Luby-Rackoff Ciphers from Weak Round Functions?

Ueli Maurer[1], Yvonne Anne Oswald[1],
Krzysztof Pietrzak[2,⋆], and Johan Sjödin[1,⋆⋆]

[1] Department of Computer Science, ETH Zurich, CH-8092 Zurich, Switzerland
{maurer, sjoedin}@inf.ethz.ch, yoswald@student.ethz.ch
[2] Département d'informatique, Ecole Normale Supérieure, Paris, France
pietrzak@di.ens.fr

Abstract. The Feistel-network is a popular structure underlying many block-ciphers where the cipher is constructed from many simpler rounds, each defined by some function which is derived from the secret key.

Luby and Rackoff showed that the three-round Feistel-network – each round instantiated with a pseudorandom function secure against adaptive chosen plaintext attacks (CPA) – is a CPA secure pseudorandom permutation, thus giving some confidence in the soundness of using a Feistel-network to design block-ciphers.

But the round functions used in actual block-ciphers are – for efficiency reasons – far from being pseudorandom. We investigate the security of the Feistel-network against CPA distinguishers when the only security guarantee we have for the round functions is that they are secure against non-adaptive chosen plaintext attacks (nCPA). We show that in the information-theoretic setting, four rounds with nCPA secure round functions are sufficient (and necessary) to get a CPA secure permutation. Unfortunately, this result does not translate into the more interesting pseudorandom setting. In fact, under the so-called Inverse Decisional Diffie-Hellman assumption the Feistel-network with four rounds, each instantiated with a nCPA secure pseudorandom function, is in general not a CPA secure pseudorandom permutation.

1 Introduction

FEISTEL-NETWORK. The Feistel-network is a popular design approach for block-ciphers where the cipher over $\{0,1\}^{2n}$ is constructed by cascading simpler permutations, each constructed from a round function $\{0,1\}^n \to \{0,1\}^n$. The secret key of the cipher is only used to choose the particular round functions.

⋆ Most of this work was done while the author was a PhD student at ETH where he was supported by the Swiss National Science Foundation, project No. 200020-103847/1. Currently the author is partially supported by the Commission of the European Communities through the IST program under contract IST-2002-507932 ECRYPT.

⋆⋆ This work was partially supported by the Zurich Information Security Center. It represents the views of the authors.

S. Vaudenay (Ed.): EUROCRYPT 2006, LNCS 4004, pp. 391–408, 2006.

Luby-Rackoff Ciphers. In their celebrated paper [LR86] Luby and Rackoff prove that the three-round Feistel-network is an adaptive chosen plaintext (CPA) secure block-cipher – i.e. a pseudorandom permutation (PRP) – if each round is instantiated with an independent CPA secure pseudorandom function (PRF), and with one extra round even adaptive chosen ciphertext (CCA) security is achieved.

Besides reducing PRPs to PRFs, this result also gives some confidence in the soundness of using a Feistel-network to design block-ciphers. But unlike in the Luby-Rackoff ciphers, in most block-ciphers based on Feistel-networks the round functions are not independent (in order to keep the secret key short) and also far from being pseudorandom (for efficiency reasons). Instead, the number of rounds is much larger than four. (which was sufficient for the Luby-Rackoff constructions).

In order to achieve more efficient constructions of PRPs from PRFs, many researcher have investigated the security of weakened versions of the Luby-Rackoff ciphers. Several variations of the ciphers were proven to be pseudorandom where for example the round functions were not required to be independent [Pie90], some round functions were replaced by weaker primitives than PRFs [Luc96, NR02] or the distinguisher was given direct oracle access to some of the round functions [RR00]. These results further fortify the confidence in using Feistel-networks to design block ciphers.

All these relaxed constructions need at least some of the round functions to be CPA secure PRFs in order to get a CPA secure PRP. In this paper, we investigate for the first time – to the best of our knowledge – the CPA security of the permutation one gets by a Feistel-network where none of the round functions is guaranteed to be CPA secure. In particular, we investigate the security of the Feistel-network where each round is instantiated with a *non-adaptive* chosen plaintext (nCPA) secure round function. Although nCPA security is still a strong requirement, this was the weakest natural class of attacks we could imagine which does not make the Feistel-network trivially insecure against CPA attackers. For example round functions which are only secure against known-plaintext attacks (KPA), i.e. look random on random inputs, are too weak.[1]

Pseudo- and Quasirandomness. Informally, a pseudorandom function PRF is a family of functions which can be efficiently computed, and where a random member from the family cannot be distinguished from a uniform random function (URF) by any efficient adversary. Pseudorandom permutations (PRP) are defined analogously. As usual in cryptography, an adversary is efficient if he is in P/poly, i.e. in non-uniform polynomial time (but almost all our results also hold when considering uniform adversaries; the only exception is addressed in Footnote 13). A *quasirandom* function (QRF) (similarly for a quasirandom permutation (QRP)) is defined similar to a *pseudorandom* one but where one does not require the distinguisher or the function to be efficient, only the number of

[1] Just consider a function f which satisfies $f(0\ldots0) = 0\ldots0$ but otherwise looks random. This f is KPA secure as a random query is unlikely to be the all zero string. But a Feistel-network build from such functions will output $0\ldots0$ on input $0\ldots0$ and thus is easily seen not to be CPA (or even nCPA) secure.

queries the distinguisher is allowed to make is bounded. Quasirandomness can be seen as an extension of the concept of statistically close distributions to systems which can be queried interactively.

In order to prove that some system – which is built from *pseudorandom* components – is pseudorandom itself, it is often enough to prove it to be *quasirandom* when the components are replaced by the corresponding ideal systems. In particular, to prove the security of the original three-round Luby-Rackoff cipher it is enough to prove – the purely information-theoretic result – that the network instantiated with URFs is a CPA secure QRP. It then immediately follows that the construction is a CPA secure PRP when the URFs are replaced by CPA secure PRFs, since if it was not a CPA secure PRP, we could use the distinguisher for it to build a distinguisher for the CPA secure PRF (via a standard hybrid argument). Similarly one can easily show that if the round functions are only nCPA or only KPA secure PRFs, the construction is a PRP, but only against the same class of attacks – i.e. nCPA or KPA.

2 Contributions

Our results and related work are summarized in Fig. 2 on page 395. Due to space limitations, some proofs are provided in the full version of this paper only [MOPS06].

(In)secure Relaxations of the Three-Round Luby-Rackoff Cipher. In the pseudo- and quasirandom setting, the three-round Feistel-network is – as mentioned above – $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}, \mathsf{KPA}\}$ secure when the round functions are ATK secure. Moreover it is known that one can replace the first round with a pairwise independent permutation [Luc96, NR99].[2] We further relax this by showing that the function in the last round only needs to be secure against known plaintext attacks (KPA). This resolves an open question posed by Minematsu and Tsunoo in [MT05]. Furthermore, for $\mathsf{ATK} = \mathsf{KPA}$ we show that the first round is not necessary – as opposed to when $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}\}$ – and that it is sufficient to instantiate the (two) round functions with a single instantiation of a KPA secure function.

But the second round seems to be the crucial one for $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}\}$. We show that for constructing a CPA secure permutation – i.e. PRP or QRP depending on the setting – one cannot in general instantiate the second round with a function which is only nCPA secure by constructing a counter-example, i.e. a nCPA secure function such that the three-round Feistel-network with this function in the second, and any random functions in the first and third round can easily be distinguished from a uniformly random permutation (URP) with only three adaptively chosen queries. Similarly, if one instantiates the second round with a KPA secure function, then the construction will in general not even be nCPA secure.

[2] In fact, the permutation must only be such that on any two values, the collision probability on one half of the domain is small. For example one can use one normal Feistel round instantiated with an almost XOR-universal function.

Four Rounds with non-adaptive Round Functions. As a consequence, three rounds with nCPA secure round functions are not enough to get CPA security. On the positive side, we show that one extra nCPA secure round is sufficient (and necessary) in the quasirandom setting. Note that for the translation of a security proof from quasi- to pseudorandom systems – as described at the end of the previous section – it is crucial that we can construct a distinguisher for the components from a distinguisher for the whole system. But here the components have a weaker security guarantee (i.e. nCPA) than what we prove for the whole system (i.e. CPA). So even when we have a CPA distinguisher for the four-round Feistel-network, we cannot construct a nCPA distinguisher for any round function. This is not just a shortcoming of the used approach, but indeed, in the pseudorandom setting the situation is different: we show that here four rounds are not enough to get CPA security. To show this we construct a nCPA secure PRF, such that the four-round Feistel-network with such round functions can easily be distinguished from URP with only three adaptive queries.

This phenomenon – i.e. that some construction implies adaptive security for quasirandom but not for pseudorandom systems – has already been proven [MP04, MPR06, Pie05] for two simple constructions: the sequential composition $f \triangleright g(.) \stackrel{\text{def}}{=} g(f(.))$ and the parallel composition $f \star g(.) \stackrel{\text{def}}{=} f(.) \star g(.)$ (where $\star$ stands for any group operation). The security proofs from [MP04] in the quasirandom setting crucially use the fact that the sequential composition of two permutations is a URPs whenever at least one of the permutations is a URP, similarly the parallel composition of two functions is a URF whenever one of the components is a URF. The Feistel-network does not have this nice property of being ideal whenever one of the components is ideal, and we have to work harder here (using a more general approach from [MPR06]). Our counter-example for the pseudorandom setting – i.e. a four-round Feistel-network with nCPA secure PRFs as round functions that is not a CPA secure PRP – is similar to the counter-examples for sequential and parallel composition shown in [Pie05, Ple05]. In [Ple05], it is shown that the sequential composition of arbitrarily many nCPA secure PRFs will not be a CPA secure PRF in general, whereas for the parallel composition only a counter-example with two components is known [Pie05]. For the Feistel-network we also could only find a counter-example for four rounds. So we cannot rule out the possibility that five or more rounds imply adaptive security. However, if this was the case, then it seems likely that – like for sequential composition [Mye04] – there is no black-box proof for this fact.[3]

Unconditional vs. Conditional Counter-examples. The counter-example showing that the three-round Feistel-network with a nCPA secure PRF

[3] Myers [Mye04] constructs an oracle relative to which there exist PRPs that are nCPA secure, but for which their sequential composition is not a CPA secure PRP. The idea behind this oracle is quite general, and we see no reason (besides being technically challenging) why one should not be able to construct a similar oracle for the Feistel-network, and thus also rule out a black-box proof for showing that the Feistel-network with nCPA secure PRFs as round functions is a CPA secure PRP.

Construction	Quasirandom	Pseudorandom	Reference
$\psi[RRR]$	CPA		[LR86, Mau02]
$\psi[NNN]$	nCPA		§4
$\psi[KKK]$	KPA		§4
$H \triangleright \psi[RR]$	CPA		[Luc96, NR02]
$H \triangleright \psi[RK]$	CPA		§4
$H \triangleright \psi[NK]$	nCPA		§4
$H \triangleright \psi[KK]$	KPA		§4
$\psi[RR]$	KPA (and NOT nCPA)		[MT05] (and §4)
$\psi[K^2]$	KPA		§4
$\psi[RNR]$	NOT CPA		§5
$\psi[RKR]$	NOT nCPA		§5
$\psi[NNNN]$	CPA	NOT CPA (under IDDH)	§6 and §7

Fig. 1. Security of the Feistel-network ψ with various security guarantees on the round functions. Here $\psi[f_1 \cdots f_k](\cdot)$ denotes the k-round Feistel-network with f_i in the i'th round, and $\psi[f^2] \stackrel{\text{def}}{=} \psi[ff]$ – i.e. the same function f in both rounds. Each occurrence of R, N, and K stands for an independent CPA, nCPA, and KPA secure function (i.e. a PRF or a QRF depending on the setting) respectively. The same holds for H which is any "lightweight" permutation from which we only require that the collision probability be small on the left half of the output, an almost pairwise independent permutation or a Feistel round instantiated with an almost XOR-universal function is thus sufficient.

F in the second round is not adaptively secure is unconditional[4] and black-box; with this we mean that we can construct F starting from any (nCPA secure) PRF via a reduction which uses this PRF only as a black-box.[5] As four rounds are enough to get adaptive security for quasirandom systems, there cannot be a black-box counter-example (like for three rounds) for the four (or more) round case. Thus it is not surprising that our counter-example for four rounds is not unconditional. It relies on the so-called Inverse Decisional Diffie-Hellman assumption. The fact that there is no black-box counter-example can be used to show that there is in some sense no "generic" adversary which breaks the adaptive security of the four-round Feistel-network with non-adaptive round functions. What "generic" actually means will not be the topic of this paper, but see Sect. 4 from [Pie06] (in this proceedings) for the corresponding statement for sequential composition.

[4] I.e. we make no other assumption besides the trivially necessary one that pseudorandom functions – which are equivalent to one-way functions [HILL99, GGM86] – exist at all.

[5] We build F from a pseudorandom involution (PRI), how to construct a PRI from a PRP (via a black-box reduction) has been shown in [NR02].

3 Basic Definitions and Random Systems

We use capital calligraphic letters like $\mathcal{X}$ to denote sets, capital letters like X to denote random variables and small letters like x denote concrete values. To save on notation we write X^i for $(X_1, X_2, \ldots, X_i)$.

For $x \in \{0,1\}^{2n}$ we denote with ${}_Lx$ and ${}_Rx$ the left and right half of x respectively, so $x = {}_Lx \| {}_Rx$. Similarly for any function f with range $\{0,1\}^{2n}$, we denote with ${}_Lf$ (${}_Rf$) the function one gets by ignoring the right (left) half of the output of f. For two functions $f(.)$ and $g(.)$ we denote with $f \triangleright g(.) \stackrel{\text{def}}{=} g(f(.))$ the sequential composition of f and g.[6] For a (randomized) function f we denote with $\mathsf{coll}_k(f)$ the collision probability of any fixed k-tuple of distinct inputs, i.e.

$$\mathsf{coll}_k(f) = \max_{x_1,\ldots,x_k} \mathsf{P}(\exists i, j; 1 \le i < j \le k : f(x_i) = f(x_j)).$$

If f denotes a uniform random function with range $\{0,1\}^n$, then $\mathsf{coll}_k(f) \le k^2/2^{n+1}$, this is called the *birthday bound* which we will use quite often.

Definition 1 (Feistel-network). *The (one round) Feistel-network $\psi[f] : \{0,1\}^{2n} \to \{0,1\}^{2n}$ based on a function $f : \{0,1\}^n \to \{0,1\}^n$ is defined as*

$$\psi[f](x) \stackrel{def}{=} (f({}_Lx) \oplus {}_Rx) \| {}_Lx.$$

With $\psi[f_1 \cdots f_k] \stackrel{def}{=} \psi[f_1] \triangleright \psi[f_2] \triangleright \cdots \triangleright \psi[f_k]$ we denote the k-round Feistel-network based on (randomized) round functions $f_1, \ldots, f_k$, here the randomness used by any function is always assumed to be independent of the randomness of the other round functions. The k round Feistel-network where the same instantiation of a function f is used for all rounds is denoted by $\psi[f^k] \stackrel{def}{=} \psi[\underbrace{f \cdots f}_{k \text{ times}}]$.

Random Systems. Many results from this paper are stated and proven in the random systems framework of [Mau02]. A *random system* is a system which takes inputs $X_1, X_2, \ldots$ and generates, for each new input X_i, an output Y_i which depends probabilistically on the inputs and outputs seen so far. We define random systems in terms of the distribution of the outputs Y_i conditioned on $X^i Y^{i-1}$ (i.e. the actual query X_i and all previous input/output pairs $X_1 Y_1, \ldots, X_{i-1} Y_{i-1}$).

Definition 2 (Random systems). *An $(\mathcal{X}, \mathcal{Y})$-*random system $\mathbf{F}$ *is a sequence of conditional probability distributions $\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$ for $i \ge 1$. Here we denote by $\mathsf{P}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}(y_i, x^i, y^{i-1})$ the probability that $\mathbf{F}$ will output y_i on input x_i conditioned on the fact that $\mathbf{F}$ did output y_j on input x_j for $j = 1, \ldots, i-1$.*

As special classes of random systems we will consider *random functions* (which are exactly the stateless random systems) and *random permutations*.

Definition 3 (Random functions and permutations). *A* random function $\mathcal{X} \to \mathcal{Y}$ *(*random permutation *on $\mathcal{X}$) is a random variable which takes as values functions $\mathcal{X} \to \mathcal{Y}$ (permutations on $\mathcal{X}$).*

[6] Note that $f \triangleright g$ is usually denoted with $g \circ f$.

A uniform random function (URF) $\mathbf{R} : \mathcal{X} \to \mathcal{Y}$ *(A* uniform random permutation (URP) $\mathbf{P}$ *on* $\mathcal{X}$*) is a random function with uniform distribution over all functions from* $\mathcal{X}$ *to* $\mathcal{Y}$ *(permutations on* $\mathcal{X}$*). Throughout, the symbols* $\mathbf{R}$ *and* $\mathbf{P}$ *are used for the systems defined above (*$\mathcal{X}, \mathcal{Y}$ *to be understood).*

Indistinguishability of random systems. The distinguishing advantage of a computationally unbounded distinguisher for two random variables A and B is simply the statistical distance of A and B. It is more intricate to define what we mean by the indistinguishability of random systems as here one must specify how the systems can be accessed. For this we define the concept of a distinguisher.

Definition 4. *A* $(\mathcal{Y}, \mathcal{X})$-distinguisher *is a* $(\mathcal{Y}, \mathcal{X})$*-random system which is one query ahead; i.e. it is defined by* $\mathsf{P}^{\mathbf{D}}_{X_i|Y^{i-1}X^{i-1}}$ *instead of* $\mathsf{P}^{\mathbf{D}}_{X_i|Y^{i}X^{i-1}}$ *for all* i*. In particular the first output* $\mathsf{P}^{\mathbf{D}}_{X_1}$ *is defined before* $\mathbf{D}$ *is fed with any input.*

We can now consider the random experiment where a $(\mathcal{Y}, \mathcal{X})$-distinguisher queries a $(\mathcal{X}, \mathcal{Y})$-random system

Definition 5. *With* $\mathbf{D}\Diamond\mathbf{F}$ *we denote the random experiment where a distinguisher* $\mathbf{D}$ *interactively queries a compatible random system* $\mathbf{F}$*.*

We divide distinguishers into classes by posing restrictions on how the distinguisher can access his inputs and produce his queries. In particular the following attacks will be of interest to us:

- CPA: Adaptively Chosen Plaintext Attack; here the adversary can choose the i'th query after receiving the $(i-1)$'th output.
- nCPA: Non-Adaptively Chosen Plaintext Attack; here the distinguisher must choose all queries in advance.
- KPA: Known Plaintext Attack; the queries are chosen uniformly at random.

If $\mathbf{F}$ is a permutation, its inverse $\mathbf{F}^{-1}$ is well-defined and we can consider a

- CCA: Chosen Ciphertext Attack.

which is defined like a CPA but where the attacker can additionally make queries from the inverse direction.

Definition 6. *For* $k \geq 1$*, the two random experiments* $\mathbf{D}\Diamond\mathbf{F}$ *and* $\mathbf{D}\Diamond\mathbf{G}$ *define a distribution over* $\mathcal{X}^k \times \mathcal{Y}^k$*. The* advantage *of* $\mathbf{D}$ *after* k *queries in distinguishing* $\mathbf{F}$ *from* $\mathbf{G}$*, denoted* $\Delta^{\mathbf{D}}_k(\mathbf{F}, \mathbf{G})$*, is the statistical difference between those distributions*[7]

$$\Delta^{\mathbf{D}}_k(\mathbf{F}, \mathbf{G}) \stackrel{def}{=} \frac{1}{2} \sum_{\mathcal{X}^k \times \mathcal{Y}^k} \left| \mathsf{P}^{\mathbf{D}\Diamond\mathbf{F}}_{X^k Y^k} - \mathsf{P}^{\mathbf{D}\Diamond\mathbf{G}}_{X^k Y^k} \right|. \tag{1}$$

[7] This definition has a natural interpretation in the random experiment where we first toss a uniform random coin $C \in \{0, 1\}$. Then we let $\mathbf{D}$ (which has no a priori information on C) make k queries to a system $\mathbf{H}$ where $\mathbf{H} \equiv \mathbf{F}$ if $C = 0$ and $\mathbf{H} \equiv \mathbf{G}$ if $C = 1$. Here the expected probability that an optimal guess on C based on the k inputs and outputs of $\mathbf{H}$ will be correct is $1/2 + \Delta^{\mathbf{D}}_k(\mathbf{F}, \mathbf{G})/2$.

The advantage of the best ATK*-distinguisher making* k *queries for* $\mathbf{F}$ *and* $\mathbf{G}$ *is*

$$\Delta_k^{\mathsf{ATK}}(\mathbf{F},\mathbf{G}) \stackrel{def}{=} \max_{\mathsf{ATK}-distinguisher\ \mathbf{D}} \Delta_k^{\mathbf{D}}(\mathbf{F},\mathbf{G}).$$

Pseudorandomness. We denote with $\mathbf{Adv}_{PRP}^{\mathsf{ATK}}(\mathsf{F},t,k)$ the distinguishing advantage of the best oracle circuit for F from a URP $\mathbf{P}$ where the circuit must be of size at most t and make at most k ATK-queries to its oracle. So $\mathbf{Adv}$ is defined similarly to Δ but with an additional restriction on the size of the distinguisher. In particular $\mathbf{Adv}_{PRP}^{\mathsf{ATK}}(\mathbf{F},\infty,k) = \Delta_k^{\mathsf{ATK}}(\mathbf{F},\mathbf{P})$. $\mathbf{Adv}_{PRF}^{\mathsf{ATK}}$ is defined similarly, but with $\mathbf{P}$ replaced by $\mathbf{R}$.

Informally, a family of keyed functions F indexed by a security parameter $\gamma \in \mathbb{N}$ is an ATK-secure pseudorandom function (PRF) if F (with security parameter γ) can be computed in uniform polynomial (in γ) time, and for any polynomial $p(.)$ the distinguishing advantage $\mathbf{Adv}_{PRF}^{\mathsf{ATK}}(\mathsf{F},p(\gamma),p(\gamma))$ is negligible in γ (for a key chosen uniformly at random). Pseudorandom permutations (PRP) are defined similarly but using $\mathbf{Adv}_{PRP}^{\mathsf{ATK}}$, and where we additionally require that F (for any security parameter and key) is a permutation.

We usually use sans-serif fonts like F to denote systems which can be efficiently computed (in particular pseudorandom systems), and bold fonts like $\mathbf{F}$ to denote quasirandom and ideal systems.

4 Relaxations of the Three-Round Luby-Rackoff Cipher

Let us first state some results for the three-round Feistel-network.

Proposition 1. *For any* $\mathsf{ATK} \in \{\mathsf{CPA},\mathsf{nCPA},\mathsf{KPA}\}$ *and function* $\mathbf{F}$

$$\Delta_k^{\mathsf{ATK}}(\psi_{2n}[\mathbf{FFF}],\mathbf{P}) \leq 3 \cdot \Delta_k^{\mathsf{ATK}}(\mathbf{F},\mathbf{R}) + 2 \cdot \frac{k^2}{2^{n+1}}. \tag{2}$$

The analogous statement also holds in the computational case: for any $\mathsf{ATK} \in \{\mathsf{CPA},\mathsf{nCPA},\mathsf{KPA}\}$ *and any efficient function* F

$$\mathbf{Adv}_{PRP}^{\mathsf{ATK}}(\psi_{2n}[\mathsf{FFF}],t,k) \leq 3 \cdot \mathbf{Adv}_{PRF}^{\mathsf{ATK}}(\mathsf{F},t',k) + 2 \cdot \frac{k^2}{2^{n+1}}, \tag{3}$$

where $t' = \mathsf{poly}(t,k)$ *for some polynomial* poly *which accounts for the overhead implied by the reduction we make.*

The classical result of Luby and Rackoff [LR86], states that the Feistel-network with three independent PRF rounds is a CPA secure PRP – i.e (3) for CPA.

Luby and Rackoff proved this result directly. One gets a simpler proof by first showing that the three-round Feistel-network with URFs $\mathbf{R}$ is a CPA secure QRP as this is a purely information-theoretic statement. In particular it was shown in [Mau02] that[8]

$$\Delta_k^{\mathsf{CPA}}(\psi_{2n}[\mathbf{RRR}],\mathbf{P}) \leq 2 \cdot \frac{k^2}{2^{n+1}}, \tag{4}$$

[8] This bound has been improved – for various number of rounds – in a series of papers. The latest [Pat04] by Patarin presents the best possible security for up to $k \ll 2^n$ (and not just $k \ll 2^{n/2}$) queries, using five rounds which is also necessary.

from which Proposition 1 directly follows using a standard hybrid argument.[9] Lucks showed [Luc96] (see also [NR02]) that the first round in the three-round Luby-Rackoff cipher can be replaced with a much weaker primitive which only must provide some guarantee on the collision probability on the left half of the output (for any two fixed inputs). In particular, an almost pairwise independent permutation or a Feistel-round with an almost XOR-universal function will do.

Proposition 2. *For any* $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}, \mathsf{KPA}\}$*, any functions* $\mathbf{F}$*,* $\mathbf{G}$*, and any permutation* H

$$\Delta_k^{\mathsf{ATK}}(\mathsf{H} \triangleright \psi_{2n}[\mathbf{FG}], \mathbf{P}) \leq \Delta_k^{\mathsf{ATK}}(\mathbf{F}, \mathbf{R}) + 2 \cdot \Delta_k^{\mathsf{KPA}}(\mathbf{G}, \mathbf{R}) + \mathsf{coll}_k({}_L\mathsf{H}) + 2 \cdot \frac{k^2}{2^{n+1}}. \quad (5)$$

The analogous statement also holds in the computational case: for any $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}, \mathsf{KPA}\}$*, any efficient functions* F, G*, and any efficient permutation* H

$$\begin{aligned} &\mathbf{Adv}_{PRP}^{\mathsf{ATK}}(\mathsf{H} \triangleright \psi_{2n}[\mathsf{FG}], t, k) \quad (6) \\ &\leq \mathbf{Adv}_{PRF}^{\mathsf{ATK}}(\mathsf{F}, t', k) + 2 \cdot \mathbf{Adv}_{PRF}^{\mathsf{KPA}}(\mathsf{G}, t', k) + \mathsf{coll}_k({}_L\mathsf{H}) + 2 \cdot \frac{k^2}{2^{n+1}}, \end{aligned}$$

where $t' = t + \mathsf{poly}(n, k)$ *for some polynomial* poly *which accounts for the overhead implied by the reduction we make.*

Let us stress that (6) does *not* directly follow from (5).[10] The proof of Proposition 2 is given in the full version of this paper [MOPS06].

We relax the construction further for $\mathsf{ATK} = \mathsf{KPA}$ by showing that the first round can be removed completely (as opposed to when $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}\}$)[11]. The round functions can also be replaced by a *single* instantiation of a KPA secure function. Note that the resulting construction is an involution, i.e. has the structural property of being self inverse. This result also generalizes Lemma 2.2 of [MT05] which states that the two round Feistel-network with CPA secure PRFs is a KPA secure PRP.

---

[9] The argument goes as follows for pseudorandom systems: suppose there is an efficient $\mathsf{ATK} \in \{\mathsf{CPA}, \mathsf{nCPA}, \mathsf{KPA}\}$ distinguisher A for $\psi_{2n}[\mathsf{FFF}]$ and $\mathbf{P}$, then by (4) this A will also distinguish $\psi_{2n}[\mathsf{FFF}]$ from $\psi_{2n}[\mathbf{RRR}]$. Consider the hybrids $H_0 = \psi_{2n}[\mathsf{FFF}], H_1 = \psi_{2n}[\mathbf{R}\mathsf{FF}], \ldots, H_3 = \psi_{2n}[\mathbf{RRR}]$. By the triangle inequality there is an $0 \leq i \leq 2$ (say $i = 1$) such that A can distinguish H_i from H_{i+1}. Now, the distinguisher which – with access to an oracle G (implementing either F or $\mathbf{R}$) – simulates $A \Diamond \psi_{2n}[\mathbf{R}\mathsf{GF}]$ and outputs the output of A is an efficient ATK-distinguisher for F with the same advantage as A's advantage for H_1 and H_2. The corresponding argument also holds in the quasirandom setting.

[10] The reason why a reduction – like the simple argument to show that Proposition 1 follows from (4) – fails here, is that the KPA security guarantee for one of the components is weaker than the CPA security for the whole construction. But fortunately the *proof* of (5) is such that it easily translates to the pseudorandom setting.

[11] $\psi_{2n}[\mathbf{RR}]$ can be distinguish from $\mathbf{P}$ with two non-adaptively chosen queries: query $0^n \| 0^n \mapsto {}_L y \| {}_R y$ and $0^n \| 1^n \mapsto {}_L y' \| {}_R y'$, and output 1 if ${}_R y \oplus {}_R y' = 1^n$ and 0 otherwise.

Proposition 3. *For any function* $\mathbf{F}$

$$\Delta_k^{\mathsf{KPA}}(\psi_{2n}[\mathbf{F}^2], \mathbf{P}) \leq \Delta_{2k}^{\mathsf{KPA}}(\mathbf{F}, \mathbf{R}) + 4 \cdot \frac{k^2}{2^{n+1}}. \tag{7}$$

The analogous statement also holds in the computational case: for any function $\mathbf{F}$ *(in particular any efficient function* F*)*

$$\mathbf{Adv}_{PRP}^{\mathsf{KPA}}(\psi_{2n}[\mathbf{F}^2], t, k) \leq \mathbf{Adv}_{PRF}^{\mathsf{KPA}}(\mathbf{F}, t', 2k) + 4 \cdot \frac{k^2}{2^{n+1}}, \tag{8}$$

where $t' = t + \mathsf{poly}(n, k)$ *for some polynomial* poly *which accounts for the overhead implied by the reduction we make.*

The proof is in the full version of this paper [MOPS06]. Note that unlike in the previous propositions, here we do not require the round function $\mathbf{F}$ to be efficient in the computational case (the reason is that in the proof we do not need the distinguisher to simulate any round function).

5 The Second Round Is Crucial

In the previous section we have seen that in the classical three-round Luby-Rackoff cipher the first and third round function need not be CPA secure. In this section we will see that the security requirements for the second round cannot be relaxed. We only give proof sketches for the propositions of this section. Detailed proofs can be found in the full version.

The following proposition states that to achieve CPA security in general it is not sufficient that the second round function is nCPA secure. There exists a nCPA secure function, such that the three-round Feistel-network with this function in the second, and any random functions in the first and third round, is not CPA secure.

Proposition 4. *There exists a function* $\mathbf{F}$ *such that for any functions* $\mathbf{G}$ *and* $\mathbf{G}'$ *(in particular for* $\mathbf{G} = \mathbf{R}$ *and* $\mathbf{G}' = \mathbf{R}$*)*

$$\Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{R}) \leq 4 \cdot \frac{k^2}{2^{n+1}} \quad \textit{and} \quad \Delta_2^{\mathsf{CPA}}(\psi_{2n}[\mathbf{GFG}'], \mathbf{P}) \geq 1 - 2^{-n+1}.$$

The analogous statement also holds in the computational case: (informal) there is a nCPA *secure PRF* F *such that* $\psi_{2n}[\mathbf{GFG}']$ *is not a* CPA *secure PRP for any (not necessarily efficient) functions* $\mathbf{G}$ *and* $\mathbf{G}'$*.*

Proof (sketch). Let us first consider the quasirandom statement. Let $\mathbf{I}$ be a uniform random involution, i.e. $\mathbf{I}(\mathbf{I}(x)) = x$ for all x. Now, $\mathbf{F}$ is simply defined as $\mathbf{F}(x) = x \oplus \mathbf{I}(x)$, note that this $\mathbf{F}$ satisfies $\mathbf{F}(x) = \mathbf{F}(x \oplus \mathbf{F}(x))$ for all x.

The nCPA security of $\mathbf{F}$ (which is simply the nCPA security of $\mathbf{I}$) can be bounded as stated in the proposition by standard techniques. Furthermore, $\psi_{2n}[\mathbf{GFG}']$ can easily be distinguished from $\mathbf{P}$ with two adaptively chosen

queries as follows. After a first query $0^n\|0^n$, the output ${}_LY\|Z$ contains the output Z of the internal function $\mathbf{F}$. Now make a second query $0^n\|Z$. If the (unknown) input to $\mathbf{F}$ in the first query was some value V, then in this query it will be $V \oplus Z$, and as $\mathbf{F}$ satisfies $\mathbf{F}(V) = \mathbf{F}(V \oplus \mathbf{F}(V)) = \mathbf{F}(V \oplus Z)$, the output of $\mathbf{F}$ will again be Z, and the overall output will be $({}_LY \oplus Z)\|Z$. The proposition follows as the output of $\mathbf{P}$ will satisfy such a relation with probability at most 2^{-n+1}.

The corresponding statement for the pseudorandom setting is proven almost identically. The only difference is that we need to use a CPA secure pseudorandom involution instead of the uniform random involution. It is shown in [NR02] how to construct a pseudorandom involution from any CPA secure PRF. □

The next proposition states that the network will in general not (even) be nCPA secure when the second round function is only secure against KPAs.

Proposition 5. *There exists a function* $\mathbf{F}$ *such that for any functions* $\mathbf{G}$ *and* $\mathbf{G}'$

$$\Delta_k^{\mathsf{KPA}}(\mathbf{F}, \mathbf{R}) \leq \frac{k^2}{2^{n+1}}, \quad \textit{and} \quad \Delta_2^{\mathsf{nCPA}}(\psi_{2n}[\mathbf{GFG}'], \mathbf{P}) \geq 1 - 2^{-n+2}.$$

The analogous statement also holds in the computational case: (informal) there is a KPA *secure PRF* F *such that* $\psi_{2n}[\mathbf{GFG}']$ *is not a* nCPA *secure PRP for any (not necessarily efficient) functions* $\mathbf{G}$ *and* $\mathbf{G}'$*.*

Proof (sketch). Let us first consider the statement in the quasirandom setting. Let $\mathbf{F}$ be a URF which ignores the first input bit, i.e. for all $x \in \{0,1\}^{n-1}$ we have $\mathbf{F}(0\|x) = \mathbf{F}(1\|x)$. The KPA security of $\mathbf{F}$ follows from the fact that $\mathbf{F}$ looks completely random unless we happen to query two queries of the form $0\|x$ and $1\|x$. By the birthday bound the probability that this happens after k queries is at most $\frac{k^2}{2^{n+1}}$. Furthermore, $\psi_{2n}[\mathbf{GFG}']$ can be distinguished from $\mathbf{P}$ with two non-adaptively chosen queries. For instance on input $0^n\|0^n$ and $0^n\|(1\|0^{n-1})$, the right half of the output will be identical.

The corresponding statement in the pseudorandom setting is proven exactly as above, except that we have to use a PRF F instead of $\mathbf{F}$. □

6 Four nCPA Secure Rounds, the Quasirandom Case

In this section we will show that the four-round Feistel-network with nCPA secure QRFs is a CPA secure QRP. This is also the best possible as in Sect. 5 we showed that four rounds are also necessary. The theorem is even stronger as the third and fourth round function must only be KPA secure QRFs.

Theorem 1. *For any functions* $\mathbf{F}$ *and* $\mathbf{G}$

$$\Delta_k^{\mathsf{CPA}}(\psi_{2n}[\mathbf{FFGG}], \mathbf{P}) \leq 4 \cdot \Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{R}) + 3 \cdot \Delta_k^{\mathsf{KPA}}(\mathbf{G}, \mathbf{R}) + 9 \cdot \frac{k^2}{2^{n+1}}.$$

To prove this theorem we use Theorem 2 from [MPR06] which, for the special case of the four-round Feistel-network, is given as Proposition 6 below. The proposition bounds the security of a composition against a "strong" attacker sATK (in particular CPA) in terms of the security of the components against "weak" attackers wATK_i (in particular nCPA or KPA).

The proposition uses the concept of conditions defined for random systems which we only define informally here (see [MPR06] for a formal definition): With $\mathbf{F}^{\mathcal{A}}$ we denote the random system $\mathbf{F}$, but which additionally defines an internal binary random variable after each query (called a condition). Let $A_i \in \{0,1\}$ denote the condition after the i'th query. We set $A_0 = 0$ and require the condition to be monotone which means that $A_i = 1 \Rightarrow A_{i+1} = 1$ (i.e. when the condition failed, it will never hold again). Let $\overline{a}_i$ denote the event $A_i = 1$, then

$$\nu^{\mathsf{ATK}}(\mathbf{F}^{\mathcal{A}}, \overline{a}_k) \stackrel{\text{def}}{=} \max_{\mathsf{ATK}-distinguisher\ \mathbf{D}} \mathsf{P}_{\overline{a}_k}^{\mathbf{D}\Diamond\mathbf{F}^{\mathcal{A}}}, \tag{9}$$

denotes the advantage of the best ATK distinguisher to make the condition fail after at most k queries to $\mathbf{F}^{\mathcal{A}}$.

Proposition 6. *If for any* $(\{0,1\}^n, \{0,1\}^n)$*-random system with a condition* $\mathbf{F}^{\mathcal{A}}$

$$\nu^{\mathsf{sATK}}(\psi_{2n}[\mathbf{F}^{\mathcal{A}}\mathbf{RRR}], \overline{a}_k) \leq \nu^{\mathsf{wATK}_1}(\mathbf{F}^{\mathcal{A}}, \overline{a}_k) + \alpha_1 \tag{10}$$
$$\nu^{\mathsf{sATK}}(\psi_{2n}[\mathbf{RF}^{\mathcal{A}}\mathbf{RR}], \overline{a}_k) \leq \nu^{\mathsf{wATK}_2}(\mathbf{F}^{\mathcal{A}}, \overline{a}_k) + \alpha_2 \tag{11}$$
$$\nu^{\mathsf{sATK}}(\psi_{2n}[\mathbf{RRF}^{\mathcal{A}}\mathbf{R}], \overline{a}_k) \leq \nu^{\mathsf{wATK}_3}(\mathbf{F}^{\mathcal{A}}, \overline{a}_k) + \alpha_3 \tag{12}$$
$$\nu^{\mathsf{sATK}}(\psi_{2n}[\mathbf{RRRF}^{\mathcal{A}}], \overline{a}_k) \leq \nu^{\mathsf{wATK}_4}(\mathbf{F}^{\mathcal{A}}, \overline{a}_k) + \alpha_4 \tag{13}$$

for some attacks wATK_1*,* wATK_2*,* wATK_3*,* wATK_4*,* sATK *and some* $\alpha_1, \alpha_2, \alpha_3$*,* $\alpha_4 \geq 0$*, then for any* $\mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3, \mathbf{F}_4$

$$\Delta_k^{\mathsf{sATK}}(\psi_{2n}[\mathbf{F}_1\mathbf{F}_2\mathbf{F}_3\mathbf{F}_4], \psi_{2n}[\mathbf{RRRR}]) \leq \sum_{i=1}^{4} (\Delta_k^{\mathsf{wATK}_i}(\mathbf{F}_i, \mathbf{R}) + \alpha_i).$$

To apply this proposition we must show that equations (10), (11), (12) and (13) hold for some attack wATK_i and α_i for $i = 1, 2, 3, 4$.

In the full version [MOPS06] we prove the following claim, from which Theorem 1 now follows.

Claim 1. *Equation (10) - (13) are satisfied for any function with a condition* $\mathbf{F}^{\mathcal{A}}$*,* $\mathsf{sATK} = \mathsf{CPA}$*, and*

$$\Big(\mathsf{wATK}_i, \alpha_i\Big) = \begin{cases} \Big(\mathsf{nCPA},\ 2 \cdot \frac{k^2}{2^{n+1}}\Big) & \text{if} \quad i = 1 \\ \Big(\mathsf{nCPA},\ 2 \cdot \frac{k^2}{2^{n+1}} + 2 \cdot \Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{R})\Big) & \text{if} \quad i = 2 \\ \Big(\mathsf{KPA},\ 3 \cdot \frac{k^2}{2^{n+1}} + \Delta_k^{\mathsf{KPA}}(\mathbf{F}, \mathbf{R})\Big) & \text{if} \quad i = 3 \\ \Big(\mathsf{KPA},\ 2 \cdot \frac{k^2}{2^{n+1}}\Big) & \text{if} \quad i = 4\ . \end{cases}$$

7 Four nCPA Secure Rounds, the Pseudorandom Case

In this section we again investigate the CPA security of the four-round Feistel-network with nCPA secure round functions, but this time for *pseudorandom* systems. We show that here the situation is dramatically different from the quasirandom setting by constructing a nCPA secure PRF where the four-round Feistel-network with this PRF as round function is not CPA secure.

This PRF is defined over some group, and to prove the nCPA security we assume that the so-called *inverse decisional Diffie-Hellman* (IDDH) is hard in this group. Informally, the IDDH assumption requires that for a generator g and random x, y it is hard do distinguish the triple (g, g^x, g^y) from $(g, g^x, g^{x^{-1}})$.

Theorem 2. *(Informal) Under the IDDH assumption there exists a* nCPA *secure PRF* F *such that the four-round Feistel-network where each round is instantiated with* F *(with independent keys) is* not *a* CPA *secure pseudorandom permutation.*

This theorem follows from Lemma 1 below which states that there exist nCPA secure PRFs $\mathsf{F}_1, \mathsf{F}_2, \mathsf{F}_3$ such that the left half of the *three* round Feistel-network ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$ is not a CPA secure PRF. This implies that also $\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3\mathsf{G}]$ is not a CPA secure PRP for any G (and thus proves Theorem 2) as follows. By the so-called PRF/PRP Switching Lemma any CPA secure PRP P is also a CPA secure PRF. Clearly, then also ${}_L\mathsf{P}$ must be a CPA secure PRF. Now, by Lemma 1 ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3] = {}_R\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3\mathsf{G}]$ is not a CPA secure PRF, so $\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3\mathsf{G}]$ cannot be a CPA secure PRP.[12]

Lemma 1. *Under the IDDH-assumption there exist* nCPA *secure PRFs* $\mathsf{F}_1, \mathsf{F}_2, \mathsf{F}_3$ *such that* ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$ *is not a* CPA *secure PRF: it can be distinguished efficiently from a URF with only three (adaptive) queries with high probability.*

Outline For this Section. In §7.1 we give a more formal definition of the IDDH assumption. Then, in §7.2 we first show the construction from [Ple05] of a nCPA secure PRF whose sequential composition will not be CPA secure. This extremely simple and intuitive construction is the basis for the (more involved) counter-example for the Feistel-network (i.e. Lemma 1) given in §7.3.

7.1 The Non-uniform IDDH Assumption

Below we define the IDDH assumption which is similar (and easily seen to imply) the well known decisional Diffie-Hellman assumption. Throughout, we will work with hardness assumptions in a non-uniform model of computation (i.e. we

[12] The lemma talks about three different F_i's (and in the proof we really construct a different F_i for every round), but the theorem is stated for a single F. This does not really make a difference. For example this single F can be defined as behaving like F_i with probability $1/3$ for $i \in \{1, 2, 3\}$. Then with constant probability 3^{-3} the $\psi_{2n}[\mathsf{FFF}]$ behaves like $\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$.

require hardness against polynomial size circuit families and not just any fixed Turing machine).[13]

Let $\mathcal{G}$ denote an efficiently computable family of groups indexed by a security parameter $n \in \mathbb{N}$. By efficiently computable we mean that one can efficiently (i.e. in time polynomial in n) sample a group (together with a generator) from the family, and efficiently compute the group operations therein. Abusing notation we denote with $(G, g) = \mathcal{G}(n)$ any group/generator pair for security parameter n.

The IDDH assumption is hard in $\mathcal{G}$ if for $(G, g) = \mathcal{G}(n)$ polynomial size circuits have negligible advantage guessing whether for a given triple (g, g^x, g^y) the y is random or computed as $y = x^{-1}$, more formally

Definition 7 (non-uniform IDDH). *For a group G and a generator g of G*

$$\mathbf{Adv}_{IDDH}(G, g, s) \stackrel{def}{=} \max_{C, |C| \leq s} \left| \Pr_{x} \left[C(g, g^x, g^{x^{-1}}) = true \right] - \Pr_{x,y} \left[C(g, g^x, g^y) = true \right] \right|,$$

where the probability is over the random choice of $x, y \in [1, \ldots, |G|]$. We say that IDDH is hard in $\mathcal{G}$ if for any polynomial $p(.)$

$$\mathbf{Adv}_{IDDH}(\mathcal{G}(n), p(n)) = \mathsf{negl}(n).$$

7.2 Counter-Example for Sequential Composition from [Ple05]

In this section we construct a simple PRF F, but where the sequential composition of (arbitrary many) such F (with independent keys) is not CPA secure.

F is based on some prime order cyclic group $(G, g) = \mathcal{G}(n)$ where the IDDH problem is hard and where the elements of the group can be efficiently and densely encoded into $\{0,1\}^n$ (with dense we mean that all but a negligible fraction of the strings should correspond to an element of the group).[14] For example we can take the subgroup of prime order q of $\mathbb{Z}_p^*$ where p is a safe prime (i.e. $2q + 1$) and q is close to 2^n ([Dam04] describes how to embed such a G into $\{0,1\}^n$).

Let $[.] : \mathcal{G}(n) \to \{0,1\}^n$ denote an (efficient) embedding of $\mathcal{G}$ into bitstrings (to save on notation we let $[a, b]$ denote the concatenation of $[a]$ and $[b]$). Let

[13] In cryptography security usually means security against non-uniform (and not just uniform) adversaries, and thus also the hardness assumptions used are usually non-uniform, though this is sometimes not explicitly stated as the security proofs work in both settings – i.e. a uniform (non-uniform) assumption implies hardness against uniform (non-uniform) adversaries. But here this is not quite the case, we do not know how to prove a uniform version of Lemma 1. (But one can do so under a somewhat stronger assumption than IDDH. Loosely speaking, this assumption is IDDH but where the attacker can also choose the generators to be used in the challenge.)

[14] For this construction we actually do not need this embedding, we could define F directly over the group. But we will need it (or more precisely, the fact that if X is in the range of F, also $X \oplus R$ for a random bitstring R is in the range with overwhelming probability) when we extend this construction to get the counter-example for the Feistel-network in the next section.

$\mathsf{R} : \mathcal{K} \times \{0,1\}^{4n} \to \mathbb{Z}_q^4$ be any nCPA secure PRF. Now consider the following definition of a nCPA secure PRF $\mathsf{F} : \{0,1\}^{4n} \to \{0,1\}^{4n}$ with secret key ($k_1 \in \mathcal{K}, x \in \mathbb{Z}_q^*$).

The first thing F does on input $(\alpha, \beta, \gamma, \delta) \in \{0,1\}^{4n}$ is to generate some pseudorandom values using R, i.e.

$$(r_1, r_2, r_3, r_4) \leftarrow \mathsf{R}(k_i, \alpha, \beta, \gamma, \delta). \tag{14}$$

Further, if there exists $(a, b, c, d) \in G^4$ s.t. $\alpha = [a], \beta = [b], \gamma = [c], \delta = [d]$ then F outputs (here x^{-1} is the inverse of x in $\mathbb{Z}_q^*$)

$$\mathsf{F}([a, b, c, d]) \to ([a^{xr_1}, b^{r_1}, c^{x^{-1}r_2}, d^{r_2}]), \tag{15}$$

with r_1, r_2 generated as in (14). On the remaining inputs (which are a negligible fraction of $\{0,1\}^{4n}$) F outputs just the (pseudo) random values $[g^{r_1}, g^{r_2}, g^{r_3}, g^{r_4}]$.

Now consider the cascade $\mathsf{F}' \triangleright \mathsf{F}'' \triangleright \mathsf{F}'''$ of three independent F's (with corresponding keys (x_1, k_1), (x_2, k_2), and (x_3, k_3)). Make a first query $[g, g, g, g]$

$$\mathsf{F}' \triangleright \mathsf{F}'' \triangleright \mathsf{F}'''([g, g, g, g]) \to [g^{x_1 x_2 x_3 r}, g^r, g^{x_1^{-1} x_2^{-1} x_3^{-1} r'}, g^{r'}].$$

Then the output will have the form $g^{x_1 x_2 x_3 r}, g^r, g^{x_1^{-1} x_2^{-1} x_3^{-1} r'}, g^{r'}$ for some r, r'. Now exchange the right and the left half of this output and use it as the second query

$$\mathsf{F}' \triangleright \mathsf{F}'' \triangleright \mathsf{F}'''([g^{x_1^{-1} x_2^{-1} x_3^{-1} r'}, g^{r'}, g^{x_1 x_2 x_3 r}, g^r]) \to [g^{r''}, g^{r''}, g^{r'''}, g^{r'''}]$$

so the output is of the form $[u, u, v, v]$ for some u, v and thus can be distinguished from random. Therefore $\mathsf{F}' \triangleright \mathsf{F}'' \triangleright \mathsf{F}'''$ is not a CPA secure PRF. This proves that the sequential composition of nCPA secure PRFs does not yield a CPA secure function in general. Note that this distinguishing attack works for any number of rounds, not just three. In the full version of this paper [MOPS06] we prove the following lemma which states that F is an nCPA secure PRF if IDDH is hard in $\mathcal{G}$ and R is a nCPA secure PRF.

Lemma 2. *Let g be any generator of the group over which F is defined, then*

$$\mathbf{Adv}_{PRF}^{\mathsf{nCPA}}(\mathsf{F}, k, s) \leq 6k \cdot \mathbf{Adv}_{IDDH}(\mathsf{F}, g, s') + \mathbf{Adv}_{PRF}^{\mathsf{nCPA}}(\mathsf{R}, k, s'),$$

where $s' = s + \mathsf{poly}(k, n)$ for some polynomial poly which accounts for the overhead implied by the reduction we make.

7.3 Proof of Lemma 1

The Feistel-network can be seen as a sequential composition of the round functions, but where one additionally XORs the input to the i'th round function to the output of the $(i+1)$'th round function. So it is not surprising that we can use F_i's similar to the F from the previous section to prove Lemma 1. But the $\mathsf{F}_1, \mathsf{F}_2$, and F_3 (from the statement of the lemma) are a bit more complicated as

we have to "work around" this additional XORs. Like F, each F_i has a $k_i \in \mathcal{K}$ as part of its secret key. Moreover F_1 has a $x \in \mathbb{Z}_q^*$ and $s, t \in \{0,1\}^n$, F_2 has a $y \in \mathbb{Z}_q^*$, and F_3 a $z \in \mathbb{Z}_q^*$ as keys. On input $(\alpha, \beta, \gamma, \delta) = [a, b, c, d]$ the F_i's are defined as (with the r_i's generated as in (14))

$$\mathsf{F}_1([a,b,c,d]) \to \begin{cases} [g^{xr_1}, g^{r_1}], s, t & \text{if } [a,b,c,d] = [0,0,0,0]; \\ [0,0,0,0] & \text{elseif } c = d^x; \\ [g^{xr_1}, g^{r_1}, ([\gamma \oplus s]^{-1})^{x^{-1}r_2}, ([\delta \oplus t]^{-1})^{r_2}] & \text{elseif } [a,b] = [0,0]; \\ [g^{r_1}, g^{r_2}, g^{r_3}, g^{r_4}] & \text{otherwise.} \end{cases}$$

$$\mathsf{F}_2([a,b,c,d]) \to \quad [c^{y^{-1}r_1}, d^{r_1}, a^{yr_2}, b^{r_2}]$$

$$\mathsf{F}_3([a,b,c,d]) \to \begin{cases} [0,0,0,0] & \text{if } b^z = a; \\ [a^{z^{-1}r_1}, b^{r_1}, c^{zr_2}, d^{r_2}] & \text{otherwise.} \end{cases}$$

Proof (of Lemma 1). The lemma follows from Claim 2 and 3 below. □

Claim 2. *One can distinguish* ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$ *from a URF with three adaptively chosen queries with advantage almost* 1.

Proof (sketch). In Fig. 2 we demonstrate an adaptive three query distinguishing attack on ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$. In the figure, values which are not relevant for the attack are denoted by $*$. All r_i' values are random, but not necessarily equal to a random value generated by a round function (i.e. as in (14)).[15] To see that this is a legal attack note that every query Q_i can be computed from the previous output O_{i-1}. That the values will really have the form as described in the attack can be verified from the definition of the F_i's.[16] Since the third output starts with $[0,0]$ it can be distinguished from a random output with high probability. □

Claim 3. $\mathsf{F}_1, \mathsf{F}_2$, *and* F_3 *are* nCPA *secure PRFs if IDDH is hard in* $\mathcal{G}$.

Proof (sketch). The nCPA security of the F_i's follows from the nCPA security of F from the previous section as stated in Lemma 2: F_2 is exactly F, so there is nothing else to prove here. The function F_3 behaves exactly as F unless it is queried on an input $[a,b,c,d]$ which satisfies $b^z = a$ for a random z. The probability that this happen on any (non-adaptive) query is just $|G|^{-1}$ (and thus exponentially small even after taking the union bound over all polynomially many queries). For the somewhat longer argument for F_1, we refer to the full version [MOPS06]. □

[15] For instance, r_1' is the first random value generated by F_1 and r_2' is the product of r_1' and the second random value generated by F_2.

[16] Actually, there is an exponentially small probability that the values will not have that form, namely when the input to some round function "by chance" satisfies a condition that is checked. E.g. when R_1^3 is of the form $[b^z, b, c, d]$, then the "$b^z = a$" case of F_3 applies, which is only supposed to happen in the second and third query.

$$
\begin{array}{ll}
{}_LQ_1 : [0,0,0,0] & {}_RQ_1 : [0,0,0,0] \\
R_1^2 : [g^{xr'_1}, g^{r'_1}], s, t & \\
R_1^3 : *, *, [g^{xyr'_2}, g^{r'_2}] & \\
O_1 : *, *, [g^{xyzr'_3}] \oplus s, [g^{r'_3}] \oplus t & \\
 & \\
{}_LQ_2 : [0,0], [g^{xyzr'_3}] \oplus s, [g^{r'_3}] \oplus t & {}_RQ_2 : [0,0,0,0] \\
R_2^2 : [g^{xr'_4}, g^{r'_4}, g^{yzr'_5}, g^{r'_5}] & \\
R_2^3 : [g^{zr'_6}, g^{r'_6}], *, * & \\
O_2 : [g^{xr'_4}, g^{r'_4}, g^{yzr'_5}, g^{r'_5}] & \\
 & \\
{}_LQ_3 : [0,0,g^{xr'_4}, g^{r'_4}] & {}_RQ_3 : [0,0,g^{yzr'_5}, g^{r'_5}] \\
R_3^2 : [0,0,g^{yzr'_5}, g^{r'_5}] & \\
R_3^3 : [g^{zr'_7}, g^{r'_7}], *, * & \\
O_3 : [0,0,g^{yzr'_5}, g^{r'_5}] &
\end{array}
$$

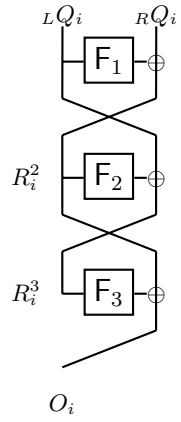

Fig. 2. An adaptive three query distinguishing attack for ${}_L\psi_{2n}[\mathsf{F}_1\mathsf{F}_2\mathsf{F}_3]$

8 Some Remarks on CCA Security

We have shown that the four-round Feistel-network with nCPA secure round functions is CPA secure in the information-theoretic, but in general not in the computational setting. A natural question is to ask how many rounds are necessary/not sufficient to achieve CCA security. In this section we state some observations. The full version of this paper addresses this question in more detail.

In order to get a CCA secure quasirandom permutations (QRP), it is enough – by the following statement (taken from [MPR06]) – to cascade two nCPA secure QRPs (the second in inverse direction)

$$\Delta_k^{\mathsf{CCA}}(\mathbf{F} \triangleright \mathbf{G}^{-1}, \mathbf{P}) \leq \Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{P}) + \Delta_k^{\mathsf{nCPA}}(\mathbf{G}, \mathbf{P}).$$

With this and Proposition 1 we directly get that six rounds with nCPA secure QRFs give a CCA secure QRP, i.e.

$$\Delta_k^{\mathsf{CCA}}(\psi_{2n}[\mathbf{FFFFFF}], \mathbf{P}) \leq 6 \cdot \Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{R}) + \frac{k^2}{2^{n-1}}.$$

So six nCPA secure round functions are sufficient to get CCA security, and by Proposition 4 we know that at least four rounds are necessary. Using Proposition 5 we can further relax the requirements for the round functions as

$$\begin{aligned}
&\Delta_k^{\mathsf{CCA}}(\mathsf{H} \triangleright \psi_{2n}[\mathbf{FGGF}] \triangleright \mathsf{H}^{-1}, \mathbf{P}) \\
&\leq 2 \cdot \Delta_k^{\mathsf{nCPA}}(\mathbf{F}, \mathbf{R}) + 4 \cdot \Delta_k^{\mathsf{KPA}}(\mathbf{G}, \mathbf{R}) + 2 \cdot \mathsf{coll}_k({}_LH) + 2 \cdot \frac{k^2}{2^{n+1}}.
\end{aligned}$$

As to the (in)security of the Feistel-network with nCPA secure round-functions in the computational setting, we do not know anything beyond what is already implied by CPA security alone, i.e. four rounds are not enough to get CCA security (as it is not enough to get CPA security by Theorem 2).

References

[Dam04] Ivan Damgård. Discrete log based cryptosystems, 2004. Manuscript, `www.daimi.au.dk/ivan/DL.pdf`.

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.

[LR86] Michael Luby and Charles Rackoff. Pseudo-random permutation generators and cryptographic composition. In *Proc, 18th ACM Symposium on the Theory of Computing (STOC)*, pages 356–363, 1986.

[Luc96] Stefan Lucks. Faster Luby-Rackoff ciphers. In *Fast Software Encryption*, volume 3557 of *LNCS*, pages 189–203. Springer-Verlag, 1996.

[Mau02] Ueli Maurer. Indistinguishability of random systems. In *Advances in Cryptology — EUROCRYPT '02*, volume 2332 of *LNCS*, pages 110–132. Springer-Verlag, 2002.

[MOPS06] For the full version of this paper see `www.crypto.ethz.ch/publications`

[MP04] Ueli Maurer and Krzysztof Pietrzak. Composition of random systems: When two weak make one strong. In *Theory of Cryptograpy — TCC '04*, volume 2951 of *LNCS*, pages 410–427. Springer-Verlag, 2004.

[MPR06] Ueli Maurer, Krzysztof Pietrzak, and Renato Renner. Indistinguishability amplification, 2006. Manuscript.

[MT05] Kazuhiko Minematsu and Yukiyasu Tsunoo. Hybrid symmetric encryption using known-plaintext attack-secure components. In *ICISC '05*, *LNCS*. Springer-Verlag, 2005.

[Mye04] Steven Myers. Black-box composition does not imply adaptive security. In *Advances in Cryptology — EUROCRYPT '04*, volume 3027 of *LNCS*, pages 189–206. Springer-Verlag, 2004.

[NR99] Moni Naor and Omer Reingold. On the construction of pseudorandom permutations: Luby-Rackoff revisited. *J. Cryptology*, 12(1):29–66, 1999.

[NR02] Moni Naor and Omer Reingold. Constructing pseudo-random permutations with a prescribed structure. *J. Cryptology*, 15(2):97–102, 2002.

[Pat04] Jacques Patarin. Security of random feistel schemes with 5 or more rounds. In *Advances of Cryptology — CRYPTO '04*, volume 3152 of *LNCS*, pages 106–122. Springer-Verlag, 2004.

[Pie90] Josef Pieprzyk. How to construct pseudorandom permutations from single pseudorandom functions. In *Advances in Cryptology — EUROCRYPT '90*, volume 537 of *LNCS*, pages 140–150. Springer-Verlag, 1990.

[Pie05] Krzysztof Pietrzak. Composition does not imply adaptive security. In *Advances in Cryptology — CRYPTO '05*, volume 3621 of *LNCS*, pages 55–65. Springer-Verlag, 2005.

[Pie06] Krzysztof Pietrzak. Composition implies adaptive security in minicrypt. In *Advances in Cryptology — EUROCRYPT '06*, *LNCS*. Springer-Verlag, 2006.

[Ple05] Patrick Pletscher. Adaptive security of composition, 2005. Semester Thesis. `www.pletscher.org/eth/minor/adapt_sec.pdf`

[RR00] Zulfikar Ramzan and Leonid Reyzin. On the round security of symmetric-key cryptographic primitives. In *Advances in Cryptology — CRYPTO '00*, volume 1880 of *LNCS*, pages 376–393. Springer-Verlag, 2000.