. $(L, R) \leftarrow \text{Adapt}^{-1}(S, T, X)$ 3. Return R | $\text{Adapt}^{-1}(S, T, X)$: <ol style="list-style-type: none"> 1. $L \ R \leftarrow P^{-1}(S \ T)$ 2. Store $E_2(X, R) = S$ in history. 3. Store $E_1(R, L) = X$ in history. 4. Return (L, R) |

Finally, the simulator aborts if for some E_i and some key K , it has not defined a permutation for $E_i(K, \cdot)$; that is the simulator aborts if it has defined $E_i(K, X) = E_i(K, Y)$ for some $X \neq Y$ or it has defined $E_i^{-1}(K, X) = E_i^{-1}(K, Y)$ for some $X \neq Y$. This completes the description of the simulator.

As a consistency check, it is easy to see that if the distinguisher makes a single query for $P(L \| R)$ and then queries the simulator for $X \leftarrow E_1(R, L)$, $S \leftarrow$

$E_2(X, R)$ and $T \leftarrow E_3(S, X)$, then the distinguisher obtains $S\|T = P(L\|R)$ as required.

We now proceed to prove that the systems (Ψ_3, E) and (P, S) are indistinguishable. We consider a distinguisher \mathcal{D} making at most q queries to the system (Ψ_3, E) or (P, S) and outputting a bit γ . We define a sequence $\text{Game}_0, \text{Game}_1, \dots$ of modified distinguisher games. In the first game the distinguisher interacts with the system (Ψ_3, E) . We incrementally modify the system so that in the last game the distinguisher interacts with the system (P, S) , where S is the previously defined simulator. We denote by S_i the event that in game i the distinguisher outputs $\gamma = 1$.

- **Game₀**: the distinguisher interacts with Ψ_3 and the ideal ciphers E_i .
- **Game₁**: we modify the way E_i queries are answered, without actually changing the value of the answer. We also maintain an history of already answered queries for E_1, E_2 and E_3 . We proceed as follows:

| | |
|--|--|
| $E_1(R, L)$ query: <ol style="list-style-type: none"> 1. Let $X \leftarrow E_1(R, L)$ 2. $(S, T) \leftarrow \text{Adapt}'(L, R, X)$ 3. Return X | $E_1^{-1}(R, X)$ query <ol style="list-style-type: none"> 1. Let $L \leftarrow E_1^{-1}(R, X)$ 2. $(S, T) \leftarrow \text{Adapt}'(L, R, X)$ 3. Return L |
| $E_2(X, R)$ query: <ol style="list-style-type: none"> 1. Let $L \leftarrow E_1^{-1}(R, X)$ 2. $(S, T) \leftarrow \text{Adapt}'(L, R, X)$ 3. Return S | $\text{Adapt}'(L, R, X)$: <ol style="list-style-type: none"> 1. $S\ T \leftarrow \Psi_3(L\ R)$ 2. Store $E_2(X, R) = S$ in history. 3. Store $E_3(S, X) = T$ in history. 4. Return (S, T) |

The queries to $E_2^{-1}(X, S), E_3(S, X)$ and $E_3^{-1}(S, T)$ are answered symmetrically.

For example, when given a query to $E_1(R, L)$, we first query ideal cipher E_1 for $X \leftarrow E_1(R, L)$; then instead of X being returned immediately as in **Game₀**, we let $S\|T = \Psi_3(L\|R)$, which gives $S = E_2(X, R)$ and $E_3(S, X) = T$; we then store $(2, X, R, S)$ and $(3, S, X, T)$ in history. Therefore, the value that get stored in history is exactly the same as the value from ideal ciphers E_2 and E_3 ; the only difference is that this value was obtained indirectly by querying Ψ_3 instead of directly by querying E_2 and E_3 . It is easy to see that this holds for any query made by the distinguisher, who receives exactly the same answers in **Game₀** and **Game₁**; this implies:

$$\Pr[S_1] = \Pr[S_0]$$

As illustrated in Fig. 4, we have actually constructed a simple simulator \mathcal{S}' that makes queries to a subsystem \mathcal{T} that comprises the construction Ψ_3 and the ideal ciphers E_1, E_2 and E_3 . The difference between \mathcal{S}' in **Game₁** and the main simulator \mathcal{S} defined previously is that 1) \mathcal{S}' calls ideal cipher $E_1(R, L)$ instead of simulating it and 2) \mathcal{S}' makes calls to $\Psi_3(L\|R)$ instead of $P(L\|R)$.

- **Game₂**: we modify the way the permutation queries are answered. Instead of using Ψ_3 as in system \mathcal{T} , we use the random permutation P in the new system \mathcal{T}' (see Fig. 4).




Fig. 4. Sequence of games for proving indifferentiability

We must show that the distinguisher’s view has statistically close distribution in **Game**₁ and **Game**₂. For this, we consider the subsystem \mathcal{T} with the 3-round Feistel Ψ_3 and the ideal ciphers E_i ’s in **Game**₁, and the subsystem \mathcal{T}' with the random permutation P and ideal ciphers E_i ’s in **Game**₂. We show that the output of systems \mathcal{T} and \mathcal{T}' is statistically close; this in turn shows that the distinguisher’s view has statistically close distribution in **Game**₁ and **Game**₂. Note that the indistinguishability of \mathcal{T} and \mathcal{T}' only holds for the particular set of queries made by the distinguisher and the simulator; it could not hold for any possible set of queries.

In the following, we assume that the distinguisher eventually makes a sequence of E_i queries corresponding to all previous Ψ_3 queries that he has made. More precisely, if the distinguisher has made a $\Psi_3(L, R)$ query, then eventually the distinguisher makes the sequence of queries $X \leftarrow E_1(R, L)$, $S \leftarrow E_2(X, R)$ and $T \leftarrow E_3(S, X)$ to the simulator; the same holds for $\Psi_3^{-1}(S, T)$ queries. This is without loss of generality, because from any distinguisher \mathcal{D} we can build a distinguisher \mathcal{D}' with the same output that satisfies this property.

The outputs to E_i queries provided by subsystem \mathcal{T} in **Game**₁ and by subsystem \mathcal{T}' in **Game**₂ are the same, since in both cases these queries are answered by ideal ciphers E_i . Therefore, we must show that the output to P/P^{-1} queries provided by \mathcal{T} and \mathcal{T}' have statistically close distribution, when the outputs to E_i queries provided by \mathcal{T} or \mathcal{T}' are fixed.

We consider a forward permutation query $L\|R$ made by either the distinguisher or the simulator S' . If this $L\|R$ query is made by the distinguisher, since we have assumed that the distinguisher eventually makes the E_i queries corresponding to all his permutation queries, this $L\|R$ query will also be made by the simulator S' , by definition of S' . Therefore we can consider $L\|R$ queries made by the simulator S' only.

We first consider the answer to $S\|T = \Psi_3(L\|R)$ in **Game**₁. In this case the answer $S\|T$ is computed as follows:

$$\begin{aligned} X &= E_1(R, L) \\ S &= E_2(X, R) \\ T &= E_3(S, X) \end{aligned}$$

By definition of the simulator \mathcal{S}' , when the simulator \mathcal{S}' makes a query for $\Psi_3(L\|R)$, it must have made an ideal cipher query to $E_1(R, L)$ before, or an ideal cipher query to $E_1^{-1}(R, X)$ before, with $L = E_1^{-1}(R, X)$.

If the simulator \mathcal{S}' has made an ideal cipher query for $E_1(R, L)$ to subsystem \mathcal{T} , then from the definition of the simulator a call to $\text{Adapt}'(L, R, X)$ has occurred, where $X = E_1(R, L)$; in this Adapt' call the values $E_2(X, R)$ and $E_3(S, T)$ are defined by the simulator; therefore the simulator does not make these queries to sub-system \mathcal{T} . This implies that the values of $E_2(X, R)$ and $E_3(S, X)$ are not included in the subsystem \mathcal{T} output; therefore these values are not fixed in the probability distribution that we consider; only the value $X = E_1(R, L)$ is fixed.

Moreover, for fixed X, R the distribution of $S = E_2(X, R)$ is uniform in $\{0, 1\}^n \setminus \mathcal{B}$, where \mathcal{B} is the set of already defines values for $E_2(X, \cdot)$. Since there are at most q queries, the statistical distance between the distribution of $E_2(X, R)$ and the uniform distribution in $\{0, 1\}^n$ is at most $2q/2^n$; the same holds for the distribution of $T = E_3(S, X)$. Therefore, we obtain that for a fixed X , the distribution of (S, T) is statistically close to the uniform distribution in $\{0, 1\}^{2n}$, with statistical distance at most $4q/2^n$.

If the simulator has made an ideal cipher query for $E_1^{-1}(R, X)$, then the same analysis applies and we obtain that for a fixed $L = E_1^{-1}(R, X)$ the distribution of (S, T) is statistically close to the uniform distribution in $\{0, 1\}^{2n}$, with statistical distance at most $4q/2^n$. Therefore we obtain that in Game_1 the statistical distance of $S\|T = \Psi_3(L\|R)$ with the uniform distribution is always at most $4q/2^n$.

In Game_2 , the output to permutation query $L\|R$ is $S\|T = P(L\|R)$; since there are at most q queries to P/P^{-1} , the statistical distance between $P(L\|R)$ and the uniform distribution in $\{0, 1\}^{2n}$ is at most $2q/2^{2n}$.

Therefore the statistical distance between $\Psi_3(L, R)$ in Game_1 and $P(L\|R)$ in Game_2 is at most $4q/2^n + 2q/2^{2n} \leq 5q/2^n$. The same argument applies to inverse permutation queries. This holds for a single permutation query; since there are at most q such queries, we obtain that the statistical distance between outputs of systems \mathcal{T} and \mathcal{T}' to permutation queries and E_i queries, is at most $5q^2/2^n$; this implies:

$$|\Pr[S_2] - \Pr[S_1]| \leq \frac{5q^2}{2^n}$$

- **Game₃**: eventually the distinguisher interacts with system (P, S) . The only difference between the simulator \mathcal{S}' in Game_2 and the simulator \mathcal{S} in Game_3 is that instead of querying ideal ciphers E_i in Game_2 , these ideal ciphers are simply simulated in Game_3 , while the answer to permutation queries are exactly the same. Therefore, the distinguisher's view has the same distribution in Game_2 and Game_3 , which gives:

$$\Pr[S_2] = \Pr[S_3]$$

and finally:

$$|\Pr[S_3] - \Pr[S_0]| \leq \frac{5q^2}{2^n}$$

which terminates the proof of Theorem \square

\square

We note that the security bound in $q^2/2^n$ for our 3-round ideal cipher based construction is much better than the security bound in $q^{16}/2^n$ obtained for the 6-round Feistel construction in [10] (based on random oracles).

4.1 Practical Considerations

EXTENDING THE KEY. So far, we showed how to construct an ideal cipher Ψ_3 with $2n$ -bit message and k -bit key from three ideal ciphers E_1, E_2, E_3 on n -bit message and $(n + k)$ -bit key. As already mentioned, we can actually implement E_1, E_2, E_3 from a single n -bit ideal cipher E whose key length is $n + k + 2$.

However, if only a block-cipher with n -bit key and n -bit message is available (for example AES-128), we need a procedure to extend the key size. To handle such cases, we notice that it suffices to first hash the key using a random oracle, and the resulting block cipher remains indifferentiable from an ideal cipher.

Lemma 1. *Assume $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is an ideal cipher and $H : \{0, 1\}^t \rightarrow \{0, 1\}^k$ is a random oracle. Define $E' : \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ by $E'(K', X) = E(H(K'), X)$, $E'^{-1}(K', Y) = E^{-1}(H(K'), Y)$. Then E' is $(t_D, t_S, q, \varepsilon)$ -indifferentiable from an ideal cipher, where $t_S = \mathcal{O}(q(n + t))$ and $\varepsilon = \mathcal{O}(q^2/2^k)$.*

Proof. See the full version of the paper [8].

Using this observation, given a single ideal cipher E on n -bit messages and k -bit key and a random oracle H with output size k bits, we can first build an ideal cipher E' with n -bit message and $(n + k' + 2)$ -bit key, and then from Theorem 2 we can obtain an ideal cipher Ψ_3 on $2n$ -bit messages and k' -bit key. It remains to remove the assumption of having random oracle H ; this can easily be accomplished by sacrificing 1 key bit from E , and then using one of the two resulting (independent) ideal ciphers to efficiently implement H using any of the methods from [9].

GOING BEYOND DOUBLE? Another natural question is to extend the domain of the ideal cipher beyond doubling it. One way to accomplish this task is to apply our 3-round construction recursively, each time doubling the domain. However, in this case it is not hard to see that, to extend the domain by a factor of t , the original block cipher E will have to be used $\mathcal{O}(t^{\log_2 3})$ times [2]. This makes the resulting constructions somewhat impractical for large t .

In contrast, assume that we use the 2-step construction: first build a length-preserving random oracle H on $nt/2$ bits (using [9]), and then use the 6-round Feistel construction [10] to get a nt -bit permutation. To construct a random oracle from $nt/2$ -bit to $nt/2$ -bit, only $\mathcal{O}(t)$ calls to the n -bit ideal cipher are

² In essence, this is because we call E three times for each doubling. Actually, this is not counting the calls to the independent variable length random oracle H to hash down the key, as above. However, because the constructions of such an H in [9] are so efficient, it is not hard to see that, even when implementing H using E itself, the dominant term remains $\mathcal{O}(t^{\log_2 3})$ (although the constant is slightly worse).

required (first hash from $nt/2$ -bit to n -bit using [9], then expand back to $nt/2$ -bits using counter mode). Therefore the 2-step construction requires only $\mathcal{O}(t)$ calls to E , instead of $\mathcal{O}(t^{\log_2 3})$ when iterating our construction. This implies that for large t , the 2-step construction is more efficient.

To give a practical example, let us consider the applications of [15,27], where one needs to apply a random permutation to the domain of an RSA modulus. We take the length of modulus N to be 1024 bits and the underlying block-cipher E to be $n = 128$ with 128-bit key (as in AES-128). One can see that to obtain a 1024-bit permutation from E , only 48 calls to E are required for the 2-step construction, instead of 243 when iterating our construction. However for 1024-bit, the exact security of the 2-step construction is dominated by the term $\mathcal{O}(q^{16}/2^{512})$ from [10], which requires $q \ll 2^{32}$, whereas the exact security of the recursive construction is $\mathcal{O}(q^2/2^{128})$, which requires $q \ll 2^{64}$. Therefore, for a 1024-bit permutation our recursive construction still provides a better security bound; however, for any size larger than 2048 bits, the two constructions have the same $q \ll 2^{64}$ bound [3].

To summarize, our construction is more efficient than the 2-step construction when doubling only once ($t = 2$). However for a large expansion factor t the 2-step construction is more efficient than the recursive method.

4.2 Indifferentiability for 2 Rounds in the Honest-But-Curious Model

In the full version of the paper we also consider the *honest-but-curious* model of indifferentiability introduced by Dodis and Puniya [12], which is a variant of the general indifferentiability model. We show that in the honest-but-curious model, 2 rounds as depicted in Fig 3 are actually sufficient to get indifferentiability.

5 Domain Extension of Tweakable Block Cipher

In this section, we also analyse our construction in the standard model, and we use a *tweakable* block-cipher as the underlying primitive. The main result of this section is that a 3-round Feistel enables to get a security guarantee beyond the birthday paradox.

Tweakable block-ciphers were introduced by Liskov, Rivest and Wagner in [22] and provide an additional input - the tweak - that enables to get a *family* of independent block-ciphers. Efficient constructions of tweakable block-ciphers were described in [22], given ordinary block-ciphers.

Definition 2. A *tweakable block-cipher* is an efficiently computable function $\tilde{E} : \{0, 1\}^k \times \{0, 1\}^\omega \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ that takes as input a key $K \in \{0, 1\}^k$, a tweak $W \in \{0, 1\}^\omega$ and a message $m \in \{0, 1\}^n$ and returns a ciphertext $c \in \{0, 1\}^n$. For every $K \in \{0, 1\}^k$ and $W \in \{0, 1\}^\omega$, the function $\tilde{E}(K, W, \cdot)$ is a permutation over $\{0, 1\}^n$.

³ The length-preserving random oracle used in the 6-round Feistel has the birthday bound of $q^2/2^{128}$.

The security notion for a tweakable block-cipher is a straightforward extension of the corresponding notion for block-ciphers. A classical block-cipher E is a strong pseudo-random permutation if no adversary can distinguish $E(K, \cdot)$ from a random permutation, where \mathcal{A} can make calls to both E and E^{-1} , and $K \leftarrow \{0, 1\}^k$. For tweakable block-ciphers, the adversary can additionally choose the tweak, and $E(K, \cdot, \cdot)$ should be indistinguishable from a family of random permutations, parametrised by $W \in \{0, 1\}^\omega$:

Definition 3. A tweakable block-cipher is said to be (t, q, ε) -secure if for any adversary \mathcal{A} running in time at most t and making at most q queries, the adversary’s advantage in distinguishing $\tilde{E}(K, \cdot, \cdot)$ with $K \leftarrow \{0, 1\}^k$ from a family of independent random permutation $\tilde{\Pi}(\cdot, \cdot)$ is at most ε , where \mathcal{A} can make calls to both \tilde{E} and \tilde{E}^{-1} .

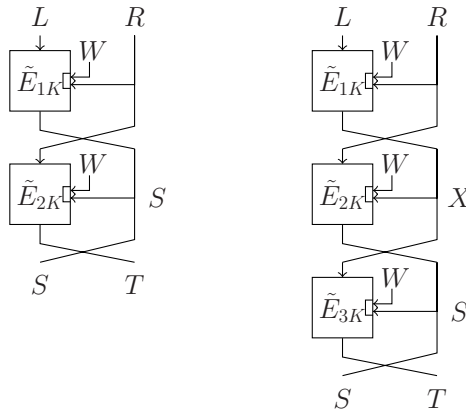


Fig. 5. The tweakable block ciphers $\tilde{\Psi}_2$ (left) and $\tilde{\Psi}_3$ (right), with key K and tweak W

We first show that 2 rounds are enough to get a $2n$ -bit tweakable block-cipher from a n -bit tweakable block-cipher (see Fig. 5, left). Formally, our 2-round domain extender for tweakable block-cipher works as follows. Let E_1 and E_2 be two tweakable block-ciphers with the same signature:

$$\tilde{E}_i : \{0, 1\}^k \times \{0, 1\}^\omega \times \{0, 1\}^n \rightarrow \{0, 1\}^n$$

The tweakable block cipher $\tilde{\Psi}_2 : \{0, 1\}^k \times \{0, 1\}^{\omega-n} \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ is then defined as follows; the difference with Fig. 3 is that the R and S inputs go to the tweak (concatenated with the main tweak W) instead of the key.

$$\begin{aligned} S &= E_1(K, W \| R, L) \\ T &= E_2(K, W \| S, R) \\ \tilde{\Psi}_2(K, W, (L, R)) &= (S, T) \end{aligned}$$

Theorem 3. *The tweakable block-cipher $\tilde{\Psi}_2$ is a (t', q, ε') -secure tweakable block-cipher, if \tilde{E}_1 and \tilde{E}_2 are both (t, q, ε) -secure tweakable block-ciphers, where $\varepsilon' = 2 \cdot \varepsilon + q^2/2^n + q^2/2^{2n}$ and $t' = t - \mathcal{O}(qn)$.*

Proof. See the full version of the paper [8].

Now we consider the 3 round tweakable block cipher $\tilde{\Psi}_3$, defined in a similar manner as $\tilde{\Psi}_2$ (see Fig. 5 for an illustration). The 3-round construction enables to go beyond the birthday security bound. Namely instead of having a bound in $q^2/2^n$ as in the 2-round construction, the bound for the 3-round construction is now $q^2/2^{2n}$, which shows that the construction remains secure until $q < 2^n$ instead of $q < 2^{n/2}$.

Theorem 4. *The tweakable block-cipher $\tilde{\Psi}_3$ is a (t', q, ε') -secure tweakable block-cipher, if \tilde{E}_1, \tilde{E}_2 and \tilde{E}_3 are all (t, q, ε) -secure tweakable block-ciphers, where $\varepsilon' = 3 \cdot \varepsilon + q^2/2^{2n}$ and $t' = t - \mathcal{O}(qn)$.*

Proof. See the full version of the paper [8].

One drawback of our construction is that it shrinks the tweak size from ω bits to $\omega - n$ bits. We show a simple construction that extends the tweak size, using a keyed universal hash function; this construction can be of independent interest.

Definition 4. *A family \mathcal{H} of functions with signature $\{0, 1\}^{\omega'} \rightarrow \{0, 1\}^\omega$ is said to be ε -almost universal if $\Pr_h[h(x) = h(y)] \leq \varepsilon$ for all $x \neq y$, where the probability is taken over h chosen uniformly at random from \mathcal{H} .*

Let \tilde{E} be a tweakable block-cipher with tweak in $\{0, 1\}^\omega$. Given a family \mathcal{H} of hash functions h with signature $\{0, 1\}^{\omega'} \rightarrow \{0, 1\}^\omega$ and $\omega' > \omega$, our tweakable block-cipher \tilde{E} with extended tweak length ω' is defined as:

$$\tilde{E}'((K, h), W', m) = \tilde{E}(K, h(W'), m)$$

Theorem 5. *The tweakable block cipher \tilde{E}' is a (q, t', ε') -secure tweakable block cipher if \tilde{E} is a (q, t, ε_1) -secure tweakable block cipher and the hash function family \mathcal{H} is ε_2 -almost universal, with $\varepsilon' = \varepsilon_1 + q^2 \cdot \varepsilon_2$ and $t' = t - \mathcal{O}(q)$.*

Proof. See the full version of the paper [8].

We note that many efficient constructions of universal hash function families are known, with $\varepsilon_2 \simeq 2^{-\omega}$. Therefore the new tweakable block-cipher can have the same level of security as the original one, up to the birthday bound for the tweak, i.e. for $q \leq 2^{\omega/2}$.

6 Conclusion

We have described the first domain extender for ideal ciphers, i.e. we have showed a construction that is indifferentiable from a $2n$ -bit ideal cipher, given a n -bit

ideal cipher. Our construction is based on a 3-round Feistel, and is more efficient and more secure than first building a n -bit random oracle from a n -bit ideal cipher (as in [9]) and then a $2n$ -bit ideal cipher from a n -bit random oracle (as in [10]). We have also shown that in the standard model, our construction with 2 rounds enables to get a $2n$ -bit tweakable block-cipher from a n -bit tweakable block-cipher and that with 3 rounds we get a security guarantee beyond the birthday paradox.

References

1. Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security, pp. 62–73 (1993)
2. Biryukov, A., Khovratovich, D., Nikolic, I.: Distinguisher and Related-Key Attack on the Full AES-256. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 231–249. Springer, Heidelberg (2009)
3. Biryukov, A., Khovratovich, D.: Related-key Cryptanalysis of the Full AES-192 and AES-256. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 1–18. Springer, Heidelberg (2009)
4. Black, J.: The Ideal-Cipher Model, Revisited: An Uninstantiable Blockcipher-Based Hash Function. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 328–340. Springer, Heidelberg (2006)
5. Black, J., Rogaway, P., Shrimpton, T.: Black-Box Analysis of the Block Cipher-Based Hash-Function Constructions from PGV. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, p. 320. Springer, Heidelberg (2002)
6. Chakraborty, D., Sarkar, P.: A new mode of encryption providing a tweakable strong pseudo-random permutation. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 293–309. Springer, Heidelberg (2006)
7. Chakraborty, D., Sarkar, P.: HCH: A new tweakable enciphering scheme using the hash-encrypt-hash approach. In: Barua, R., Lange, T. (eds.) INDOCRYPT 2006. LNCS, vol. 4329, pp. 287–302. Springer, Heidelberg (2006)
8. Coron, J.S., Dodis, Y., Mandal, A., Seurin, Y.: A Domain Extender for the Ideal Cipher. Full version of this paper. Cryptology ePrint Archive, Report 2009/356, <http://eprint.iacr.org/>
9. Coron, J.S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård Revisited: How to Construct a Hash Function. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer, Heidelberg (2005)
10. Coron, J.S., Patarin, J., Seurin, Y.: The Random Oracle Model and the Ideal Cipher Model are Equivalent. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 1–20. Springer, Heidelberg (2008), Full version available at Cryptology ePrint Archive, Report 2008/246, <http://eprint.iacr.org/>
11. Desai, A.: The security of all-or-nothing encryption: Protecting against exhaustive key search. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, p. 359. Springer, Heidelberg (2000)
12. Dodis, Y., Puniya, P.: On the Relation Between the Ideal Cipher and the Random Oracle Models. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 184–206. Springer, Heidelberg (2006)

13. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. In: Matsumoto, T., Imai, H., Rivest, R.L. (eds.) ASIACRYPT 1991. LNCS, vol. 739, pp. 210–224. Springer, Heidelberg (1993)
14. Fluhrer, S.R., McGrew, D.A.: The extended codebook (XCB) mode of operation. Technical Report 2004/078, IACR eprint archive (2004)
15. Granboulan, L.: Short signature in the random oracle model. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 364–378. Springer, Heidelberg (2002)
16. Halevi, S., Rogaway, P.: A tweakable enciphering mode. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 482–499. Springer, Heidelberg (2003)
17. Halevi, S., Rogaway, P.: A parallelizable enciphering mode. In: Okamoto, T. (ed.) CT-RSA 2004. LNCS, vol. 2964, pp. 292–304. Springer, Heidelberg (2004)
18. Halevi, S.: Invertible Universal hashing and the TET Encryption Mode. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 412–429. Springer, Heidelberg (2007)
19. Jonsson, J.: An OAEP variant with a tight security proof, <http://eprint.iacr.org/2002/034/>
20. Kilian, J., Rogaway, P.: How to protect DES against exhaustive key search (An analysis of DESX). *Journal of Cryptology* 14(1), 17–35 (2001)
21. Krovetz, T.: Message Authentication on 64-Bit Architectures. In: Biham, E., Youssef, A.M. (eds.) SAC 2006. LNCS, vol. 4356, pp. 327–341. Springer, Heidelberg (2007)
22. Liskov, M., Rivest, R., Wagner, D.: Tweakable Block Ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, p. 31. Springer, Heidelberg (2002)
23. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal of Computing* 17(2), 373–386 (1988)
24. Maurer, U., Renner, R., Holenstein, C.: Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer, Heidelberg (2004)
25. Minematsu, K.: Beyond-Birthday-Bound Security Based on Tweakable Block Cipher. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 308–326. Springer, Heidelberg (2009)
26. Naor, M., Reingold, O.: On the construction of pseudorandom permutations: Luby-Rackoff revisited. *J. of Cryptology* (1999); Preliminary Version: STOC 1997
27. Phan, D.H., Pointcheval, D.: Chosen-Ciphertext Security without Redundancy. In: Lai, C.-S. (ed.) ASIACRYPT 2003. LNCS, vol. 2894, pp. 1–18. Springer, Heidelberg (2003)
28. Rogaway, P., Bellare, M., Black, J.: OCB: A block-cipher mode of operation for efficient authenticated encryption. In: ACM Conference on Computer and Communication Security, pp. 196–205 (2001)
29. Shoup, V.: Sequences of games: a tool for taming complexity in security proofs, <http://eprint.iacr.org/2004/332/>
30. Wang, P., Feng, D., Wu, W.: HCCTR: A variable-input-length enciphering mode. In: Feng, D., Lin, D., Yung, M. (eds.) CISC 2005. LNCS, vol. 3822, pp. 175–188. Springer, Heidelberg (2005)