

New Attacks on Feistel Structures with Improved Memory Complexities

Itai Dinur¹, Orr Dunkelman^{2,4,*},
Nathan Keller^{3,4,**}, and Adi Shamir⁴

¹ Département d'Informatique, École Normale Supérieure, Paris, France

² Computer Science Department, University of Haifa, Israel

³ Department of Mathematics, Bar-Ilan University, Israel

⁴ Computer Science department, The Weizmann Institute, Rehovot, Israel

Abstract. Feistel structures are an extremely important and extensively researched type of cryptographic schemes. In this paper we describe improved attacks on Feistel structures with more than 4 rounds. We achieve this by a new attack that combines the main benefits of meet-in-the-middle attacks (which can reduce the time complexity by comparing only half blocks in the middle) and dissection attacks (which can reduce the memory complexity but have to guess full blocks in the middle in order to perform independent attacks above and below it). For example, for a 7-round Feistel structure on n -bit inputs with seven independent round keys of $n/2$ bits each, a MITM attack can use $(2^{1.5n}, 2^{1.5n})$ time and memory, while dissection requires $(2^{2n}, 2^n)$ time and memory. Our new attack requires only $(2^{1.5n}, 2^n)$ time and memory, using a few known plaintext/ciphertext pairs. When we are allowed to use more known plaintexts, we develop new techniques which rely on the existence of multicollisions and differential properties deep in the structure in order to further reduce the memory complexity.

Our new attacks are not just theoretical generic constructions — in fact, we can use them to improve the best known attacks on several concrete cryptosystems such as round-reduced CAST-128 (where we reduce the memory complexity from 2^{111} to 2^{64}) and full DEAL-256 (where we reduce the memory complexity from 2^{200} to 2^{144}), without affecting their time and data complexities. An extension of our techniques applies even to some non-Feistel structures — for example, in the case of FOX, we reduce the memory complexity of all the best known attacks by a factor of 2^{16} .

Keywords: Cryptanalysis, block cipher, Feistel structure, dissection, meet-in-the-middle, splice-and-cut, CAST-128, DEAL.

* The second author was supported in part by the Israeli Science Foundation through grant No. 827/12.

** The third author was supported by the Alon Fellowship.




Fig. 1. The i 'th Round of a Feistel Structure

1 Introduction

Feistel structures were first used in the design of DES [18], and had a major influence on the development of both the theory and the practice of cryptography (e.g., in the Luby-Rackoff [17] construction of pseudo random permutations and in the design of numerous block cipher proposals). In this paper we will primarily consider generic Feistel structures whose i -th round is depicted in Figure 1. They divide their n -bit blocks into two equal parts (L_i, R_i) , use independent $n/2$ -bit subkeys in their ℓ rounds, and have round functions F_i over $n/2$ -bit inputs, outputs and subkeys which are perfect in the sense that they cannot be broken with attacks that are faster than exhaustive search. This choice of parameters allows us to consider any two consecutive rounds in a Feistel structure as a single round in a regular (non-Feistel) structure that has n -bit inputs outputs and subkeys. However, when we describe attacks on concrete schemes which have a Feistel structure, we will consider the relevant key and block sizes, and exploit some of the weaknesses of the actual round functions.

A major type of low-data attacks which can be applied to multi-round constructions is the Meet-In-The-Middle (abbreviated as MITM) attack, which was proposed by Diffie and Hellman [8] in 1977 as a method for cryptanalyzing double encryption schemes. It gained additional fame in 1985, when Chaum and Evertse [7] applied it to reduced-round variants of DES [18], and it is now considered as an essential part in any course in cryptanalysis. In the last few years, research of MITM techniques had expanded in diverse directions and numerous new extensions of the basic MITM appeared, including partial matching [4], probabilistic matching [15, 20], bicliques [3], sieve-in-the-middle [6], and many others. One of the recent approaches is the dissection attack, which was introduced by Dinur et al. [9] at CRYPTO 2012. Dissection can solve a wide variety of combinatorial search problems with improved combinations of time and memory complexities. In its cryptanalytic application, dissection significantly improved the time/memory tradeoff achievable by MITM attacks on multiple encryption schemes with more than 3 rounds.

The main difference between these two types of low-data attacks can be described in the following way: In basic MITM the adversary starts from the known plaintexts and ciphertexts at the endpoints, and works from both endpoints towards the middle by guessing some keys and building appropriate lookup tables. The equality of the pairs of values in the middle is used just as a filtering con-

dition to identify the correct keys, and there is no need to know them in order to start the attack. In dissection attacks the adversary starts by guessing the relevant values in the middle, and works from the middle towards the endpoints. In fact, the knowledge of the middle values enables the adversary to break the cryptanalytic problem into two independent smaller problems with new known plaintext and ciphertext pairs at their endpoints, which can be solved recursively by another dissection, or with MITM at the leaves of the recursion tree.

When we compare the two types of attacks on a Feistel structure with an odd number of $\ell = 2r + 1$ rounds, we notice that each one of them offers different advantages and disadvantages. The MITM attack can ignore the middle round by comparing only the $n/2$ -bit half blocks which are not affected by this round in the Feistel structure. This enables the MITM attack to be more time-efficient in the case of Feistel structures, since it does not have to guess the middle subkey. Even though dissection is naturally more efficient than MITM, it loses in its time complexity in this case since it has to guess the full n -bit middle value in order to be able to encrypt and decrypt this guessed value through multiple rounds.

In this paper we present new techniques which enable us to combine the MITM and dissection approaches, along with additional ingredients, such as iterating over values that are not used later in the MITM attack, and using multi-collisions and differential properties in the middle of the Feistel structure. We first consider the case of Feistel structures with an odd number of rounds, and try to reduce the memory complexity of the most time-efficient attacks on them. We show that the memory complexity of the MITM attack on $\ell = 2r + 1$ rounds for $r \geq 3$ can be reduced all the way from $2^{0.5rn}$ to about $2^{\lceil r/2 \rceil 0.5n}$ (like in dissection) without increasing the time complexity, at the expense of increasing the data complexity to about $2^{\lceil \frac{r-3}{r+1} \rceil 0.5n}$ known plaintexts. If no additional plaintexts are allowed, we are still able to reduce the memory, but only to about $2^{\lceil 2r/3 \rceil 0.5n}$. In particular, we can reduce the memory complexity of the standard MITM attack on 7-round Feistel from $2^{1.5n}$ to 2^n with no effect on the data and time complexities.¹

A different goal is to reduce the time complexity of the most memory-efficient nontrivial attacks (in which the memory available to the adversary is restricted to $2^{0.5n}$, which makes it possible to store in memory all the possible values of a single half-block or a single subkey, but not more). Here, we assume that the round functions can be inverted efficiently when their subkey is given. In [9], Dinur et al. considered low-memory dissection attacks on general (non-Feistel) structures, in which the available memory is restricted to 2^n (where n is the length of a single block or a single subkey). They defined the *gain* of a dissection attack of time complexity T over a standard MITM attack with 2^n memory (whose time complexity is $2^{(\ell-1)n}$ for ℓ rounds) by $(\ell - 1) - \log_2(T)/n$. Then, they computed the sequence of round numbers ℓ for which the gain of the best dissection attack increases by one — $\{4, 7, 11, 16, 22, 29, 37, \dots\}$. We use our techniques to compute

¹ We alert the reader that similarly to [9], we concentrate on asymptotic complexity and ignore small polynomial factors in n . Moreover, we treat ℓ as a (possibly big) constant, and therefore disregard all multiplicative factors related to ℓ (and r).

a similar sequence of round numbers ℓ of Feistel structures for which the gain (over MITM, whose time complexity is $2^{(\ell-2)0.5n}$) increases by one — it turns out to be $\{5, 10, 15, 22, 29, 38, 47, \dots\}$, and the asymptotic complexity of an attack on ℓ -round Feistel using a minimal amount of $2^{0.5n}$ memory turns out to be $2^{0.5n(\ell-2-\sqrt{2\ell}+o(\sqrt{\ell}))}$. In particular, we present an attack on 5-round Feistel with time complexity of 2^n and memory complexity of $2^{0.5n}$ (compared to $(2^n, 2^n)$ or $(2^{1.5n}, 2^{0.5n})$ which are the best that can be obtained by previous attacks).

To deal with an even number of rounds without having to guess the extra key, we show that our algorithms can be combined with a recent algorithm presented by Isobe and Shibutani [12] at Asiacrypt 2013. This algorithm extends MITM attacks on Feistel structures by one round, at the expense of increasing the time complexity by $2^{0.25n}$ and using $2^{0.25n}$ chosen plaintexts. As a result, we obtain an attack on a Feistel structures with $\ell = 2r$ rounds (for $r \geq 4$) that requires $2^{(0.5r-0.25)n}$ time, about $2^{(0.5\lceil r/2 \rceil - 0.25)n}$ memory, and $\max\{2^{0.25n}, 2^{\lceil \frac{r-4}{r} \rceil 0.5n}\}$ chosen plaintexts. Alternatively, we can use only $2^{0.25n}$ chosen plaintexts like in [12], with $2^{(0.5r-0.25)n}$ time and about $2^{(\lceil 2(r-1)/3 \rceil 0.5 + 0.25)n}$ memory. In particular, we reduce the memory complexity of Isobe and Shibutani’s attack [12] on 8-round Feistel structures from $2^{1.75n}$ to $2^{1.25n}$, with no effect on the data and time complexities.

In Table 1 we compare the complexity of our attacks with previous results for certain numbers of rounds which have “clean” exponents.

While all the techniques described so far are completely generic, they allow us to significantly improve the best known attacks on several concrete block ciphers. In particular, we reduce the memory complexity of the best known attack on 8-round CAST-128 [1] from 2^{111} to 2^{64} , and the memory complexity of the best known attack on 8-round DEAL [16] with 256-bit keys from 2^{200} to 2^{144} , both without affecting the time and data complexities. It is interesting to note that an extension of our techniques can even be applied to certain non-Feistel cryptosystems, such as FOX [14], in which the best MITM attack uses the *partial matching* technique.

This paper is organized as follows: In Section 2 we describe the improved memory complexities we obtain when we consider the most time-efficient attacks, and in Section 3 we describe the improved time complexities which can be obtained when we consider the most memory-efficient attacks. In Section 4 we sketch the application of our results to concrete block ciphers. We conclude the paper in Section 5.

2 Improving the Memory Complexity of the most Time-Efficient Attacks on Feistel Structures

Consider a standard Feistel structure with an odd number $\ell = 2r + 1$ of rounds. The generic dissection attack on this construction (that does not exploit the

² In the case of 5-round Feistel, the dissection attack is not better than the meet in the middle attack.

Rounds	Complexity			Attack
	Time	Memory	Data	
5	2^n	2^n	3 KP	Meet in the Middle ²
	$2^{1.5n}$	$2^{0.5n}$	3 KP	Meet in the Middle
	2^n	$2^{0.5n}$	3 KP	New (Section 3.1)
7	$2^{1.5n}$	$2^{1.5n}$	4 KP	Meet in the Middle
	2^{2n}	2^n	4 KP	Dissection
	$2^{1.5n}$	2^n	4 KP	New (Section 2.1)
8	2^{2n}	$2^{1.5n}$	4 KP	Meet in the Middle
	2^{2n}	2^n	4 KP	Dissection
	$2^{1.75n}$	$2^{1.75n}$	$2^{0.25n}$ CP	Splice-and-cut [12]
	$2^{1.75n}$	$2^{1.25n}$	$2^{0.25n}$ CP	New (Section 2.3)
15	$2^{3.5n}$	$2^{3.5n}$	8 KP	Meet in the Middle
	$2^{6.5n}$	$2^{0.5n}$	8 KP	Meet in the Middle
	2^{4n}	2^{2n}	8 KP	Dissection
	2^{5n}	$2^{0.5n}$	8 KP	New (Section 3.2)
	$2^{3.5n}$	2^{2n}	$2^{0.25n}$ KP	New (Section 2.2)
31	$2^{7.5n}$	$2^{7.5n}$	16 KP	Meet in the Middle
	$2^{14.5n}$	$2^{0.5n}$	16 KP	Meet in the Middle
	2^{8n}	2^{4n}	16 KP	Dissection
	2^{12n}	$2^{0.5n}$	16 KP	New (Section 3.2)
	$2^{7.5n}$	2^{5n}	16 KP	New (Section 2.1)
	$2^{7.5n}$	2^{4n}	$2^{0.375n}$ KP	New (Section 2.2)
32	2^{8n}	$2^{7.5n}$	16 KP	Meet in the Middle
	2^{15n}	$2^{0.5n}$	16 KP	Meet in the Middle
	2^{8n}	2^{4n}	16 KP	Dissection
	$2^{7.75n}$	$2^{7.75n}$	$2^{0.25n}$ CP	Splice-and-cut [12]
	$2^{12.5n}$	$2^{0.5n}$	16 KP	New (Section 3.2)
	$2^{7.75n}$	$2^{7.25n}$	$2^{0.25n}$ CP	New (Section 2.3)

KP — Known plaintext, CP — Chosen plaintext

Table 1. Comparing and Summarizing Some of our Results

Feistel structure) requires $2^{0.5(r+1)n}$ time and roughly $2^{0.25(r+1)n}$ memory.³ The standard MITM attack can exploit the Feistel structure to reduce the time complexity to $2^{0.5rn}$, at the expense of enlarging the memory complexity to $2^{0.5rn}$. No attacks faster than $2^{0.5rn}$ are known (unless additional assumptions are made on the round functions or on the key schedule) and thus we concentrate in this section on attacks which have this time complexity. Our goal is to combine the

³ The precise generic formula for the memory complexity of dissection is less relevant here. Note that such an attack has to treat every two consecutive Feistel rounds as a single round with an n -bit block and an n -bit key, and it is the last “half-round” which makes it suboptimal.

benefits of both MITM and dissection attacks in order to reduce the memory complexity to $2^{0.25(r+1)n}$. We show that this is indeed possible, but at the expense of somewhat enlarging the data complexity of the attack.

First, we present a basic 7-round attack⁴ that requires $2^{1.5n}$ time and 2^n memory (compared to $(2^{1.5n}, 2^{1.5n})$ and $(2^{2n}, 2^n)$ in generic MITM and dissection, respectively), and extend it to an attack on $2r + 1$ rounds that requires $2^{0.5rn}$ time and about $2^{0.33rn}$ memory. Then, we present a more sophisticated attack that requires $2^{0.5rn}$ time and about $2^{0.25rn}$ memory as desired, but at the expense of enlarging the data complexity to $2^{\lceil (r-4)/r \rceil \cdot 0.5n}$ known plaintexts. Finally, we show that our attacks can be combined with a technique of Isobe and Shibutani [12] that allows extending MITM attacks on Feistel structures by one round using a splice-and-cut technique [2]. We obtain an attack on $\ell = 2r$ -round Feistel that requires $2^{(0.5r-0.25)n}$ time, about $2^{0.25(r+1)n}$ memory, and $2^{\max(\lceil \frac{r-4}{r} \rceil, 0.5) \cdot 0.5n}$ chosen plaintexts.

2.1 Attacks with a Low Data Complexity

In this section we present attacks that are time-efficient (i.e., have a time complexity of $2^{0.5rn}$ for $2r + 1$ rounds) and also data-efficient (i.e., require only a few known plaintexts, like the standard MITM attack).

A standard MITM attack In order to put our attacks in context, we begin by describing a standard MITM attack on a 7-round Feistel structure (see left side of Figure 2).

A 7-Round MITM Attack

1. Obtain 4 plaintext-ciphertext pairs (P^i, C^i) ($i = 1, 2, 3, 4$).
2. For each value of K_1, K_2, K_3 :
 - (a) Partially encrypt P^i for $i \in \{1, 2, 3, 4\}$ through the first three rounds and obtain suggestions for R_3^i . Store the suggestions in a list *List* sorted by the R_3^i values.
3. For each value of K_5, K_6, K_7 :
 - (a) Partially decrypt C^i for $i \in \{1, 2, 3, 4\}$ through the last three rounds, obtain suggestions for R_3^i and search the suggestions in *List*. For each match, retrieve K_1, K_2, K_3 , guess⁵ K_4 , and test the full key using trial encryptions.

The time complexity of Step 2 is about $2^{1.5n}$, which is also the size of *List*. In order to calculate the time complexity of Step 3, we note that we have a total of $2^{1.5n}$ key suggestions from each side of the encryption, each associated with 4 values of R_3^i (filtering conditions). Thus, the expected total number of key

⁴ We consider attacks on less than 7 rounds in Section 3.

⁵ We note that K_4 can also be found by a precomputed table instead of guessing, but this will not make a big difference as will be explained in the sequel.

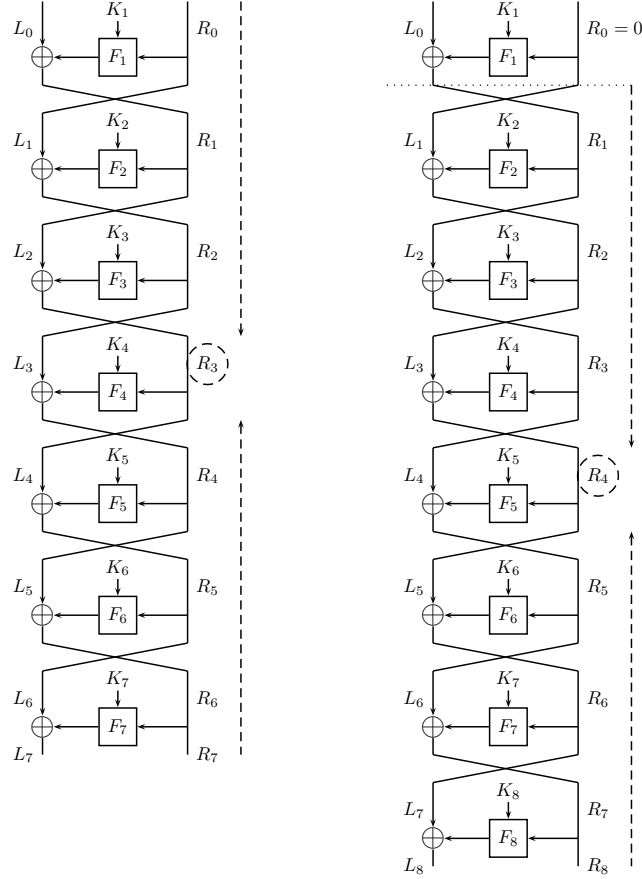


Fig. 2. On the Left: 7-Round MITM Attack. On the Right: 8-Round Splice-and-Cut Attack.

suggestions that remain after the $2n$ -bit match in Step 3 is $2^{1.5n+1.5n-2n} = 2^n$. For each such suggestion, we guess K_4 , and thus we expect to perform about $2^{1.5n}$ trial encryptions. Consequently, the time complexity of Step 3 is also about $2^{1.5n}$, which is the time complexity of the full attack.

A 7-round attack The basic idea of our reduced memory attack is to guess the $n/2$ -bit value R_3^1 and to iterate over all the possible guesses as an outer loop. Each guess imposes an $n/2$ -bit constraint on the key suggestions for K_1, K_2, K_3 and K_5, K_6, K_7 , and thus, allows reducing the expected size of *List* to 2^n . In order to compute the reduced lists efficiently, we prepare auxiliary tables T_{upper}, T_{lower} that allow retrieving the subkey K_3 (resp., K_5) instantly given the input (L_2, R_2) of round 3 (resp., the output (L_5, R_5) of round 5).

The table T_{upper} is computed as follows. We guess the intermediate value R_2^1 and the subkey K_3 . Since $(R_2^1, R_3^1) = (L_3^1, R_3^1)$ form the full state after the 3'rd round in the encryption process of P^1 , the guesses enable us to partially decrypt through round 3 and obtain $(R_1^1, R_2^1) = (L_2^1, R_2^1)$. We store the triple (L_2^1, R_2^1, K_3) in T_{upper} , sorted by (L_2^1, R_2^1) . The table T_{lower} is constructed similarly.

7-Round Attack with Reduced Memory Complexity

1. Obtain 4 plaintext-ciphertext pairs (P^i, C^i) .
2. For each value of R_3^1 :
 - (a) For each value of K_3 and $I_3^1 = R_2^1$, compute $L_2^1 = F_3(K_3, I_3^1) \oplus R_3^1$, and store the triplet (L_2^1, R_2^1, K_3) in a table T_{upper} , sorted according to (L_2^1, R_2^1) .
 - (b) For each value of K_5 and $I_5^1 = R_4^1$, compute $R_5^1 = F_5(K_5, I_5^1) \oplus R_3^1$, and store the triplet (L_5^1, R_5^1, K_5) in a table T_{lower} , sorted according to (L_5^1, R_5^1) .
 - (c) For each value of K_1, K_2 :
 - i. Partially encrypt P^1 through the first two rounds to obtain suggestions for R_2^1 and L_2^1 .
 - ii. Search for the pair (L_2^1, R_2^1) in T_{upper} and obtain suggestions for K_3 . For each suggestion, given K_1, K_2, K_3 , partially encrypt P^i for $i \in \{2, 3, 4\}$ through the first three rounds and obtain suggestions for R_3^i . Store the suggestions in a list $List1$, sorted by the values R_3^2, R_3^3, R_3^4 .
 - (d) For each value of K_6, K_7 :
 - i. Partially decrypt C^1 through the last two rounds to obtain suggestions for R_5^1 and $L_5^1 = R_4^1$.
 - ii. Search for the pair (L_5^1, R_5^1) in T_{lower} and obtain suggestions for K_5 . For each suggestion, given K_5, K_6, K_7 , partially decrypt C^i for $i \in \{2, 3, 4\}$ through the last three rounds, obtain suggestions for R_3^i and search the suggestions in $List1$. For each match, retrieve K_1, K_2, K_3 , guess K_4 , and test the full key using trial encryptions.

In Steps 2a and 2b, a single round function (either F_3 or F_5) is called once for each guess of (R_3^1, I_3, K_3) (or (R_3^1, I_5, K_5) , respectively), and thus, their time complexity is $2^{1.5n}$. The memory complexity of the tables T_{upper} and T_{lower} is 2^n . In Steps 2(c)ii and 2(d)ii there is an average of one match in T_{upper} and T_{lower} , respectively. Thus, on average, we perform a constant number of partial encryption and decryption operations per guess of K_1, K_2 and K_6, K_7 in Steps 2c and 2d, respectively. The expected number of matches in Step 2(d)ii is $2^{n+n-1.5n} = 2^{0.5n}$, and the expected number of trial encryptions (after guessing K_4) is $2^{0.5n+0.5n} = 2^n$. Therefore, the time complexity of each of Steps 2c and 2d is about 2^n each, and the total time complexity of the attack is about $2^{1.5n}$, as in the standard MITM attack. On the other hand, the memory complexity of the attack is reduced from $2^{1.5n}$ to about 2^n , which is the expected number of

elements in *List1* (based on standard randomness assumptions on the round functions).

We note that the time complexity of the attack can be slightly reduced by precomputing a table for F_4 , which allows to avoid guessing K_4 in Step 2(d)ii. However, this requires an additional table of size 2^n , i.e., maintaining the 2^n total memory complexity.

Extension to $6r + 1$ rounds. The 7-round attack presented above can be extended to an attack on a $6r + 1$ -round Feistel structure, with time complexity of $2^{1.5rn}$ (as in standard MITM) and memory complexity of 2^{rn} (instead of $2^{1.5rn}$ in standard MITM). As the attack is similar to the 7-round attack, we describe it briefly. The reader can follow this attack by verifying that the case $r = 1$ reduces exactly to the attack described above.

First, we obtain $3r + 1$ plaintext/ciphertext pairs (P^i, C^i) . Then, the outer loop is performed for all guesses of the r intermediate values $R_{3r}^1, R_{3r}^2, \dots, R_{3r}^r$. In the inner loop, we guess the $2r$ values $R_{3r-1}^1, R_{3r-1}^2, \dots, R_{3r-1}^r, K_{3r}, \dots, K_{2r+1}$. Since for each i , $(R_{3r-1}^i = L_{3r}^i, R_{3r}^i)$ forms the full state after the $3r$ 'th round in the encryption process of P^i , we can partially decrypt this state through rounds $3r, \dots, 2r + 1$ to obtain the corresponding values $(R_{2r-1}^i = L_{2r}^i, R_{2r}^i)$. This allows us to prepare a table T_{upper} of size 2^{rn} of the values $((L_{2r}^i, R_{2r}^i)_{i=1}^r, K_{2r+1}, \dots, K_{3r})$, sorted by $((L_{2r}^i, R_{2r}^i)_{i=1}^r)$. The table T_{lower} is prepared similarly. Note that the $2r$ values guessed from each side are used only for preparing the tables and not in the rest of the attack.

After preparing the tables, we guess the subkeys K_1, K_2, \dots, K_{2r} , obtain the intermediate values $(L_{2r}^i, R_{2r}^i)_{i=1}^r$ and access the table T_{upper} to obtain a suggestion for the subkeys K_{2r+1}, \dots, K_{3r} . For each suggestion, given K_1, K_2, \dots, K_{3r} , we partially encrypt P^i for $i \in \{r+1, \dots, 3r+1\}$ through the first $3r$ rounds and obtain suggestions for R_{3r}^i . We store the suggestions in a list *List1*, sorted by the values $R_{3r}^{r+1}, \dots, R_{3r}^{3r+1}$. Then, we guess the subkeys $K_{6r+1}, \dots, K_{4r+2}$, access the table T_{lower} to obtain a suggestion for $K_{4r+1}, \dots, K_{3r+2}$, partially decrypt the ciphertexts to get suggestions for R_{3r}^i ($i = r + 1, \dots, 3r + 1$), and search them in *List1*. For each match, we retrieve K_1, \dots, K_{3r} , guess K_{3r+1} , and test the full key using trial encryptions.

The analysis of the attack is similar to that of the 7-round attack described above, and yields time complexity of $2^{1.5rn}$ and memory complexity of 2^{rn} . The same attack applies for a general odd number $2r' + 1$ of rounds. The time complexity is $2^{r'n}$ (like in MITM), but the memory complexity has to be rounded up to $2^{\lceil 2r'/3 \rceil \cdot 0.5n}$, due to lack of balance between the part of table creation and the rest of the attack.

2.2 Using Multi-Collisions to Further Reduce the Memory Complexity

We now present a more sophisticated variant of the attacks described above, that allows to reduce the memory complexity to the “desired” $2^{0.25(r+1)n}$ (simi-

larly to dissection on multiple encryption schemes), with no increase in the time complexity, but at the expense of some increase in the data complexity.

Consider the $6r + 1$ -round attack described above. In the course of preparing the table T_{upper} , we make an auxiliary guess of the values $R_{3r-1}^1, \dots, R_{3r-1}^r, K_{3r}, \dots, K_{2r+1}$, and in the course of preparing the table T_{lower} , we guess $R_{3r+1}^1, \dots, R_{3r+1}^r, K_{3r+2}, \dots, K_{4r+1}$. If there was some relation between the guessed values, we could have used this relation to enumerate over some “common relative” in the outer loop of the attack, and thus reduce the memory complexity. We cannot hope for such a relation between the subkeys, as they are assumed to be independent.⁶ However, some relation between R_{3r-1}^i and R_{3r+1}^i may exist.

We observe that such a relation can be “created”, using multi-collisions. Assume that the partial encryption of the plaintexts P^1, P^2, \dots, P^r considered in the attack results in an r -multi-collision at the state R_{3r} , i.e., that $R_{3r}^1 = R_{3r}^2 = \dots = R_{3r}^r$. In such a case, the r values $R_{3r-1}^i \oplus R_{3r+1}^i = O_{3r+1}^i$ ($i = 1, \dots, r$) are all equal! This allows us to enumerate over the $r - 1$ values $R_{3r-1}^i \oplus R_{3r-1}^1$ ($i = 2, \dots, r$) in the outer loop, such that in the inner loop, a single guess of R_{3r-1}^1 provides all the values $\{R_{3r-1}^i\}_{i=2, \dots, r}$, while a single guess of R_{3r+1}^1 provides all the values $\{R_{3r+1}^i\}_{i=2, \dots, r}$. In order to obtain the multi-collision, we consider $2^{((r-1)/r) \cdot 0.5n}$ known plaintexts (which guarantees that an r -multi-collision exists in the data with a constant probability⁷), and repeat the attack for all r -tuples of plaintext/ciphertext pairs in the data set.

In the description of the algorithm below, we switch from $6r + 1$ rounds to $8r - 1$ rounds, in order to balance the complexities of all the steps of the attack. Hence, the external guesses are performed at state R_{4r-1} , instead of R_{3r} . Note that for $r = 1$, the algorithm reduces to the 7-round attack presented above.

An $8r - 1$ -Round Attack Using Multi-Collisions

1. Obtain $2^{((r-1)/r) \cdot 0.5n}$ plaintext-ciphertext pairs (P^i, C^i) .
2. For each r -tuple $(P_{i_1}, C_{i_1}), (P_{i_2}, C_{i_2}), \dots, (P_{i_r}, C_{i_r})$ of plaintext-ciphertext pairs in the data set (hereinafter denoted for simplicity by $(P^1, C^1), \dots, (P^r, C^r)$), for each possible value of R_{4r-1}^1 , and for all possible differences $R_{4r-2}^1 \oplus R_{4r-2}^i$ ($i = 2, 3, \dots, r$):
 - (a) For each $I_{4r-1}^1 = R_{4r-2}^1$ and the subkeys $K_{4r-1}, K_{4r-2}, \dots, K_{2r+1}$, compute⁸ (L_{2r}^i, R_{2r}^i) for all $i = 1, \dots, r$, and store the vector $((L_{2r}^i, R_{2r}^i)_{i=1, \dots, r}, K_{2r+1}, \dots, K_{4r-1})$ in a table T_{upper} , sorted according to $(L_{2r}^i, R_{2r}^i)_{i=1, \dots, r}$.
 - (b) For each $I_{4r+1}^1 = R_{4r}^1$ and the subkeys $K_{4r+1}, \dots, K_{6r-1}$, compute (L_{6r-1}^i, R_{6r-1}^i) for all $i = 1, \dots, r$, and store the vector

⁶ Some useful relations between subkeys may exist if they are derived from a master key using a simple key schedule. However, in this paper, we focus on the most general Feistel constructions and do not assume the existence of such relations.

⁷ Roughly speaking, given D data, we have about D^r different r -tuples of the $0.5n$ -bit value R_{3r} . For each r -tuple, the probability that all the values of R_{3r} are equal is $2^{-0.5n \cdot (r-1)}$. Therefore, we require $D^r = 2^{0.5n \cdot (r-1)}$ or $D = 2^{((r-1)/r) \cdot 0.5n}$.

- $((L_{6r-1}^i, R_{6r-1}^i)_{i=1, \dots, r}, K_{4r+1}, \dots, K_{6r-1})$ in a table T_{lower} , sorted according to $(L_{6r-1}^i, R_{6r-1}^i)_{i=1, \dots, r}$.
- (c) For each value of K_1, K_2, \dots, K_{2r} :
- i. Partially encrypt P^1, \dots, P^r through the first $2r$ rounds to obtain suggestions for $(L_{2r}^i, R_{2r}^i)_{i=1, \dots, r}$.
 - ii. Search for $(L_{2r}^i, R_{2r}^i)_{i=1, \dots, r}$ in T_{upper} and obtain suggestions for $K_{2r+1}, \dots, K_{4r-1}$. For each suggestion, given K_1, \dots, K_{4r-1} , partially encrypt $2r+1$ additional plaintexts P^j for $j \in \{1, \dots, 2r+1\}$ through the first $4r-1$ rounds and obtain suggestions for R_{4r-1}^j . Store the suggestions in a list $List1$, sorted by the values $\{R_{4r-1}^j\}_{j=1, \dots, 2r+1}$.
- (d) For each value of K_{8r-1}, \dots, K_{6r} :
- i. Partially decrypt C^1, \dots, C^r through the last $2r$ rounds to obtain suggestions for $(L_{6r-1}^i, R_{6r-1}^i)_{i=1, \dots, r}$.
 - ii. Search for $(L_{6r-1}^i, R_{6r-1}^i)_{i=1, \dots, r}$ in T_{lower} and obtain suggestions for $K_{4r+1}, \dots, K_{6r-1}$. For each suggestion, given $K_{4r+1}, \dots, K_{8r-1}$, partially decrypt the additional ciphertexts C'^j for $j \in \{1, \dots, 2r+1\}$ through the last $4r-1$ rounds, obtain suggestions for $\{R_{4r-1}^j\}_{j=1, \dots, 2r+1}$, and search the suggestions in $List1$. For each match, retrieve K_1, \dots, K_{4r-1} , guess K_{4r} , and test the full key using trial encryptions.

The inner loop of the algorithm is repeated for each of the $2^{(r-0.5)n}$ values of the external guess (there are $2^{(r-1) \cdot 0.5n}$ r -tuples in the data set, $2^{0.5n}$ possible values of R_{4r-1}^1 , and $2^{(r-1) \cdot 0.5n}$ possible differences $R_{4r-2}^1 \oplus R_{4r-2}^i$). In each of Steps 2(a) and 2(b), we perform 2^{rn} partial encryptions/decryptions and construct a table of size 2^{rn} . In Steps 2.(c).ii. and 2.(d).ii. there is an average of one match in T_{upper} and T_{lower} , respectively. Thus, on average, we perform a constant number of partial encryption and decryption operations per guess of K_1, \dots, K_{2r} and K_{6n}, \dots, K_{8n-1} in Steps 2.(c) and 2.(d), respectively. The expected number of matches in Step 2.(d).ii is $2^{rn+rn-(r+0.5)n} = 2^{(r-0.5)n}$, and the expected number of trial encryptions (after guessing K_4) is $2^{(r-0.5)n+0.5n} = 2^{rn}$. Therefore, the time complexity of Steps 2.(c) and 2.(d) is about 2^{rn} and the total time complexity of the attack is about $2^{(2r-0.5)n}$, as in the standard MITM attack. On the other hand, the memory complexity of the attack is reduced from $2^{(2r-0.5)n}$ to about 2^{rn} , which is the expected number of elements in $List1$.

The same attack applies for a general odd number $2r'+1$ of rounds. The time complexity is $2^{r'n}$ (like in MITM), the memory complexity has to be rounded up to $2^{\lceil 0.5(r'+1) \rceil \cdot 0.5n}$, due to lack of balance between the part of table creation and the rest of the attack, and the data complexity is $2^{\lceil \frac{r'-3}{r'+1} \rceil \cdot 0.5n}$ known plaintexts.

⁸ The computation of (L_{2r}^i, R_{2r}^i) for all $i = 1, \dots, r$ is feasible, since by the assumption that $(P^1, C^1), \dots, (P^r, C^r)$ is an r -multi-collision, the value R_{4r-2}^1 along with the externally guessed values are sufficient for obtaining the values $R_{4r-2}^2, \dots, R_{4r-2}^r$.

2.3 Attacks on Feistel Structures with an Even Number of Rounds

In this section we show that all the attacks presented above can be combined with the recent technique of Isobe and Shibutani [12] that allows to extend MITM attacks on Feistel structures by one round, at the expense of a relatively small increase in the time complexity and of using $2^{0.25n}$ chosen plaintexts. The generic attack of [12] on a $2r$ -round Feistel structures requires $2^{(0.5r-0.25)n}$ time, $2^{(0.5r-0.25)n}$ memory, and $2^{0.25n}$ chosen plaintexts. Our attacks allow to either reduce the memory complexity to about $2^{0.33(r-1)n+0.25n}$ with no effect on the time and data complexities or to reduce the memory complexity all the way to about $2^{0.25(r+1)n}$, while increasing the data complexity to $2^{\max(\lceil \frac{r-4}{r} \rceil, 0.5) \cdot 0.5n}$ chosen plaintexts, with no effect on the time complexity. In the specific case of the 8-round Feistel structures considered in [12], our attack reduces the memory complexity significantly from $2^{1.75n}$ to $2^{1.25n}$, without affecting the other complexities.

Consider a MITM attack on a $2r$ -round Feistel structure. In a standard application (skipping the guess of the middle subkey), the attack is not balanced, as r subkeys are guessed on one side of the MITM, while $r-1$ subkeys are guessed on the other side. The attack of [12] aims to rebalance the attack, by “splitting” the guess of one subkey between the two sides. The basic idea behind the attack is as follows. If in all plaintexts used in the attack, the right half is equal to a fixed value R_0 (e.g., $R_0 = 0$), then in all encryptions, we have $R_1 = Const \oplus L_0$, where $Const$ is an unknown constant that depends on K_1 . This allows to replace the $2r$ -round Feistel with an equivalent construction that consists of a $2r-1$ -round Feistel, prepended by an addition of $Const$ to the right half of the plaintext that can be treated as a subkey addition (see right side of Figure 2 for a sketch of an 8-round attack). This, in turn, allows to use the splice-and-cut technique [2] to “split” the guess of the $n/2$ -bit value $Const$ between the two sides of the MITM, at the price of using $2^{n/4}$ chosen plaintexts (each associated with an $n/4$ -bit value of $Const$, while the remaining $n/4$ bits of $Const$ are guessed from the other side of the attack). As a result, the attack becomes balanced and the time complexity is reduced from $2^{0.5rn}$ to $2^{(0.5r-0.25)n}$. For a full description of the attack, see [12].

In order to incorporate the splice-and-cut procedure of [12] into our attacks, we consider the equivalent $2r-1$ -round variant, perform one of our attacks against it, and insert the splice-and-cut procedure into the “key guessing” part of the attack (i.e., Steps 2(c) and 2(d)), without changing the “table construction” part (Steps 2(a) and 2(b)). As a result, the time complexity of our $2r-1$ -round attack is increased by a factor of $2^{0.25n}$ to $2^{(0.5r-0.25)n}$ (just like the complexity of the attack of [12]), and the memory complexity is increased by a factor of $2^{0.25n}$ to either $2^{0.33(r-1)n+0.25n}$ (in the low data complexity attack) or to about $2^{0.25(r+1)n}$ (in the attack using multi-collisions). As for the data complexity, in our low data complexity attack the data complexity increases to $2^{0.25n}$ chosen plaintexts (required for the splice-and-cut procedure), and in the multi-collision based attack the data complexity increases to $2^{\max(\lceil \frac{r-4}{r} \rceil, 0.5) \cdot 0.5n}$,

as the plaintexts required for the multi-collision can be chosen in such a way that they will contain the structures required for the splice-and-cut attack.

3 Memory-Restricted Attacks on Feistel Structures

After analyzing the most time-efficient attacks, a natural question to explore is what are the most memory-efficient attacks one can devise against an r -round Feistel structure. Specifically, we shall concentrate on the problem of devising attacks with $2^{0.5n}$ memory complexity, since it is the smallest amount of memory that enables us to list all the values of a single subkey or of half a block.

With such a restriction, a standard meet in the middle attack takes $2^{(\ell-2)\cdot 0.5n}$ on an ℓ -round Feistel, and one can trade time for memory. One can also try to consider the original dissection attack of [9]. However, as noted before, dissection takes at least 2^n memory to store all the possible values of a full block, which implies that it cannot be used in this context, even though we adopt several concepts from it.

Section 3.1 presents our new attack on 5-round Feistel structures that uses 2^n time and $2^{0.5n}$ memory. This is to be compared with meet in the middle attacks that use time of $2^{1.5n}$ with $2^{0.5n}$ memory or time of 2^n with 2^n memory. We then generalize the attack to more rounds, and show in Section 3.2 how to increase the gain over meet in the middle attacks as the number of rounds increases. The attacks in this section assume that the round function is efficiently invertible given the round's subkey. We postpone the discussion of this assumption to Appendix A, but note that it holds for almost any Feistel block cipher we are aware of.

3.1 A Memory-Restricted Attack against 5-Round Feistel Constructions

The algorithm of our basic 5-round attack is as follows.

A 5-Round Attack with $2^{0.5n}$ Memory ($DF_2(5, 1)$)

1. Obtain 3 plaintext-ciphertext pairs $(P^i, C^i)_{i=1,2,3}$.
2. For each value of $R_2^1 = I_3^1$:
 - (a) Compute $O_2^1 = I_3^1 \oplus R_0^1$ and $O_4^1 = I_3^1 \oplus R_4^1$.
 - (b) For each value of K_1 , compute $R_1^1 = F_1(K_1, I_1^1) \oplus L_0^1$ and store the pair (R_1^1, K_1) in a table T_{upper} sorted according to R_1^1 .
 - (c) For each value of K_2 , compute $R_1^1 = F_2^{-1}(K_2, O_2^1)$, search for the value R_1^1 in T_{upper} and obtain suggestions for K_1 . For each suggestion, given K_1, K_2 , compute R_2^2, R_2^3 for P^2, P^3 . Store the suggestions (R_2^2, R_2^3, K_1, K_2) in a list $List1$ sorted by the value of (R_2^2, R_2^3) .
 - (d) For each value of K_5 , compute $R_3^1 = F_5(K_5, I_5^1) \oplus R_5^1$ and store the pair (R_3^1, K_5) in a table T_{lower} sorted according to R_3^1 .

- (e) For each value of K_4 , compute $R_3^1 = F_4^{-1}(K_4, O_4^1)$, search for the value R_3^1 in T_{lower} and obtain suggestions for K_5 . For each suggestion, given K_4, K_5 , compute R_2^2, R_2^3 from C^2, C^3 , and search the suggestion in $List1$. For each match, retrieve K_1, K_2 , guess K_3 , and test the full key using trial encryptions.

For reasons which will become apparent later, we call the above attack $DF_2(5, 1)$.

As before, T_{upper}, T_{lower} , are each of size $2^{0.5n}$. The memory complexity of $List1$ depends on the number of (K_1, K_2) pairs that satisfy the meet in the middle condition on the value of I_2 in Step 2c. For sufficiently random round functions, we expect about $2^{0.5n}$ such (K_1, K_2) pairs.

The time complexity of the algorithm is 2^n , as it iterates over $2^{n/2}$ values for R_2^1 , and each step of the loop takes $2^{0.5n}$ operations (besides the XOR of Step 2a, which takes less).

We note that the time complexity of the attack can be slightly reduced by precomputing a table for F_3 (given its input value I_3^1), which allows to avoid guessing K_3 . However, this requires an additional table of size $2^{0.5n}$, which increases the memory complexity by a small constant factor.

Finally, it is important to note that given only two plaintext-ciphertext pairs, the above attack finds all possible $(K_1, K_2, K_3, K_4, K_5)$ in time 2^n and memory of $2^{0.5n}$. The expected number of candidates is about $2^{0.5n}$. This observation will be used in the subsequent attacks.

3.2 Extension to More Rounds

As the time complexity of a MITM attack with $2^{0.5n}$ memory on an ℓ -round Feistel structure is $2^{(\ell-2)0.5n}$, we define the *gain* over MITM of an attack on ℓ -round Feistel that requires T time and $2^{0.5n}$ memory by $(\ell - 2) - \log_2(T)/0.5n$. Thus, the 5-round attack presented above has gain of 1. We denote by $Gain(\ell)$ the maximal gain achieved by an ℓ -round attack with $2^{0.5n}$ memory. In this section, we extend the 5-round attack to a sequence of attacks which show that asymptotically, $Gain(\ell) = \Omega(\sqrt{\ell})$, and compute the sequence of round numbers for which the gain is strictly increased.

Obviously, it is possible to attack 6-round Feistel by guessing K_6 , and applying the 5-round attack for each guess. The result is an attack of $2^{1.5n}$ time and $2^{0.5n}$ memory on 6-round Feistel structure. This approach can obviously be extended, but maintains a gain of 1.

Attacking 10-Round Feistel Constructions To increase the gain to 2, we consider the case of 10-round Feistel, and develop the following attack:

10-Round Dissection Attack with $2^{n/2}$ Memory ($DF_5(10, 4)$)

1. Obtain 5 plaintext-ciphertext pairs $(P^i, C^i)_{i=1, \dots, 5}$.
2. For each value of (L_5^1, R_5^1) and (L_5^2, R_5^2) :

- (a) Run $DF_2(5, 1)$ on the first 5 rounds, and obtain a list of $2^{0.5n}$ candidates for $(K_1, K_2, K_3, K_4, K_5)$.
- (b) For each candidate for the subkeys $(K_1, K_2, K_3, K_4, K_5)$, partially encrypt a third plaintext P^3 through the first 5 rounds, and store the suggestions (with the keys) $(L_5^3, R_5^3, K_1, \dots, K_5)$ in a list *List1* sorted by the values of (L_5^3, R_5^3) .
- (c) Run $DF_2(5, 1)$ on the last 5 rounds, and obtain a list of $2^{0.5n}$ candidates for $(K_6, K_7, K_8, K_9, K_{10})$.
- (d) For each candidate for the subkeys $(K_6, K_7, K_8, K_9, K_{10})$, partially decrypt C^3 through the last 5 rounds, and obtain suggestions for (L_5^3, R_5^3) and search the suggestion in *List1*. For each match, retrieve K_1, \dots, K_5 , and test the full key using trial encryptions.

It is easy to see that the 10-round attack calls 2^{2n} times two independent 5-round attacks, each running in time 2^n . Hence, the time complexity of the 10-round attack is 2^{3n} , and the memory complexity is $2^{0.5n}$. Hence, the gain of the 10-round attack is 2.

We use the following notations, $DF_2(5, 1)$ denotes the 5-round attack presented earlier, as it is a generalized Dissection attack on Feistel structures with 5 rounds, which guesses one $n/2$ -bit value after two rounds of encryption. Similarly, $DF_5(10, 4)$ attacks 10 rounds by guessing 4 $n/2$ -bit values after 5 rounds of encryption (i.e., two full intermediate encryption values). As in [9], we now explore how to extend the attack to more rounds.

Attacking 15-Round Feistel Constructions We can increase the gain to 3, when attacking 15-round Feistel: Guess two complete intermediate encryption values after 5 and after 10 rounds (a total of four internal values), and run the 5-round attack three times subsequently (on rounds 1–5, 6–10, and 11–15), resulting in $2^{0.5n}$ candidates for each set of corresponding subkeys. Then, the correct value can be found by an additional standard MITM. If the memory is kept at $2^{0.5n}$, this means that the last layer of the MITM takes 2^n time. Hence, the total time complexity of the 15-round attack is 2^{5n} , and thus, its gain is 3.

In other words, $DF_5(15, 4)$ is based on guessing four $n/2$ -bit internal state words after 5 rounds, and recursively running $DF_2(5, 1)$ on the first rounds, and $DF_5(10, 4)$ on the last rounds. This contrasts with the works of [9], where each new layer in the dissection was of a different size.

The different “expanding” rule is due to two inherent differences between the dissection attacks presented in [9] and our new attacks. First, in our attacks, guessing a full intermediate state adds two “units” of time complexity (as each full state contains two $n/2$ -bit values) compared with one in the case of regular dissection attacks. The second difference is more subtle, but has a larger effect on the way the attack scales up: In our attacks we can enjoy the “Feistel” advantage (meeting in the middle only on $n/2$ bits) only once in the internal recursion step (e.g., in the 5-round attack), as the external steps must rely on guessing a full internal state. The second difference is already apparent in the transition from

5-round to 10-round (comparing $DF_2(5, 1)$ and $DF_5(10, 4)$): whereas the 5-round attack guesses a single $n/2$ -bit value, the 10-round attack starts by guessing 4 such values.

Attacking 22-Round Feistel Structures We now turn our attention to 22-round Feistel structures. Due to the differences between regular dissection attacks and attacking Feistels, the extension of the 15-round attack into the 22-round attack follows a slightly different path than the extension from the 10-round to the 15-round:

22-Round Dissection Attack with $2^{n/2}$ Memory ($DF_7(22, 6)$)

1. Obtain 11 plaintext-ciphertext pairs $(P^i, C^i)_{i=1, \dots, 11}$.
2. For each possible value of (L_7^1, R_7^1) , (L_7^2, R_7^2) , and (L_7^3, R_7^3) :
 - (a) Run $DF_2(7, 3)$ on the first 7 rounds, and obtain a list of $2^{0.5n}$ candidates for (K_1, K_2, \dots, K_7) .
 - (b) For each candidate for the subkeys (K_1, K_2, \dots, K_7) , partially encrypt a fourth plaintext P^4 through the first 7 rounds, and store the suggestions (with the keys) $(L_7^4, R_7^4, K_1, \dots, K_7)$ in a list *List1* sorted by the values of (L_7^4, R_7^4) .
 - (c) Run $DF_5(15, 4)$ on the last 15 rounds, and obtain a list of $2^{4.5n}$ candidates^a for $(K_8, K_9, \dots, K_{22})$.
 - (d) For each candidate for the subkeys (K_8, \dots, K_{22}) , partially decrypt C^4 through the last 15 rounds, and obtain suggestions for (L_7^4, R_7^4) and search the suggestion in *List1*. For each match, retrieve K_1, \dots, K_7 , and test the full key using trial encryptions.

^a We guess $3n$ bits in Step 2, giving rise to $3n$ constraints on $7.5n$ key bits.

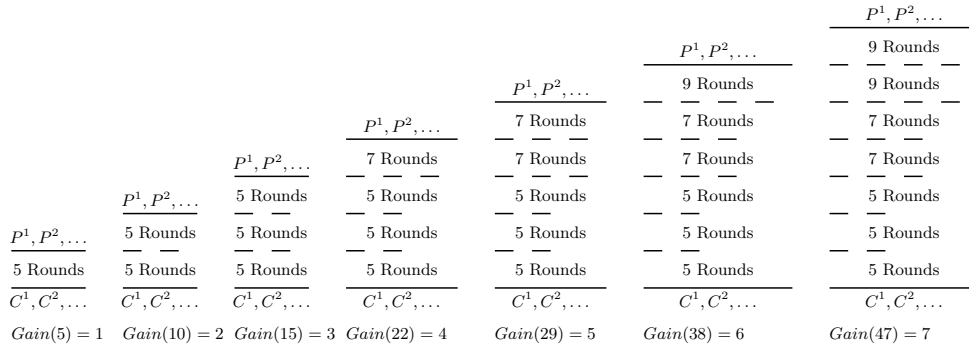
It is easy to see that the memory complexity of the attack is $2^{0.5n}$. The 7-round attack $DF_2(7, 3)$ is actually $DF_2(5, 1)$, run when K_6, K_7 are guessed, i.e., takes 2^{2n} time for $2^{0.5n}$. Both the 7-round attack and the 15-round attack are called 2^{3n} times, suggesting a total running time of 2^{8n} , i.e., the attack offers a gain of 4.

Generalization to More Rounds In the second generalization, we prepend 5 rounds to the 10-round attack (obtaining 15 rounds in total), and again, guess two full internal states in order to run two independent attacks — one on 5 rounds, and the other on 10 rounds. The 22-round attack is based on guessing three additional full internal states, and prepending 7 rounds before the 15-round attack. Similarly, a 29-round attack with a gain of 5 can be obtained by prepending 7 rounds before the 22-round attack and guessing three additional full internal states. It is now apparent that the sequence of round numbers for which the gain increases is $\{5, 10, 15, 22, 29, 38, 47, 58, 69, \dots\}$,⁹ which shows that

⁹ A gain of j is first achieved when attacking $5j + 2 \sum_{i=1}^{j-1} \lfloor i/2 \rfloor$ rounds.

for all k , $\text{Gain}(2k^2 + 6k + 2) \geq 2k$ and $\text{Gain}(2k^2 + 8k + 5) \geq 2k + 1$. It follows that asymptotically, $\text{Gain}(\ell)$ grows as $\sqrt{2\ell}$.

The general form of the recursion is described in Figure 3. We note that the recursion yields only a lower bound on $\text{Gain}(\ell)$.



Dashed lines represent internal state values which are guessed in the attack.

Fig. 3. Generalizing the 5-Round Attack and Increasing the Gain

4 Applications to Concrete Cryptosystems

While the new techniques presented in Sections 2 and 3 are generic, they can also be used to improve the memory complexity of the best known attacks on several block ciphers, including CAST-128, DEAL, and FOX. It turns out that even a straightforward application of the generic techniques already yields improvements over previously known attacks, and if we also exploit the specific properties of the analyzed cipher, the improvements become even more significant. We present in this section only a very brief description of our improved attacks on specific schemes. The detailed description of the attack on CAST-128 is given in the full version of this paper [10].

4.1 Lower Memory Attacks on DEAL

DEAL [16] is a 128-bit Feistel structure, designed in 1997 by Knudsen and submitted as a candidate to the AES selection process. The round function of DEAL is extremely complex — it consists of a full keyed DES [18] encryption. (Recall that DES has 64-bit blocks and 56-bit keys.) In return, the number of rounds is rather small — 8 rounds for the 256-bit key variant and 6 rounds for the 128-bit and 192-bit key variants. The only published attack on the full 8-round DEAL is a standard MITM attack mentioned by the designers [16], with time and memory complexities of $2^{4 \cdot 56} = 2^{224}$. The generic attack of [12] on 8-round Feistel

structures (described in Section 2.3) can be used to reduce the time complexity to $2^{3 \cdot 56 + 32} = 2^{200}$, with memory complexity of 2^{200} and data complexity of 2^{32} chosen plaintexts. We show that by using our techniques, the memory complexity can be significantly reduced to 2^{144} , while maintaining the same data and time complexities.

Our generic attack on 8-round Feistel structures (presented in Section 2.3) requires $2^{1.75n}$ time, $2^{1.25n}$ memory, and $2^{0.25n}$ chosen plaintexts. A direct application of this attack to DEAL, taking into account the fact that each round key has only 56 bits rather than 64, yields time complexity of $2^{64+56+56+32} = 2^{208}$. However, the time complexity can be reduced to 2^{200} by performing the external enumeration over 56 out of the 64 bits of the intermediate value R_4^1 , rather than over the full value. In the phase of table preparation, we guess the remaining 8 bits of R_4^1 , along with the auxiliary guess of R_3^1, K_4 , and thus, the complexity of this step is increased to $2^{8+64+56} = 2^{128}$. However, this complexity is still dominated by the $2^{32+56+56} = 2^{144}$ complexity of the key guessing step. As a result, the overall time complexity remains $2^{56+56+56+32} = 2^{200}$, the memory complexity is reduced to $2^{56+56+32} = 2^{144}$, and the data complexity remains 2^{32} chosen plaintexts.

In a similar way we can reduce the memory complexity of the improved MITM attack on the full 6-round DEAL with 192-bit keys from $2^{56+56+32} = 2^{144}$ to $2^{56+32} = 2^{88}$ (while keeping the 2^{144} time complexity and 2^{32} data complexity unchanged), using a modification of the generic attack on 6-round Feistel structures presented in Section 3. The resulting attack in this case is the best known attack which uses 2^{32} data, but is outperformed (in terms of time complexity) by the impossible differential attack presented by Knudsen [16] that requires 2^{121} time but uses 2^{70} chosen plaintexts. Table 2 compares the complexities of attacks against the variants of DEAL.

Key Size	Rounds	Complexity			Attack
		Time	Memory	Data	
192	6	2^{121}	2^{64}	2^{70}	CP Impossible differential [16]
	6	2^{144}	2^{144}	2^{32}	CP Splice-and-cut ¹⁰ [12]
	6	2^{144}	2^{88}	2^{32}	CP New
256	8	2^{224}	2^{168}	4 KP	Meet in the Middle
	8	2^{200}	2^{200}	2^{32}	CP Splice-and-cut ¹⁰ [12]
	8	2^{200}	2^{144}	2^{32}	CP New

KP — Known plaintext, CP — Chosen plaintext

Table 2. Comparison of Results against DEAL

¹⁰ This attack was not really suggested in [12], but can be derived from the paper.

4.2 A Lower Memory Attack on CAST-128

CAST-128 [1] is a 16-round Feistel structure that uses 64-bit inputs and 128-bit keys, which was designed in 1996 by Adams. It is used in several real-life products, such as GPG, PGP, and SSH2. The currently best known attack on the cipher (excluding weak-key attacks such as [21]) is the MITM attack of Isobe and Shibutani [12], breaking 8 out of the 16 rounds in time complexity of about 2^{118} , using 8 chosen ciphertexts and a memory complexity of 2^{111} words. Using our techniques, the memory complexity can be reduced significantly to 2^{64} , while maintaining the same data and time complexities.

In the round function F_i of CAST-128, the 32 LSBs of the 37-bit round key K_i (denoted as K_{m_i}) are first either XORed, added (modulo 2^{32}), or subtracted (modulo 2^{32}) from R_{i-1} (depending on the round). Then, the result is rotated to the left by 0–31 bits, according to the value of the 5 MSBs of K_i (denoted as K_{r_i}). Finally, a key-less function f_i is applied to the result. The first round of CAST-128 (that uses modular addition) is shown in Figure 4.

Since each round key of CAST-128 is of 37 bits, the time complexity of a basic MITM attack on a 8-round variant is $2^{4 \cdot 37} = 2^{148}$. Using the generic attack of [12] on 8-round Feistel structures (described in Section 2.3), the time complexity can be reduced to $2^{3 \cdot 37 + 16} = 2^{127}$, which is only slightly faster than exhaustive key search. Isobe and Shibutani [12] showed that the specific structure of the round function of CAST-128 can be used to further reduce the time complexity to 2^{118} . The main idea of [12] is that by fixing most of the ciphertext bits (in all ciphertexts) to a constant value and exploiting the specific round function structure, the amount of key material required for partial decryption can be *reduced* (and not only divided between the upper and lower halves of the MITM, like in the generic attack). To achieve this, [12] consider an equivalent 7-round Feistel structure, with different round functions F'_5, F'_6, F'_7 that imitate the four round functions F_5, F_6, F_7, F_8 for the specifically chosen ciphertexts. See the full version of this paper [10] for details of the attack.

As in case of the generic attack of Isobe and Shibutani discussed in Section 2.3, we can incorporate the advanced attack procedure in the “key guessing” part of our generic memory-efficient attack on 7-round Feistel structures. As a result, the memory complexity of the attack is reduced from 2^{111} to 2^{79} , without increasing the time and data complexities. The memory complexity can be further reduced using a refined attack that exploits the relatively simple round function of CAST-128. As we show in the full version of this paper [10], it is possible to guess two intermediate values R_3^1, R_3^2 (instead of a single value in the generic 7-round attack) and to structures separate tables T_{upper1}, T_{upper2} for rounds 2,3 (and similarly, separate tables T_{lower1}, T_{lower2} for rounds 5,6). These tables make use of complex differential properties of F_2 and F_6 that simultaneously combine different operations over $GF(2)$ and over $GF(2^{32})$. As a result, the memory complexity is reduced to 2^{64} with no effect on the time complexity. The details of this (rather involved) attack are given in the full version of this paper [10].




Fig. 4. The First Round of CAST-128

4.3 Lower Memory Attacks on Other Cryptosystems

We conclude this section by mentioning briefly applications of our generic techniques to several other specific and generic structures.

1. **FOX** Fox is a non-Feistel block cipher. The memory complexity of all the attacks of [13] on round-reduced variants of the block cipher FOX (namely, on 6 and 7-round FOX-64 and FOX-128), which are currently the best known attacks on FOX, can be reduced by a factor of 2^{16} (with no change in the data or time complexities). We note that in these attacks, the 16 bits of filtering on which the attack iterates in the outer loop of the attack are not actual state bits, but rather linear combinations of state bits. However, our techniques are still applicable in this case, as in the inner loop, for each side of the attack, we simply complement these linear combinations to obtain an intermediate encryption state of FOX, and invert its round function as done in our generic attacks.
2. **Camellia** The memory complexity of the attacks of [12] on reduced variants of Camellia can be reduced by a factor of at least 2^{16} (depending on the attack) with no effect on the data or time complexities. We note however that MITM attacks are not the best known attacks on Camellia (in terms of the number of rounds).¹¹
3. **Feistel-2 Schemes** The memory complexity of the attacks of [12] on the 8 and 9-round Feistel-2 scheme (which is a more specific Feistel implementation compared to the generic Feistel-1) with a $2n$ -bit key can be reduced from about $2^{1.5n}$ to 2^n (with no change in the data or time complexities). We note that these MITM attacks are the best known attacks on this specific Feistel-2 only when the data complexity is limited. With large data complexity in the chosen plaintext model, the attacks of [11] have superior time complexities.

5 Conclusions

In this paper we introduced some new cryptanalytic techniques, and combined the known techniques of MITM and dissection in new ways which enabled us to merge their advantages and avoid their disadvantages. Taken together, these techniques allowed us to develop improved attacks on Feistel structures with

¹¹ Although they reach fewer rounds, the advantage of MITM attacks over other attacks on Camellia (namely, impossible differential attacks [5]) is their low data complexity.

more than four rounds, and to improve the best known concrete attacks on several well known block ciphers.

References

1. C. Adams. The CAST-128 Encryption Algorithm. RFC 2144, 1997. <https://tools.ietf.org/html/rfc2144>.
2. K. Aoki and Y. Sasaki. Preimage Attacks on One-Block MD4, 63-Step MD5 and More. In R. M. Avanzi, L. Keliher, and F. Sica, editors, *Selected Areas in Cryptography, 15th International Workshop, SAC 2008, Sackville, New Brunswick, Canada, August 14-15, Revised Selected Papers*, volume 5381 of *Lecture Notes in Computer Science*, pages 103–119. Springer, 2008.
3. A. Bogdanov, D. Khovratovich, and C. Rechberger. Biclique Cryptanalysis of the Full AES. In D. H. Lee and X. Wang, editors, *ASIACRYPT*, volume 7073 of *Lecture Notes in Computer Science*, pages 344–371. Springer, 2011.
4. A. Bogdanov and C. Rechberger. A 3-Subset Meet-in-the-Middle Attack: Cryptanalysis of the Lightweight Block Cipher KTANTAN. In A. Biryukov, G. Gong, and D. R. Stinson, editors, *Selected Areas in Cryptography - 17th International Workshop, SAC 2010, Waterloo, Ontario, Canada, August 12-13, 2010, Revised Selected Papers*, volume 6544 of *Lecture Notes in Computer Science*, pages 229–240. Springer, 2010.
5. C. Boura, M. Naya-Plasencia, and V. Suder. Scrutinizing and Improving Impossible Differential Attacks: Applications to CLEFIA, Camellia, LBlock and Simon. In Sarkar and Iwata [19], pages 179–199.
6. A. Canteaut, M. Naya-Plasencia, and B. Vayssière. Sieve-in-the-Middle: Improved MITM Attacks. In R. Canetti and J. A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 222–240. Springer, 2013.
7. D. Chaum and J.-H. Evertse. Cryptanalysis of DES with a Reduced Number Of Rounds: Sequences of Linear Factors in Block Ciphers. In *Advances in Cryptology, CRYPTO 85*, pages 192–211. Springer, 1986.
8. W. Diffie and M. E. Hellman. Cryptanalysis of the NBS Data Encryption Standard. *Computer*, 10(6):74–84, 1977.
9. I. Dinur, O. Dunkelman, N. Keller, and A. Shamir. Efficient Dissection of Composite Problems, with Applications to Cryptanalysis, Knapsacks, and Combinatorial Search Problems. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 719–740. Springer, 2012.
10. I. Dinur, O. Dunkelman, N. Keller, and A. Shamir. New Attacks on Feistel Structures with Improved Memory Complexities. *IACR Cryptology ePrint Archive*, 2015:146, 2015.
11. J. Guo, J. Jean, I. Nikolic, and Y. Sasaki. Meet-in-the-Middle Attacks on Generic Feistel Constructions. In Sarkar and Iwata [19], pages 458–477.
12. T. Isobe and K. Shibutani. Generic Key Recovery Attack on Feistel Scheme. In K. Sako and P. Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 - 19th International Conference on the Theory and Application of Cryptology and Information Security, Bengaluru, India, December 1-5, 2013, Proceedings, Part I*, volume 8269 of *Lecture Notes in Computer Science*, pages 464–485. Springer, 2013.

13. T. Isobe and K. Shibutani. Improved All-Subkeys Recovery Attacks on FOX, KATAN and SHACAL-2 Block Ciphers. Presented at FSE 2014. To appear in Lecture Notes in Computer Science., 2014.
14. P. Junod and S. Vaudenay. FOX : A New Family of Block Ciphers. In H. Handschuh and M. A. Hasan, editors, *Selected Areas in Cryptography, 11th International Workshop, SAC 2004, Waterloo, Canada, August 9-10, 2004, Revised Selected Papers*, volume 3357 of *Lecture Notes in Computer Science*, pages 114–129. Springer, 2004.
15. D. Khovratovich, C. Rechberger, and A. Savelieva. Bicliques for Preimages: Attacks on Skein-512 and the SHA-2 Family. In A. Canteaut, editor, *Fast Software Encryption - 19th International Workshop, FSE 2012, Washington, DC, USA, March 19-21, 2012. Revised Selected Papers*, volume 7549 of *Lecture Notes in Computer Science*, pages 244–263. Springer, 2012.
16. L. Knudsen. DEAL - A 128-bit Block Cipher, 1998. NIST AES Proposal.
17. M. Luby and C. Rackoff. How to Construct Pseudorandom Permutations from Pseudorandom Functions. *SIAM J. Comput.*, 17(2):373–386, 1988.
18. National Bureau of Standards. Data encryption standard. Federal Information Processing Standards Publications (FIPS) 46, 1977.
19. P. Sarkar and T. Iwata, editors. *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*. Springer, 2014.
20. L. Wang and Y. Sasaki. Finding Preimages of Tiger Up to 23 Steps. In S. Hong and T. Iwata, editors, *Fast Software Encryption, 17th International Workshop, FSE 2010, Seoul, Korea, February 7-10, 2010, Revised Selected Papers*, volume 6147 of *Lecture Notes in Computer Science*, pages 116–133. Springer, 2010.
21. M. Wang, X. Wang, K. Chow, and L. C. K. Hui. New Differential Cryptanalytic Results for Reduced-Round CAST-128. *IEICE Transactions*, 93-A(12):2744–2754, 2010.

A On the Invertibility of the Round Function

We first note that our 5-round only needs two round functions to be efficiently invertible — namely, $F_2(\cdot)$ and $F_4(\cdot)$. As noted before, we are not aware of any Feistel cipher which does not possess this property. Even DEAL [16], whose round functions are full DES encryptions, has invertible round functions given the subkey, as our attack described in Section 4.1 shows.

Moreover, we can relax a bit the requirement over the invertibility of the round functions $F_2(\cdot)$ and $F_4(\cdot)$. We remind the reader that we are allowed $2^{0.5n}$ memory, which can help in inverting the round functions. For example, if the cipher is a Feistel-2 structure (i.e., the round function is $F_i(K_i, I_i) = G_i(K_i \oplus I_i)$, for some completely one-way function $G_i(\cdot)$), a simple enumeration of all input/output pairs of $G_i(\cdot)$ is sufficient to invert the round function.

Finally, we note that when we discuss the general Feistel-2 structure, the memory complexity can be slightly reduced, as no memory is needed for the meet in the middle step in itself. For example, instead of structure T_1 , given O_2^1 , we invert G_2 , to obtain $I_2^1 \oplus K_2$. Hence, for any K_1 value, it is possible to immediately obtain the corresponding K_2 .