Exponential attacks on 6-round Luby-Rackoff and on 5-round Lai-Massey

Jean-Philippe Aumasson

Nagravision SA, Cheseaux, Switzerland

Abstract. The random oracle model and the ideal cipher model were proven equivalent after Coron et al. (CRYPTO '08) showed that six Feistel rounds are indifferentiable from an ideal cipher. This result, however, does not imply the inexistence of superpolynomial-time attacks outperforming generic (exponential-time) attacks. The finding of such attacks was left open by Coron et al., and is of utmost importance to evaluate the security of concrete fixed-parameters systems, as deployed in practice, for which the superpolynomial guarantee is an insufficient security argument. In addressing this issue, this paper proposes an exponential attack on six Feistel rounds, thus showing that at least seven rounds are necessary for optimal security guarantees. We then consider the Lai-Massey construction, as used in the block ciphers IDEA and FOX, for which we present an efficient attack on four rounds and an exponential attack on five. As a consequence, at least five Lai-Massey rounds are necessary to achieve indifferentiability in the general model.

Keywords: block ciphers, indifferentiability, Luby-Rackoff, Lai-Massey

1 Introduction

The Luby-Rackoff scheme is with the substitution-permutation network the most common block cipher construction. Examples of Luby-Rackoff ciphers (henceforth LR ciphers) are DES, CAST5, KASUMI, RC5, or Twofish. LR ciphers transform an internal state (L, R) through r rounds that set

$$(L, R) \leftarrow (R, L \oplus F_i(R)), \quad i = 1, \ldots, r \ ,$$

where the F_i's are distinct key-dependent functions. A less common construction is the Lai-Massey (LM) scheme, whose round function transforms the state by doing $S \leftarrow F_i(L \oplus R)$ followed by

$$(L, R) \leftarrow \big(\sigma(L \oplus S), R \oplus S\big) \ ,$$

where σ is a function that breaks the scheme's symmetry. LM can achieve full diffusion in one round, against two for LR. Examples of LM ciphers are IDEA [1] and FOX [2].

From PRPs to (super-) indifferentiability. Luby and Rackoff showed [3] that three Feistel rounds are sufficient to construct a pseudorandom permutation (PRP) and that four suffice to construct a strong PRP (sPRP) if the F_i's are pseudorandom. Vaudenay proved similar results for three and four LM rounds [4]. The PRP and sPRP notions, however, fail to model the ability of the cipher to instantiate an ideal cipher. The "right" notion for this—and more generally, to substitute and ideal primitive by a construction based on another ideal primitive—is that of *indifferentiability* [5], wherein attackers interact with the subprimitives rather than with the main primitive.

The notion of *indifferentiable construction* is best illustrated by the 2005 result of Coron et al. [6], who showed that certain block cipher-based constructions of hash functions can securely replace a random oracle if the block cipher is ideal. This result implies that schemes proven secure in the random oracle model (ROM) remain secure in the ideal cipher model (ICM).

A major progress was presented at CRYPTO 08 with Coron et al.'s proof [7] that 6-round LR is indifferentiable from an ideal cipher. This showed that the ROM and the ICM model rely on equivalent assumptions, and so that they are asymptotically of equal strength.

Six Feistel rounds were proven to be *necessary* [7] by showing an attack on 5-round LR. Attacks with respect to indifferentiability generalize attacks under the classical notion of indistinguishability. Roughly speaking, attacks specific to indifferentiability efficiently find input/output pairs satisfying some "evasive" property, i.e. a property provably difficult to satisfy for an ideal cipher. Here, "efficiently" and "difficult" respectively stand for "polynomial-time" and "superpolynomial-time".

Nevertheless, investigating superpolynomial-time attacks that outperform generic methods is of great practical interest, for asymptotic bounds are irrelevant to fixed-parameters algorithms. For example, meet-in-the-middle attacks are exponential in nature, yet they reduce the security of three-keys Triple DES from 168 to 112 bits, or the preimage resistance of the SHA-3 candidate hash function CubeHash from 512 to 384 bits. Also, the devastating attacks on MD5 have complexity essentially exponential, but with a dangerously low exponent. Exponential distinguishing attacks also significantly assist the security evaluation of SHA-3 candidates. The problem of finding such *exponential attacks* on 6-round LR was left open by Coron et al. [7, §5].

To study that problem, we propose the term *super-indifferentiability* to denote the practice-oriented version of indifferentiability, in the spirit of Bellare and Rogaway's provable security framework (such a "concrete" version of indifferentiability was previously defined in [8, §2]). In the super-indifferentiability model, any method achieving a goal more efficiently than generic methods constitutes an attack, and thus refutes the super-indifferentiability of the scheme considered[1].

[1] Admittedly, the notion of super-indifferentiability (which has been previously defined, but not named) is of modest theoretical interest and prone to cause technical

To our best knowledge, the most related works are the series of papers by Patarin [9–11], which report exponential attacks on up to six LR rounds, using the "coefficients H" technique [12]. However, these results are in the black-box model and rely on statistical arguments rather than on the construction of evasive relations in a white-box setting, as in the present work. These works thus study the security of LR in a different attack model but reach conclusions similar to ours.

Our results. This paper presents an algorithm that finds many solutions to the 4-sum problem—an instance of the generalized birthday problem [13]—with respect to 6-round LR more efficiently than the best known method [13, 14] and that any generic attack. It follows that seven LR rounds are necessary to achieve super-indifferentiability. As the attack runs in exponential time, it does not contradict the indifferentiability of 6-round LR.

We then present an efficient attack on 4-round LM, and an exponential attack on 5-round LM. Therefore, at least five (resp., six) rounds are necessary to achieve (resp., super-) indifferentiability. To the best of our knowledge, this is the first published analysis of the LM scheme with respect to indifferentiability.

The strategy of our exponential attacks is to precompute collisions for a subprimitive of the construction, then to exploit freedom degrees available from intermediate state values so as to construct multiple input/output pairs satisfying some evasive property. For 6-round LR, this property is a set of 4-sums, and for LM a linear relation satisfied by one input/output pair.

Table 1 summarizes necessary and sufficient number of rounds to achieve each security notion, and Fig. 1 depicts the constructions considered. We do not claim that our attacks on LR and LM are optimal, nor that no such shortcut attacks exist for higher number of rounds (i.e. we do not claim upper bounds).

Table 1. Necessary and sufficient number of rounds to achieve pseudorandomness (PRP), strong pseudorandomness (sPRP), indifferentiability (IND), and super-indifferentiability (sIND). The three lower bounds are results from this paper.

Construction	PRP	sPRP	IND	sIND
Luby-Rackoff	3 [3]	4 [3]	6 [7]	$\geq \mathbf{7}$
Lai-Massey	3 [4, 15]	4 [4, 15]	$\geq \mathbf{5}$	$\geq \mathbf{6}$

2 Attacks on Luby-Rackoff

This section presents an attack on 6-round LR. To expose our strategy, we first show a simple attack on 5-round LR in §2.2, which we generalize in §2.3 to attack

difficulties when constructing security proofs. Nonetheless, it is more relevant in practice, in particular for the analysis of hash function algorithms (see, e.g. [8]).

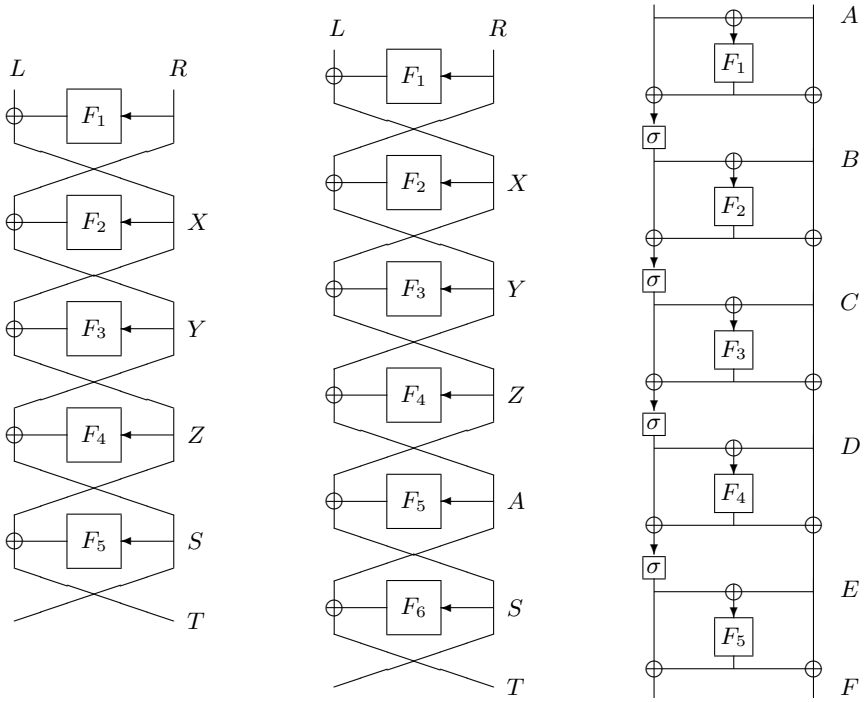

Fig. 1. 5-round Luby-Rackoff, 6-round Luby-Rackoff, and 5-round Lai-Massey.

six rounds. Note that a more efficient, constant-time attack on 5-round LR was described in [7].

2.1 Notations

We consider an LR construction processing $2n$-bit values, thus with two n-bit branches, with inner functions $F_i : \{0,1\}^n \to \{0,1\}^n$. We denote the n-bit intermediate values using the same notations as Coron et al. [7], namely $L\|R$ for the input values, then X, Y, Z, S for the subsequent inputs for the F_i functions, and the output is $S\|T$ (as depicted on Fig. 1).

2.2 Exponential attack on 5-round LR

This attack starts by finding a collision for F_3, which is done in $O(2^{n/2})$ queries[2] to F_3. One obtains Y and Y' such that

$$F_3(Y) = F_3(Y'),\ Y \neq Y' .$$

[2]The reader may legitimately object that we use here asymptotic notations, whereas we previously argued that the motivation of our work was the study of concrete, fixed parameters constructions. We respond that our use of such notations only hides small constant factors coming from the collision search overhead, and is thus admissible for our purpose.

Let $\Delta := Y \oplus Y'$, and let $X = X'$ be an arbitrary value. We denote $L, R, \ldots, T$ the intermediate values of the (X, Y) instance and $L', R', \ldots, T'$ those of the (X', Y') instance (see Fig. 1). We thus have

$$R \oplus R' = \big(F_2(X) \oplus Y\big) \oplus \big(F_2(X') \oplus Y'\big) = \Delta \ .$$

Since $Z = X \oplus F_3(Y) = X' \oplus F_3(Y') = Z'$, we also have

$$S \oplus S' = \big(F_4(Z) \oplus Y\big) \oplus \big(F_4(Z') \oplus Y'\big) = \Delta \ .$$

Therefore, the two input/outputs $(L, R)/(S, T)$ and $(L', R')/(S', T')$ satisfy the relation $R \oplus R' = S \oplus S'$. The complexity of our method is essentially that of finding a collision for F_3, i.e. $O(2^{n/2})$.

Clearly, finding such values for an ideal cipher has complexity $\Theta(2^{n/2})$, using a collision search method. Indeed, the problem is equivalent to that of finding a collision for the function $R \mapsto R \oplus E(L\|R)$, where E is the 5-round LR block cipher with output truncated to S (this construction is similar to the Davies-Meyer construction of compression functions).

We now generalize the above method: by choosing $1 < m \leq 2^{n/2}$ distinct values of $X = X'$, one can find m pairs of input/outputs $\big((L_i, R_i), (L'_i, R'_i)\big)$, $i = 0, \ldots, m-1$, such that

$$R_0 \oplus R'_0 = S_0 \oplus S'_0 = R_1 \oplus R'_1 = S_1 \oplus S'_1 = \cdots = R_{m-1} \oplus R'_{m-1} = S_{m-1} \oplus S'_{m-1} \ .$$

Since the cost of the search is dominated by that of the collision search for F_3 (which needs be done only once), the complexity of finding those m pairs remains in $O(m+2^{n/2})$, whereas the complexity for an ideal cipher is in $\Omega(m2^{n/2})$ (tighter bounds can be found). Our method thus sets 5-round LR apart from an ideal cipher.

2.3 Exponential attack on 6-round LR

The attack on 6-round LR is based on the attack on five rounds, but it exploits multiple internal collisions for two internal functions, rather than one collision for one function. We first describe the method when exploiting one collision for F_3 and one for F_4. We then generalize to $p \geq 1$ collisions for F_3 and $q \geq 1$ for F_4, in order to attack 6-round LR.

Exploiting 1+1 collisions. Let (Y, Y') be a collision for F_3 and (Z, Z') a collision for F_4, with nonzero $Y \oplus Y' = \Delta$ and $Z \oplus Z' = \nabla$:

$$\begin{aligned} F_3(Y) &= F_3(Y') \\ F_4(Z) &= F_4(Z') \ . \end{aligned}$$

For each of the four (Y, Z), (Y, Z'), (Y', Z), and (Y', Z'), one can determine the corresponding input to the LR construction, which we denote (L_0, R_0), (L_1, R_1),

(L_2, R_2), (L_3, R_3), respectively. Similarly, we denote outputs (S_i, T_i), and other intermediate values X_i and A_i, $i = 0, \ldots, 3$.

Note that, for all four pairs, both F_3 and F_4 return the same four values:

$$F_3(Y_0) = F_3(Y_1) = F_3(Y_2) = F_3(Y_3)$$
$$F_4(Z_0) = F_4(Z_1) = F_4(Z_2) = F_4(Z_3) \ .$$

We thus distinguish two cases:

1. Collision for F_3 with a constant Z (instances 0 and 2, 1 and 3).
2. Collision for F_4 and a constant Y (instances 0 and 1, 2 and 3).

In the first case, no difference propagates to X (i.e. $X_0 = X_2$, $X_1 = X_3$), thus

$$R_0 \oplus R_2 = R_1 \oplus R_3 = \Delta \ .$$

In the second case, no difference propagates to A (i.e. $A_0 = A_1$, $A_2 = A_3$), thus

$$S_0 \oplus S_1 = S_2 \oplus S_3 = \nabla \ .$$

To summarize, we have four input/output pairs $(L_i, R_i)/(S_i, T_i)$, $i = 0, \ldots, 3$, such that

$$R_0 \oplus R_1 \oplus R_2 \oplus R_3 = S_0 \oplus S_1 \oplus S_2 \oplus S_3 = 0 \ .$$

For an ideal cipher, finding such values is equivalent to the 4-sum problem, as considered by Wagner in [13] (see also [16]). It has previously been shown that truncated pseudorandom permutation behave (to some extent) as pseudorandom function [8,17], thus the analysis made for random functions in [13] is applicable: the generalized birthday method has complexity in $O(2^{n/3})$ time and space. Note that a lower bound for the 4-sum problem is $\Omega(2^{n/4})$ [13, Th.1].

Our method is thus *less efficient* than Wagner's method for finding a single 4-sum, as it runs in $O(2^{n/2})$. However, we will show that it outperforms generic methods when finding *many* 4-sums.

Exploiting $p+q$ collisions. Suppose that we precompute more than one collision for both F_3 and F_4; let p be the number of distinct colliding pairs for F_3, and q the number of colliding pairs for F_4. We assume that all $2p$ (resp., $2q$) inputs to F_3 (resp., F_4) are distinct, which can be guaranteed by adapting the generic collision search with negligible computation overhead. Note that we only search for collisions, not multicollisions.

We now have pq distinct quadruplets (Y, Y', Z, Z') such that

$$F_3(Y) = F_3(Y')$$
$$F_4(Z) = F_4(Z') \ .$$

Thus, the above method can be used to find pq distinct 4-sums in $O((p+q)2^{n/2})$. This should be compared with the best generic method known (pq applications

of the generalized birthday), with complexity $O(pq2^{n/3})$ in time and $O(2^{n/3})$ in space.

The dedicated method is thus faster than the generalized birthday when

$$\frac{pq}{p+q} \geq 2^{n/6} \ .$$

In particular, if $p = q$ our method runs in $O(p2^{n/2})$, against $O(p^2 2^{n/3})$ for the generalized birthday; it thus has lower complexity when p is greater than $2^{n/6}$.

Furthermore, when

$$\frac{pq}{p+q} \geq 2^{n/4} \ ,$$

for example when $p = q = 2^{n/4}$, then our method outperform any generic attack on an ideal cipher, as finding pq 4-sums has complexity $\Omega(pq2^{n/4})$. Our method thus constitutes an attack on 6-round LR with respect to super-indifferentiability.

For example, if $n = 32$ as in DES, our attack outperforms the generalized birthday if one finds 1024 collisions for F_3 and F_4 and exploits them to find $1024^2 = 2^{20}$ 4-sums, with complexity approximately 2^{27}, instead of approximately $2^{30.66}$ with the standard method and of 2^{28} as per the $O(pq2^{n/4})$ lower bound (note that these estimates ignore constant multiplicative factors).

3 Attacks on Lai-Massey

This section presents attacks on 4- and 5-round LM: §3.2 shows how to attack 4-round LM in $O(1)$, then §3.3 describes an exponential attack on 5-round LM, exploiting a collision for an inner function.

3.1 Notations

We consider a LM construction processing $4n$-bit values, with inner functions $F_i : \{0,1\}^{2n} \to \{0,1\}^{2n}$. We view each branch as a two n-bit subbranches, and we consider the function $\sigma : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, such that

$$\sigma : (X, Y) \mapsto (Y, X \oplus Y) \ .$$

A similar σ function is used in the block cipher FOX [2]. The only criterion is that σ be an orthomorphism[3] [4]. Note that no σ is made after the final round, for σ is key-independent.

The initial state is denoted $A = (A_1, A_2, A_3, A_4)$, from left to right, and subsequent intermediate states are denoted B, C, etc. (see Fig. 1).

[3] An orthomorphism σ on a group $(\mathcal{G}, +)$ is a permutation $x \mapsto \sigma(x)$ such that $x \mapsto \sigma(x) - x$ is also a permutation.

3.2 Efficient attack on 4-round LM

This attack starts by selecting an arbitrary intermediate state $C = (C_1, C_2, C_3, C_4)$, and an arbitrary nonzero n-bit value Δ. We then consider the four states

$$\begin{aligned} C^1 &= (C_1, C_2, C_3, C_4) \\ C^2 &= C^1 \oplus (\Delta, 0, \Delta, 0) \\ C^3 &= C^1 \oplus (0, \Delta, \Delta, 0) \\ C^4 &= C^1 \oplus (\Delta, \Delta, 0, 0) \ . \end{aligned}$$

These states satisfy

$$\begin{aligned} C^1 \oplus C^2 = C^3 \oplus C^4 &= (\Delta, 0, \Delta, 0) \\ C^1 \oplus C^3 = C^2 \oplus C^4 &= (0, \Delta, \Delta, 0) \ . \end{aligned}$$

Hence, the value entering F_3 is $(C_1 \oplus C_3) \| (C_2 \oplus C_4)$ for both C^1 and C^2, and is $(C_1 \oplus C_3 \oplus \Delta) \| (C_2 \oplus C_4 \oplus \Delta)$ for both C^3 and C^4. It follows that, after the third round, the four respective states $D^1, \ldots, D^4$ satisfy

$$D^1 \oplus D^2 = D^3 \oplus D^4 = (0, \Delta, \Delta, 0) \ .$$

Since the fourth round adds a same difference to the first and to the third n-bit component of the state, we have

$$\begin{aligned} E_1^1 \oplus E_3^1 \oplus E_1^2 \oplus E_3^2 = E_2^1 \oplus E_4^1 \oplus E_2^2 \oplus E_4^2 &= \Delta \\ E_1^3 \oplus E_3^3 \oplus E_1^4 \oplus E_3^4 = E_2^3 \oplus E_4^3 \oplus E_2^4 \oplus E_4^4 &= \Delta \ . \end{aligned}$$

Going backwards from C, the value entering F_2 is $(C_1 \oplus C_2 \oplus C_3) \| (C_1 \oplus C_4)$ for both C^1 and C^3, and is $(C_1 \oplus C_2 \oplus C_3) \| (C_1 \oplus C_4 \oplus \Delta)$ for both C^2 and C^4. Thus, the states $B^1, \ldots, B^4$ before the second round satisfy

$$B^1 \oplus B^3 = B^2 \oplus B^4 = (\Delta, 0, \Delta, 0) \ .$$

After inverting the orthomorphism following the first round, we reach differences $(\Delta, \Delta, \Delta, 0)$ between the intermediate states of instances 1 and 3, and of instances 2 and 4. After inverting the first round, we thus obtain initial states $A^1, \ldots, A^4$ such that

$$\begin{aligned} A_1^1 \oplus A_3^1 \oplus A_1^3 \oplus A_3^3 = A_1^2 \oplus A_3^2 \oplus A_1^4 \oplus A_3^4 &= 0 \\ A_2^1 \oplus A_4^1 \oplus A_2^3 \oplus A_4^3 = A_2^2 \oplus A_4^2 \oplus A_2^4 \oplus A_4^4 &= \Delta \ . \end{aligned}$$

To summarize, we found four input/output pairs that satisfy $8n$ bit conditions, on both their inputs and their outputs. Clearly, realizing such a property for an ideal cipher requires exponential time, whereas our method runs in constant time. Note that the difference Δ can be chosen, contrary to the attacks on LR presented in §2.

3.3 Exponential attack on 5-round LM

For some arbitrary $B = (B_1, B_2, B_3, B_4)$, let Δ and ∇ be nonzero values and define

$$B' = B \oplus (\Delta, \nabla, \Delta, \nabla) \ .$$

We thus have after the second round C and C' such that

$$C' = C \oplus (\nabla, \Delta \oplus \nabla, \Delta, \nabla) \ .$$

The difference entering F_3 is thus $(\Delta \oplus \nabla, \Delta)$. The attack consists in searching for values of Δ and ∇ such that

$$F_3\big((C_1 \oplus C_3) \| (C_2 \oplus C_4)\big) \oplus F_3\big((C'_1 \oplus C'_3) \| (C'_2 \oplus C'_4)\big) = (\Delta \oplus \nabla, \Delta) \ ,$$

which implies that a same input enters F_4 for both D and D', since

$$D \oplus D' = (\Delta, \Delta \oplus \nabla, \Delta, \Delta \oplus \nabla) \ .$$

By the birthday paradox, finding Δ and ∇ satisfying the $2n$ bit conditions on F_3's output needs $\Omega(2^n)$ queries to F_3. Indeed, the problem is equivalent to that of finding a collision for the function $X \mapsto F_3(X) \oplus X$, which can be done with standard (memoryless) collision search techniques.

Observe that, once a solution is found, one can determine 2^{2n} distinct values of B such that B and B' have the desired difference $(\Delta, \nabla, \Delta, \nabla)$, and such that a collision occurs for F_3.

Once B, Δ, and ∇ have been chosen, we have after the fourth round

$$E \oplus E' = (\nabla, \Delta, \Delta, \Delta \oplus \nabla) \ .$$

The final states F and F' after the fifth round thus satisfy

$$\begin{aligned} F_1 \oplus F'_1 \oplus F_3 \oplus F'_3 &= \Delta \oplus \nabla \\ F_2 \oplus F'_2 \oplus F_4 \oplus F'_4 &= \nabla \ . \end{aligned}$$

Going backwards from B and B', we obtain difference $(\Delta \oplus \nabla, \Delta, \Delta, \nabla)$ after inverting σ. It follows that A and A' satisfy

$$\begin{aligned} A_1 \oplus A'_1 \oplus A_3 \oplus A'_3 &= \nabla = F_2 \oplus F'_2 \oplus F_4 \oplus F'_4 \\ A_2 \oplus A'_2 \oplus A_4 \oplus A'_4 &= \Delta \oplus \nabla = F_1 \oplus F'_1 \oplus F_3 \oplus F'_3 \end{aligned}$$

We thus found two input/ouput pairs A/F and A'/F' satisfying the above relation. Clearly, finding such pairs for an ideal cipher is in $\Theta(2^n)$, thus our $O(2^n)$ method does not constitute an attack. However, observe that by exploiting $1 \leq m \leq 2^{2n}$ distinct values of B and a single collision for $X \mapsto F_3(X) \oplus X$, our method can be used to find m distinct input/output pairs satisfying the relation in $O(m + 2^n)$. For comparison, for an ideal cipher, the optimal strategy is to reiterate the generic attack and thus finding m such input/output pairs costs $\Theta(m2^n)$.

4 Conclusion

We showed that six Feistel rounds admit an exponential-time attack, which does not contradict the result of indifferentiability in the general model, but shows that this scheme is unideal, and thus may not provide optimal security. A same result was found for five Lai-Massey rounds, which were proven necessary (but not sufficient nor insufficient) for general indifferentiability. Our work thus leaves the following questions unanswered:

- Are there better exponential attacks on 6-round LR and 5-round LM (e.g. exploiting 3-collisions [18] for the inner functions)?
- How many LM rounds are necessary and sufficient to achieve indifferentiability?
- How many LM and LR rounds are necessary and sufficient to achieve super-indifferentiability (i.e. the "concrete" version of indifferentiability, not based on the polynomial vs. superpolynomial dichotomy)?

We plan to address those issues in a sequel to this work.

References

1. Lai, X., Massey, J.L.: A proposal for a new block encryption standard. In: EUROCRYPT. (1990)
2. Junod, P., Vaudenay, S.: FOX: a new family of block ciphers. In: SAC. (2004)
3. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM Journal of Computing **17**(2) (1988)
4. Vaudenay, S.: On the Lai-Massey scheme. In: ASIACRYPT. (1999)
5. Maurer, U.M., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: TCC. (2004)
6. Coron, J.S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: How to construct a hash function. In: CRYPTO. (2005)
7. Coron, J.S., Patarin, J., Seurin, Y.: The random oracle model and the ideal cipher model are equivalent. Cryptology ePrint Archive, Report 2008/046 (2008) Full version of [19].
8. Bertoni, G., Daemen, J., Peeters, M., Assche, G.V.: On the indifferentiability of the sponge construction. In: EUROCRYPT. (2008)
9. Patarin, J.: Generic attacks on Feistel schemes. In: ASIACRYPT. (2001)
10. Patarin, J.: Luby-Rackoff: 7 rounds are enough for $2^{n(1-\epsilon)}$security. In: CRYPTO. (2003)
11. Patarin, J.: Security of random Feistel schemes with 5 or more rounds. In: CRYPTO. (2004)
12. Patarin, J.: Etude des générateurs de permutations basés sur le schéma du DES. PhD thesis, INRIA, Domaine de Voluceau, Le Chesnay, France (1991)
13. Wagner, D.: A generalized birthday problem Full version of a CRYPTO 2002 paper. Available at `http://www.cs.berkeley.edu/~daw/papers/genbday.html`.
14. Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. In: STOC. (2000)
15. Luo, Y., Lai, X., Gong, Z., Wu, Z.: Pseudorandomness analysis of the Lai-Massey scheme. Cryptology ePrint Archive, Report 2009/266 (2009) See also [20].

16. Bernstein, D.J.: Better price-performance ratios for generalized birthday attacks. In: SHARCS. (2007)
17. Dodis, Y., Reyzin, L., Rivest, R.L., Shen, E.: Indifferentiability of permutation-based compression functions and tree-based modes of operation, with applications to MD6. In: FSE. (2009)
18. Joux, A., Lucks, S.: Improved generic algorithms for 3-collisions. In: ASIACRYPT. (2009)
19. Coron, J.S., Patarin, J., Seurin, Y.: The random oracle model and the ideal cipher model are equivalent. In: CRYPTO. (2008)
20. Luo, Y., Lai, X., Gong, Z.: Pseudorandomness analysis of the (extended) Lai-Massey scheme. Information Processing Letters **111**(2) (2010)