On the Decorrelated Fast Cipher (DFC) and Its Theory

Lars R. Knudsen and Vincent Rijmen *

Department of Informatics, University of Bergen, N-5020 Bergen

Abstract. In the first part of this paper the decorrelation theory of Vaudenay is analysed. It is shown that the theory behind the proposed constructions does not guarantee security against state-of-the-art differential attacks. In the second part of this paper the proposed Decorrelated Fast Cipher (DFC), a candidate for the Advanced Encryption Standard, is analysed. It is argued that the cipher does not obtain provable security against a differential attack. Also, an attack on DFC reduced to 6 rounds is given.

1 Introduction

In [6,7] a new theory for the construction of secret-key block ciphers is given. The notion of decorrelation to the order d is defined. Let C be a block cipher with block size m and C^* be a randomly chosen permutation in the same message space. If C has a d-wise decorrelation equal to that of C^*, then an attacker who knows at most $d-1$ pairs of plaintexts and ciphertexts cannot distinguish between C and C^*. So, the cipher C is "secure if we use it only $d-1$ times" [7]. It is further noted that a d-wise decorrelated cipher for $d = 2$ is secure against both a basic linear and a basic differential attack. For the latter, this basic attack is as follows. A priori, two values a and b are fixed. Pick two plaintexts of difference a and get the corresponding ciphertexts. Repeat a number of times. The attack is successful if and only if at least one ciphertext pair with difference b can be found in a number of tries that is significantly less than 2^m. Let $P(a,b) = \Pr(C(X \oplus a) = C(X) \oplus b)$ denote the probability of the differential with plaintext difference a and ciphertext difference b, where the probability is taken over all plaintexts X. To measure the security of the constructions against the basic differential attack the probabilities of the differentials are averaged over all keys, denoted $E(P(a,b))$. It is then argued that if $E(P(a,b))$ can be upper bounded sufficiently low for all values of a and b, e.g., $E(P(a,b)) \approx 2^{-m}$, then the differential attack will not succeed.

Also, in [7] two families of ciphers are proposed both with the above proofs of security against the basic attacks.

* F.W.O. postdoctoral researcher, sponsored by the Fund for Scientific Research, Flanders (Belgium).

L. Knudsen (Ed.): FSE'99, LNCS 1636, pp. 81–94, 1999.

The Families of Ciphers

COCONUT: This is a family of ciphers parameterised by $(p(x), m)$, where m is the block size and $p(x)$ is an irreducible polynomial of degree m in $GF(2)[x]$. A COCONUT cipher is a product cipher $C_3 \circ C_2 \circ C_1$, where C_1 and C_3 are "any (possibly weak) ciphers" [7], and C_2 is defined

$$C_2(y) = Ay + B \bmod p(x),$$

where A, B and y are polynomials of degree at most $m-1$ in $GF(2)[x]$. The polynomials A and B are secret and act as round keys. Since the COCONUT family has "perfect decorrelation" to the order two it is claimed that the ciphers are secure against the linear and differential attacks.

PEANUT: This is a family of Feistel ciphers parameterised by (m, r, d, p), where m is the block size (in bits), r is the number of rounds, d is the order of the (partial) decorrelation, and p a prime greater than $2^{m/2}$. The round function F takes a text string and d subkeys each of length $m/2$,

$$F(x) = g((k_1 \cdot x^{d-1} + k_2 \cdot x^{d-2} + \ldots + k_{d-1} \cdot x + k_d \bmod p) \bmod 2^{m/2}),$$

where g is any permutation on $m/2$ bits. The DFC is a member of this family (cf. Section 3).

The PEANUT family does not have perfect decorrelation like the COCONUT family. This is due to both the use of the Feistel structure and to the round functions, which are not perfect decorrelated. The multiplications mod p and mod $2^{m/2}$ were chosen since they allow for more efficient implementations in software as compared to multiplication in $GF(2^n)$. The price to pay is that this leads to only partial decorrelated functions. However for sufficiently large values of r it is shown that the ciphers are secure against the linear and differential attacks [7].

In the first part of the paper it is shown that the above constructions based on the decorrelation theory do not necessarily result in ciphers secure against state-of-the-art differential attacks. Example ciphers from both families are shown to be weak. In the second part of this paper we analyse the Decorrelated Fast Cipher (DFC), which was submitted as a candidate for the Advanced Encryption Standard (AES). DFC is an 8-round Feistel cipher and member of the PEANUT family. It is shown that for any fixed key, there exist very high probability differentials for the round function. Also, a differential attack is given on DFC reduced to 6 rounds.

2 Analysis of the Constructions

In this section it will be shown that the constructions in the previous section will not resist differential attacks, thereby indicating a weakness of the decorrelation theory [7].

When analyzing the resistance of a cipher against differential attacks, one often computes the probabilities of differentials over all plaintexts and all keys [4]. (Also, one distinguishes between characteristics and differentials; we use the latter name for both concepts.) For one particular class of iterated ciphers, the Markov ciphers, the probabilities of r-round differentials can be computed as the product of the probabilities of the involved r one-round differentials under the assumption of independent round keys. Moreover, the probabilities are taken only over all possible round keys. However, in an attack the encrypted texts are typically encrypted under a fixed, but secret, key. To deal with this, one assumes that the *hypothesis of stochastic equivalence* holds.

Hypothesis 1 (Hypothesis of stochastic equivalence [4]) *For virtually all high-probability differentials it holds for a substantial fraction of the keys that the probability of the differential for the used key is approximately equal to the average probability of the differential, when averaged over all keys.*

The main reason for the criticism of the constructions based on the decorrelation theory, is that this hypothesis does not hold for the case of the decorrelation modules $k_1x + k_2$ in $\mathrm{GF}(2^m)$ nor for multiplication modulo p modulo $2^{m/2}$ for prime p.

It is shown in the following that the distributions of differences through the "decorrelation modules", $k_1x + k_2$, are very key-dependent. When considering multiplication in the field $\mathrm{GF}(2^m)$ with exclusive-or as the difference operation, for any given input difference $a \neq 0$ and output difference b, the probability of the differential $P(a, b)$ (notation from previous section) for a fixed key, is either 0 or 1. To see this, let x and $x + a$ be two inputs to the module. The difference in the outputs then is, $k_1x + k_2 + k_1(x + a) + k_2 = ak_1$. So, although $E(P(a, b))$ (the average probability taken over all values of the key) can be upper bounded sufficiently low, in an attack one fixed key is used, and differentials of probability 0 and 1 can be found and exploited.

Note that the proof of security against the basic differential attack of the introduction is not affected by these observations. Assume that $P(a, b) \approx 2^{-m}$ for an m-bit block cipher (notation as in the introduction). If the attacker is restricted to choose the values in the differentials before analysing the received ciphertexts the proof of security holds. However, this is not a realistic restriction in our opinion. If for every fixed key there are high probability differentials, an attacker will be able to detect this in an attack and exploit it.

Consider the COCONUT family. In [7] it is shown that C will be secure against the basic differential attack independently of the choices of the ciphers C_1 and C_3. First note that COCONUT versions where $C_1 = C_3 = \mathrm{id}$ (the identity function) have high probability differentials for any fixed key. Also, such ciphers are easily broken using two known plaintexts. One simply solves two equations in two unknowns and retrieves A and B. (This is also noted in [7].) However, we argue based on the above discussion that if a COCONUT cipher is secure against a (state-of-the-art) differential attack for a fixed key then it is because at least one or both of C_1 and C_3 contribute to this security.

In [8] Wagner cryptanalyses COCONUT'98, a member of the COCONUT family. The attack is a differential attack, which exploits that high probability differentials exist for both C_1 and C_3.

Consider next a variant of the PEANUT family of ciphers for $d = 2$, which use multiplication in $GF(2^{m/2})$ in the round function and let g be any affine mapping in $GF(2^{m/2})$. The reason goes along the same lines as the reasoning of the claim for COCONUT. All differentials through the decorrelation modules have probabilities either 0 or 1, and the same holds for differentials through the round function, since g is affine. Consequently, since this holds for all round functions, there are differentials of probability 0 and 1 for the whole cipher.

Consider now the PEANUT family. The multiplications mod p mod $2^{m/2}$ were chosen since they allow for more efficient implementations in software as compared to multiplication in $GF(2^m)$. Consider constructions for $d = 2$ with multiplication defined in GF(p), for prime $p > 2^{m/2}$, where the Feistel round function is

$$F(x) = g((k_1 \cdot x + k_2 \bmod p) \bmod 2^{m/2})$$

for any permutation g. Let first g be the identity function and let $p = 2^{m/2} + t$, where t is small. Let the difference between two $m/2$-bit texts, x_1 and x_2, be defined as $d(x_1, x_2) = x_1 - x_2 \bmod p$ (subtraction modulo p). In the following it is examined how such a difference distributes through F. First, note that for randomly chosen y, where $0 \leq y < p$, it holds that $(y \bmod p) \bmod 2^{m/2} = y \bmod p$ with probability $p_1 = 2^{m/2}/(2^{m/2} + t) \approx 1$. So,

$$d(F(x_1), F(x_2)) = d(k_1 \cdot x_1 + k_2 \bmod p, k_1 \cdot x_2 + k_2 \bmod p)$$

with probability at least $(p_1)^2$. But since the multiplication modulo p is linear with respect to the defined difference, one gets that

$$d(F(x_1), F(x_2)) = k_1(x_1 - x_2) \bmod p$$

with probability at least $(p_1)^2$. The halves in the Feistel cipher are combined using the exclusive-or operation, however it is also noted in [7, Th. 9] that the proof of security for the construction remains valid if the group operation is replaced by any other group operation. Assume therefore that the halves are combined using addition modulo $2^{m/2}$. Let w_1 and w_2 be the two $m/2$-bit halves from a previous round which are combined with the outputs $F(x_1)$ and $F(x_2)$ of the current round. Assume $d(w_1, w_2) = \beta$. Then with probability 1/2, $w_1 + F(x_1) \bmod 2^{m/2} = w_1 + F(x_1)$ in Z, thus if $d(F(x_1), F(x_2)) = \alpha$, then

$$d(F(x_1) + w_1 \bmod 2^{m/2}, F(x_2) + w_2 \bmod 2^{m/2}) = \alpha + \beta$$

with probability 1/4.

To sum up, differences modulo p in the round functions of PEANUT distribute non-uniformly. For any fixed round key, a given difference in the inputs to F results in differences in the outputs of F with very high probabilities. Above it was assumed that g was the identity function. The point which is made here is that if members of the PEANUT family are secure against differential attacks,

then it is because g resists the differential attacks, and not because of the decorrelation module by themselves. In the next section a particular member of the PEANUT family is analysed, where the function g is not the identity function.

Note that the high probability differentials described in this section are key-dependent and therefore unknown to an attacker. However, the fact that the key is fixed in an attack means, that the high probability differentials will occur. This can be exploited in a standard attack, e.g., if the attacker guesses the first-round key and/or the last-round key, the keys which produce high probability differentials for the reduced cipher will be good candidates for the correct value of the key(s). Furthermore, also differentials with probability significantly below 2^{-m} can be used in a differential attack. This is illustrated by the attack we present in Section 5.

3 The Decorrelated Fast Cipher

The Decorrelated Fast Cipher (DFC) [2] has been submitted as a candidate for the AES encryption standard [5]. DFC is a member of the PEANUT family, described above. In the following a more precise definition of the DFC is given. For a complete description of DFC the reader is referred to [2].

3.1 General Structure

DFC is a block cipher with the classical Feistel structure. It uses 8 rounds to transform a 128-bit plaintext block into a 128-bit ciphertext block, under the influence of a key that can have a length up to 256 bits. The user key is expanded to 8 128-bit round keys K_i. Every round key is split into two 64-bit halves, denoted A_i and B_i. In every round, the round function uses the right half of the text input and the two 64-bit round key halves to produce a 64-bit output. This output is exored with the left half of the text input. Subsequently, both halves are swapped, except in the last round.

3.2 The Round Function

Let X denote the 64-bit text input. First a modular multiplication is performed, followed by an additional reduction.

$$Z = \left(A_i \cdot X + B_i \bmod (2^{64} + 13)\right) \bmod 2^{64} \tag{1}$$

Subsequently, the 'confusion permutation' is applied to Z: the value is split into two 32-bit halves, denoted Z_l and Z_r. Z_l is exored with a constant KC. Z_r is exored with a table entry that is determined by the 6 least significant bits of the original Z_l. Both halves are swapped, and the result is added with a 64-bit constant KD.

$$Y = ((Z_r \oplus RT[Z_l \bmod 64]) \ll 32) + (Z_l \oplus KC) + KD \bmod 2^{64}$$

The result Y is the output of the round function.

3.3 The Key Scheduling

The key scheduling first pads the user key with a constant string until a 256-bit string K is obtained. Subsequently, K is divided into two 128-bit parts K_1 and K_2. The keys K_1 and K_2 define each an invertible transformation, denoted $E_1()$ and $E_2()$ respectively. Let RK_O denote a string of 128 zero bits. Then the round keys of DFC are defined as follows.

$$RK_i = E_1(RK_{i-1}) \text{ if } i \text{ is odd} \tag{2}$$

$$RK_i = E_2(RK_{i-1}) \text{ if } i \text{ is even} \tag{3}$$

4 The Distribution of Differences in DFC

First note that since DFC is a member of the PEANUT family, versions of DFC which use only the decorrelation modules in the round function have very high probability differentials. However, the round function of DFC is more than that. To measure the distribution of differences through the round function of DFC we first consider a simplified version. First change all exors to additions modulo 2^{64} and remove the nonlinear S-box RT. This version is hereafter called DFC'. The swapping of the 32-bit halves inside the F-function is retained. Note that the proof of security for DFC' is the same as for DFC. Consider one round of DFC'. Define the difference of two 64-bit texts as the subtraction modulo p. The following test was implemented. Randomly choose a difference (α_L, α_R), where both α's are 64 bits, in the inputs to one round. Randomly choose a pair w_1, w_2 of texts with difference α_L in the left halves. Randomly choose a pair of round keys. For n random choices of x_1 compute the differences of the outputs y_1 and y_2 of the function F for inputs x_1 and $x_2 = (x_1 - \alpha_R) \bmod p$. Compute and store the differences of $y_1 + w_1 \bmod 2^{64}$ and $y_2 + w_2 \bmod 2^{64}$. Since modulo 2^{64} operations used to combine the halves in DFC' are not completely compatible with modulo $p = 2^{64} + 13$ operations, differentials are examined for the addition of the Feistel cipher halves in addition to the round function F.

It is infeasible to do tests for all 2^{64} inputs, but as we will see, this is not necessary in order to determine the distribution of the differences. In 10 tests with $n = 10,000$ input pairs, the number of possible differences in the outputs and the probabilities of the highest one-round differential were recorded. The 10,000 pairs of inputs lead to only an average of 13.6 possible output differences. The average probability of the best one-round differential was 3/8. In similar tests with 1,000,000 pairs, the average number of possible output differences was 14.0 still with an average probability of 3/8. Thus it can be expected that the corresponding numbers for all possible inputs are close to these estimates. Note also, these tests were performed for one randomly chosen input difference, thus, by considering many (all) possible input differences higher probability differentials can be expected.

Thus, despite the fact that the round function is almost perfectly decorrelated, very high probability differentials exist for any fixed key.

Table 1. Results of the experiments on simplified versions of DFC. Probabilities of best differentials for a randomly chosen input difference in 10 tests with randomly chosen keys. (*) Average in 10 tests of best differential for 100 randomly chosen input differences.

	max. probability	# output diff.
DFC'	3/8	14
DFC"	1/128	808
DFC" (*)	0.37	370

Consider next a version of DFC where all exors are replaced by additions modulo 2^{64}, but where RT is unchanged, hereafter called DFC". Note that the proof of security for DFC" is the same as for DFC.

In 10 tests similar to the above with $n = 10,000$ input pairs, the number of possible differences in the outputs was an average of 715 and the probabilities of the highest one-round differential were 1/100 on the average. In similar tests with 1,000,000 pairs, the average number of possible output differences was 808 with an average probability of 1/128. This is no surprise, since when the 6 bits input to RT in a differential are different, the outputs of the round will look random in one half. Since these 6 bits are equal with probability 1/64, these results are in correspondence with the test results on DFC'. Moreover, for a fixed key there are input differences such that the 6 bits input to RT are equal in more than the average case, and the probability of the differential will be higher. To test this phenomenon, we implemented some further tests. In 10 tests a randomly chosen key was used. In each tests for each of 100 randomly chosen input differences, 100,000 input pairs were generated and the output differences recorded. The probabilities of the best such differentials for the 10 keys ranged from 1/22 to 3/5 with an average of 0.37 and 370 possible output differences. Table 1 summarizes the results of the experiments.

Since the only difference between the round functions of DFC" and DFC is the use of three additions mod 2^{64} *instead of three exors, it has been clearly demonstrated that if DFC for a fixed key is secure against differential attacks it is because of the use of mixed group operations and not because of the decorrelation modules.*

Estimating the uniformity of differences and computing the probabilities of differentials are much harder for real DFC. To get an indication of such results, a version of DFC with 32-bit blocks was implemented, hereafter denoted DFC_{32}. The round function takes as input a 16-bit block, uses multiplication modulo the prime $p = 2^{16} + 3$ followed by a reduction modulo 2^{16}. The RT-table has 16 entries (the size of the table is chosen as the size of the inputs (in bits) to the round function, in the spirit of DFC) with randomly chosen values, and the constants KC and KD were chosen at random.

In 100 tests, the number of possible differences in the outputs and the probabilities of the highest one-round differential were recorded for one randomly chosen input difference and for all 2^{16} inputs.

Table 2. Results of the experiments on a scaled-down version of DFC. Probabilities of best differentials for a randomly chosen input difference in 100 tests with randomly chosen keys. (*) Average in 100 tests of best differential for 100 randomly chosen input differences.

	max. probability	# output diff.
DFC_{32}	1/397	6700
DFC_{32} (*)	1/91	1750

The 2^{16} pairs of inputs lead to an average of 6700 possible output differences. The average probability of the best one-round differential was 1/397 (and 1/21 in the best case). By considering 100 input differences for every chosen key, the number of possible output differences dropped to 1750, and the average best probability increased to 1/91 (and 1/18 in the best case).

Table 2 summarizes the results of the experiments.

It can be argued that the RT-table chosen in this scaled-down version of DFC is too big relatively to DFC. Repeating the last test above, this time with a 2-bit RT-table, the number of possible output differences dropped to 1051, and the average best probability increased to 1/49 (and 1/7 in the best case).

Based on the tests conducted here, it is hard to estimate the exact effect for the proposed DFC (without any modifications). However, the tests strongly indicate that the round function of DFC distributes differences modulo p in a very non-uniform way, and that high probability differentials exist.

Summarizing, it was demonstrated that if the DFC is secure against the differential attacks it will be because of the elements that are independent of the proof of security. Also, it was clearly indicated that high probability differentials will exist for DFC for any fixed key.

5 A Differential Attack

The high probability differentials of the previous section might lead to a straightforward differential attack on DFC. However, the large block size of DFC makes it hard to perform such tests. It is left as an open problem for the time being.

In the following we present an attack on DFC when reduced to six of the proposed eight rounds. The attack does not depend directly on the findings in the previous version, but these are incorporated in a possible improvement described later. The attack uses a differential with S/N-ratio < 1. As explained in [1] and [3], this type of differentials can be used in a similar way as 'ordinary' differentials with S/N-ratio > 1 to mount a differential attack. Before the attack is explained, we mention a property of the DFC key schedule that is useful in the attack.

5.1 A Key Scheduling Weakness

The first round key is defined as $RK_1 = E_1(RK_0)$. The string RK_0 is constant and the transformation $E_1()$ depends on one half of the key bits. The consequence

is that the entropy of the first round key is at most one half of the entropy of the user key, e.g., for a 128-bit user key, the first round key has only an entropy of 64 bits. This property makes it easier for an attacker to bypass the first round by guessing the key.

5.2 The *F*-Function Is Almost Bijective

The F-function of DFC is almost bijective. The only non-invertible part of the F-function is the reduction modulo 2^{64} after the modular multiplication in (1). Let x_1, x_2 be two different inputs and let $y_i = A \cdot x_i + B$, where A and B are the secret keys. The inputs x_1, x_2 will be mapped to the same output if and only if

$$(y_1 \bmod (2^{64}+13)) \bmod 2^{64} = (y_2 \bmod (2^{64}+13)) \bmod 2^{64}.$$

If $x_1 \neq x_2$, the equality can only hold if either $y_1 \bmod (2^{64}+13) \in \{0, 1, \ldots, 12\}$ and $y_2 = y_1 + 2^{64}$, or $y_1 \bmod (2^{64}+13) \in \{2^{64}, 2^{64}+1, \ldots, 2^{64}+12\}$ and $y_2 = y_1 - 2^{64}$. For fixed values of A and B, there can be at most 26 tuples (x_1, x_2) with $0 \leq x_1, x_2 < 2^{64}$, that result in equal output values.

It follows that for any key $K = (A, B)$

$$\sum_{\alpha \neq 0} P(\alpha \to 0|K) \leq 26 \cdot 2^{-64}.$$

Because for every value of α, $P(\alpha \to 0|K)$ is a multiple of 2^{-64}, there are for every round key value K at most 26 α's such that the probability is larger than zero.

5.3 A 5-Round Differential with Low Probability

Consider the 5-round differential with both input difference and output difference equal to $(\alpha, 0)$, where α is an arbitrary value, different from zero. (We use the bitwise exor as difference operation.) In this section we will try to give an upper bound for the probability of this differential. In order for our attack to work, this upper bound should be significantly smaller than 2^{-64}.

On Figure 1 it is easy to see that a pair that follows the differential, will have a difference of $(0, \alpha)$ at the input of the second round, and $(\alpha, 0)$ at the output of the fourth round. In the second round, the input difference to the F-function will lead to a certain output difference, denoted β. Similarly, reasoning backwards, it follows that in the fourth round, the difference at the input of the F-function equals α. The output difference is denoted γ. It follows that the third round will have input difference (α, β) and output difference (γ, α). This requires that $\beta \equiv \gamma$ and that the output difference of the F-function in the third round is zero and the input difference β.

Note that the differential does not specify any particular value of β. The probability of the differential is thus given by the sum over all β-values of the probabilities of the characteristics.

$$P_{\text{dif}} = \sum_{\beta} P_{\text{char}(\beta)}$$

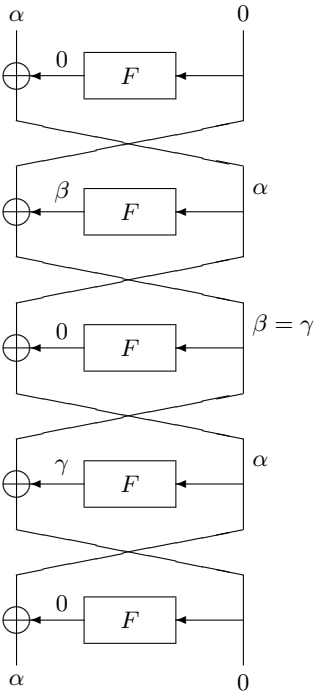

Fig. 1. A 5-round differential with very low probability.

We will approximate the probabilities of the characteristics by the product of the probabilities of the composing one-round characteristics. This may in fact cause inaccuracies, but it is a common assumption. The exact probability of the differential is not very important for our attack, and the experimental results confirm our analysis. We feel that a more exact calculation of the probability would needlessly complicate the analysis.

As already explained in Section 2 and Section 4, the probability of a one round characteristic depends heavily on the value of the round key.

$$P_{\text{dif}} \approx \sum_{\beta} P(\alpha \to \beta \mid K = K_2) P(\beta \to 0 \mid K = K_3) P(\alpha \to \beta \mid K = K_4) \quad (4)$$

When calculating the probability of a characteristic, a distinction is made between the cases $\beta = 0$ and $\beta \neq 0$. If $\beta = 0$

$$P_{\text{char}(\beta=0)} = P(\alpha \to 0 \mid K = K_2) P(\alpha \to 0 \mid K = K_4)$$

Under the assumption of independent rounds, it follows from Section 5.2 that $P_{\text{char}(\beta=0)} \leq (26/2^{64})^2 < 2^{-118}$.

If $\beta \neq 0$

$$P_{\text{char}(\beta)} = P(\alpha \to \beta \mid K = K_2) P(\beta \to 0 \mid K = K_3) P(\alpha \to \beta \mid K = K_4).$$

It follows from Section 5.2 that for every value of K_3 there are at most 26 β-values such that the probability is larger than zero. Also, from Section 4 it follows that for every value of the round key, there are (α, β) tuples such that $P(\alpha \rightarrow \beta)$ is relatively "high". Let us denote this probability by p_1. For most other values of (α, β), the probability is much lower than p_1, possibly zero. We denote this probability by p_2. The values of (α, β) which correspond to high and low probabilities depend on the value of the round keys. The worst scenario for the attack would be that the round key values are selected such that there are values of α and β with high $P(\alpha \rightarrow \beta)$ in both the second and the fourth round, and where also $P(\beta \rightarrow 0)$ is nonzero in the third round. The attack uses many different α values, therefore it can be expected that some of the α values will give a high $P(\alpha \rightarrow \beta)$ in the second *or* the fourth round, where β has nonzero $P(\beta \rightarrow 0)$ in the third round. However, it can be argued that for almost all keys it is highly unlikely that there will exist an α such that $P(\alpha \rightarrow \beta)$ is high in the second *and* the fourth round for a suitable β. For almost all keys, the probability of the differential will be at most $26 \cdot 2^{-64} \cdot p_1 \cdot p_2$ for all values of α. It is confirmed by the experiments performed, that this probability is sufficiently low for the attack to work.

5.4 The Actual Attack

The attack on 6 rounds works in almost the same way as the attack on 6-round DEAL [3]. The main differences are the following:

1. The probability of the 5-round differential is not 0, but very low.
2. The attack uses chosen ciphertexts instead of chosen plaintexts. This reason for this is that if the user key length is less than 256 bits, the first-round key of DFC has a lower entropy than the last-round key. It is therefore easier to recover.

The attack starts as follows. Choose 2^{64} ciphertexts with a fixed right half and a variable left half, say $C_i = (X_i, R)$. Obtain the corresponding plaintexts, say $P_i = (Y_i, Z_i)$. Compute $X_i \oplus Z_i$ and find matches $X_i \oplus Z_i = X_j \oplus Z_j$. About 2^{63} matches can be expected. Let $\alpha = X_i \oplus X_j = Z_i \oplus Z_j$. Guess a value for the first-round key. For all the matching pairs, encrypt the plaintexts one round. If the difference after the first round is $(\alpha, 0)$, the guessed key value is wrong with high probability. For the correct value of the first-round key, in some rare cases the probability of getting right pairs might be relatively high, but as explained earlier in by far the most cases this ratio is very low. Assuming that a wrong key produces uniformly distributed output differences, the difference $(\alpha, 0)$ will occur with probability 2^{-64} for each analysed pair. Thus, $2^{-64} \cdot 2^{63} = 0.5$ good pairs can be expected. Discarding all the key guesses that produce a good pair will eliminate about half of the wrong key values. Repeating this attack 64 times eliminates almost all the wrong key values. The attack requires about $64 \cdot 2^{64} = 2^{70}$ chosen ciphertexts and $(2^{64} + 2^{63} + \ldots + 2 + 1) \cdot 2^{64} \approx 2^{129}$ evaluations of the round function, which is roughly equivalent to 2^{126} encryptions.

Table 3. Experimental results for DFC_{32}, reduced to 6 rounds. All results are averages over 10 tests. For every structure, 2^{16} ciphertexts with a fixed right half are decrypted. The last column lists the average number of surviving wrong key guesses after all the runs. (The attack starts with 2^{16} possible values for the key.)

# structures	# wrong keys remaining
16	22.4
17	14.3
18	8.1
19	4.3
20	2.1
21	1.2

Table 4. Adjusted chosen text requirements and work load of the attack on 6 rounds of DFC.

key length	# chosen texts	work load
128	2^{71}	2^{127}
192	$1.5 \cdot 2^{71}$	2^{158}
256	2^{72}	2^{190}

5.5 Implementation

We implemented this attack on DFC_{32}, a reduced version of DFC that operates on blocks of 32 bits. All constants and tables were reduced accordingly and the prime $2^{16} + 3$ was chosen for the modular multiplication. Because of the key scheduling weakness (cf. Section 5.1), the first-round key of DFC_{32} has 16 bits entropy. The results of the previous section predict that 16 structures of 2^{16} texts should allow to determine uniquely the 16-bit key. The results of 10 tests are given in Table 3.

After repeating the basic attack sufficiently many times only a few candidates are left for the secret key. The correct value of the secret key never resulted in the 5-round differential as described above, which justifies the approximations made in Section 5.3. The experiments suggest that in practice a little more chosen plaintexts are required then predicted by the theory, because we get less good pairs than expected for the wrong key guesses. The net result is that every structure eliminates only 39% of the remaining key candidates, instead of the expected 50%. We therefore have to adjust our estimates for the plaintext requirements and the work load of our attack on 6 rounds of DFC. Increasing the plaintext requirements and the work load with a factor two is more than sufficient. The results are given in Table 4, together with the estimates for attacks on DFC with other key lengths. The estimates for the work load use one encryption as the unit operation. The estimate for 256 bit keys is pessimistic, because in that case it is easy to speed up significantly the round function evaluations since the modular multiplication does not have to be repeated for every new key guess.

5.6 A Possible Improvement

In this section a possible improvement of the attack is outlined which may reduce the complexity. The attack uses the observations in Section 4.

Consider again the 5-round differential of the previous section, but allow now the nonzero half of the difference in the last round to be different from the nonzero half of the difference in the first round. The differential starts with a difference $(\alpha, 0)$. After the first round, the difference is $(0, \alpha)$. In the second round, the inputs to F with difference α leads to an output difference β. At the output of the 5th round, a difference $(\delta, 0)$ is required. In the fourth round the inputs to F with difference δ leads to an output difference γ. It follows that $\beta \equiv \gamma$. The attack works as follows. Choose a structure of ciphertexts with equal right halves. For all values of the first-round key(s) K_1, find the plaintext pairs which yield equal inputs to F in the second round. For each such pair, one knows the input differences to F in the second and fourth rounds, and that the two output differences of F in the two rounds are equal. Since the number of output differences of F are limited for any given difference in the inputs, cf. Section 4, this gives an attacker information about the relation between one half of each of the keys K_2 and K_4. Note that the distribution of differences through F depends mostly on one half of the round keys. Repeat this procedure a number of times for each value of the first-round key. The value of the first-round key which gives rise to a frequent relation between the keys K_2 and K_4 is expected to be the correct value. It remains an open problem to what extent this variant of the attack will reduce the complexity of the previous attack.

6 Conclusions

We showed that the constructions of ciphers in the COCONUT and PEANUT families are weak against differential attacks. The main observation is that for a fixed key (which is the scenario of an attack) high probability differentials can be found. We analysed one particular member of the PEANUT-family: the DFC. It was shown that variants of DFC with only small modifications and with the same proof of security as the original, are vulnerable to a differential attack. For the proposed DFC it was indicated that differentials with high probabilities exist for the round function. Also, an attack, not directly related to these findings, on the proposed DFC reduced to six of eight rounds was given. Although the attack requires a large running time it is believed that the outlined possible improvement will be faster.

The results in this paper do not contradict the theory of decorrelation [7]. More specifically, and in accordance with [7], ciphers which are d-wise decorrelated are provably secure against the following attacks.

1. Any chosen plaintext attack using at most $d - 1$ plaintexts.
2. If $d \geq 2$, the basic differential attack, where an attacker is restricted to choose the values in the differentials before the attack.
3. If $d \geq 2$, the basic linear attack.

The point we make is that restricting the attacker to such basic attacks does not lead to a strong proof of security. More specifically, we showed how some more advanced differential techniques can be used to attack decorrelated ciphers in general and reduced versions of DFC in particular. Although the decorrelation theory may be a valuable contribution to cryptographic research, it does not guarantee resistance against state-of-the-art differential attacks.

Acknowledgments

The authors thank Serge Vaudenay and the anonymous referees of the programme committee for helpful comments.

References

1. J. Borst, L.R. Knudsen, V. Rijmen, "Two attacks on reduced IDEA," *Advances in Cryptology, Proceedings Eurocrypt '97, LNCS 1233*, W. Fumy, Ed., Springer-Verlag, 1997, pp. 1-13.
2. H. Gilbert, M. Girault, P. Hoogvorst, F. Noilhan, T. Pornin, G. Poupard, J, Stern, S. Vaudenay, "Decorrelated fast cipher: an AES candidate," *Technical report,* available from `http://www/ens.fr/~vaudenay/dfc.html`. Submitted as an AES candidate. See also `http://www.nist.gov/aes/`.
3. L.R. Knudsen. DEAL - a 128-bit block cipher. Technical Report 151, Department of Informatics,University of Bergen, Norway, February 1998. Submitted as an AES candidate. See also `http://www.nist.gov/aes/`.
4. X. Lai, J.L. Massey, and S. Murphy. Markov ciphers and differential cryptanalysis. In D.W. Davies, editor, *Advances in Cryptology - EUROCRYPT'91, LNCS 547*, pages 17–38. Springer Verlag, 1992.
5. NIST's AES homepage, `http://www.nist.gov/aes`.
6. S. Vaudenay, "Feistel ciphers with L_2-decorrelation," *Preproceedings of SAC'98*, August '98, Kingston (Canada).
7. S. Vaudenay. "Provable Security for Block Ciphers by Decorrelation," In *STACS'98, Paris, France, LNCS 1373*, Springer-Verlag, 1998, pp. 249-275.
8. D. Wagner. The boomerang attack. In these proceedings.