Differential Forgery Attack Against LAC

Gaëtan Leurent(✉)

Inria, project-team SECRET, Rocquencourt, France
Gaetan.Leurent@inria.fr

Abstract. LAC is one of the candidates to the CAESAR competition. In this paper we present a differential forgery attack on LAC. We study the collection of characteristics following a fixed truncated characteristic, in order to obtain a lower bound on the probability of a differential. We show that some differentials have a probability higher than 2^{-64}, which allows a forgery attack on the full LAC.

This work illustrates the difference between the probability of differentials and characteristics, and we describe tools to evaluate the probability of some characteristics.

Keywords: Differential cryptanalysis · Differentials · Characteristics · Forgery attack · Truncated differential · LBlock · LAC

1 Introduction

The CAESAR competition is an ongoing effort to identify new authenticated encryption primitives [3]. Authenticated encryption schemes provide both confidentiality and authenticity in a single primitive, instead of using an encryption scheme together with a MAC. The competition received 57 submissions in March 2014, and an important effort is now devoted to analyzing those candidates.

LAC is a CAESAR candidate designed by Zhang, Wu, Wang, Wu, and Zhang [10]. LAC uses the same structure as ALE [2]: it is based on a modified block cipher (the G function in LAC is based on LBlock [9]) that leaks part of its state. The main step of the algorithm is to encrypt the current state, and the leaked data is used as a keystream to produce the ciphertext. Meanwhile, a key schedule produces new keys for each encryption, and plaintext blocks are xored inside the state, so that the final state can be used to produce the tag T. This is depicted in Fig. 1.

In LAC, the main state is 64-bit wide, the key register is 80-bit wide, and the plaintext is divided in blocks of 48 bits. The G function is a modified version of LBlock. It uses 16 rounds of Feistel network, where the round function F applies a key addition, 8 parallel S-Boxes on the nibbles of the state, (the 8 S-Boxes are identical), and a nibble permutation. In addition, the inactive branch of the Feistel network is rotated by 2 nibbles; this is shown in in Fig. 2. The S-Box has a maximum differential probability of 2^{-2}, which is optimal for a 4-bit S-Box; it is defined as:

O. Dunkelman and L. Keliher (Eds.): SAC 2015, LNCS 9566, pp. 217–224, 2016.
DOI: 10.1007/978-3-319-31301-6_13

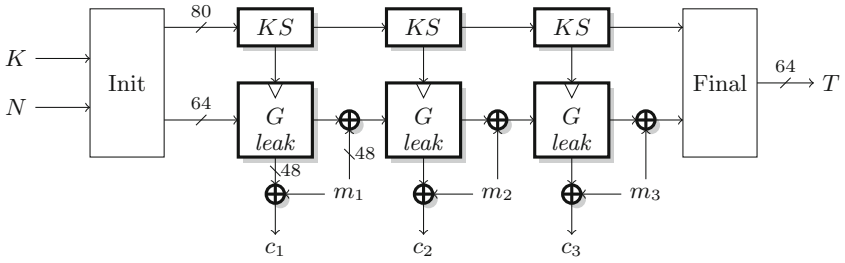

Fig. 1. LAC main structure

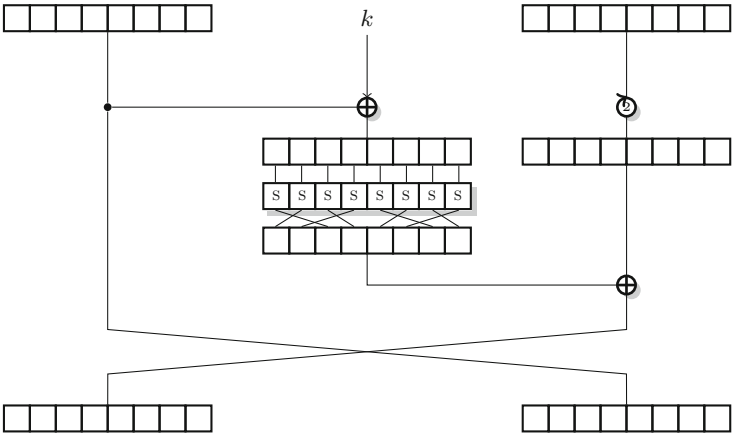

Fig. 2. A Feistel round of LAC (LBlock-s)

$$S = \{\mathtt{E}, \mathtt{9}, \mathtt{F}, \mathtt{0}, \mathtt{D}, \mathtt{4}, \mathtt{A}, \mathtt{B}, \mathtt{1}, \mathtt{2}, \mathtt{8}, \mathtt{3}, \mathtt{7}, \mathtt{6}, \mathtt{C}, \mathtt{5}\}.$$

We omit the description of the leak function and of the key schedule, because they don't affect our attack.

The security goals of LAC against forgery attacks is stated as:

> **Claim 2 (Integrity for the Plaintext).** The security claim of integrity for the plaintext is that any forgery attack with an unused tuple $(PMN^*, \alpha^*, c^*, \tau^*)$ has a success probability at most 2^{-64}.

1.1 Description of the Attack

Our attack is a differential forgery attack: given the authenticated encryption (C, T) of a message M, we build a cipher-text $(C', T') = (C \oplus \Delta, T)$ that is valid with a probability higher than 2^{-64}.

More precisely, we use a two-block difference $\Delta = (\alpha, \beta)$ so that a difference α is first injected in the state, and we predict the difference β after one evaluation

of G in order to cancel it. This will be successful if we can find a differential $\alpha \rightsquigarrow \beta$ in the function G with a probability p higher than 2^{-64}.

In Sect. 2.2, we give a differential with probability $p \approx 2^{-61.52}$. This yields a forgery attack using a single known ciphertext (of at least two blocks), with a success probability of $2^{-61.52}$.

In addition, the truncated characteristic we use does not affect the leaked output, so that, if the tag is valid, the plaintext corresponding to $(C \oplus \Delta, T)$ is $M \oplus \Delta$.

1.2 Characteristics and Differentials

We now introduce important notions for differential cryptanalysis.

A differential is given by an input difference α and an output difference β. The probability of the differential is the probability than a pair of plaintext with difference α gives a pair of ciphertext with difference β:

$$\Pr\left[\alpha \rightsquigarrow \beta\right] = \Pr_{K,x}\left[E(x \oplus \alpha) = E(x) \oplus \beta\right].$$

The probability of differentials is important to evaluate the security of a cipher against differential cryptanalysis, but it is quite challenging to compute this probability. Therefore, we introduce the notion of characteristics.

A characteristic is given by an input difference α, the difference α_i after each round, and the output difference β. Since all the intermediate difference are fixed, it is quite easy to evaluate the probability of a characteristic using the Markov cipher model (*i.e.* assuming that the rounds are independent). The probability of the differential $\alpha \rightsquigarrow \beta$ is the sum of the probability of all characteristics with input difference α and output difference β.

The designers of LAC studied its resistance against differential cryptanalysis using truncated characteristics. They show that any characteristic must have at least 35 active S-Boxes. Since the best transitions for the S-Box have a probability of 2^{-2}, any characteristic has a probability at most 2^{-70}. However, this does not imply a lower bound for the probability of *differentials*: if many good characteristics contribute to the same differential, the probability can increase significantly.

Proving an upper bound on the probability of differential is much harder than proving an upper bound for characteristics, and very few results are known in this setting. A notable example is the AES, for which an upper bound of 2^{-150} for any 4-round characteristic can easily be shown [4], and an upper bound of 1.881×2^{-114} for any 4-round differential was proved using significantly more advanced techniques [5].

In this work we give a more accurate estimation of the probability of differentials in the G function of LAC by considering more than one characteristic. Our results actually lead to a lower bound on the probability of some differential.

2 Characteristics Following the Same Truncated Trail

Truncated differential cryptanalysis was introduced by Knudsen in 1994 [6]. A truncated characteristic D does not specify the exact value of the differences at each step but uses partial information. The state is divided in words of a fixed size (usually bytes, but we use nibbles for LAC), and D only specifies whether the difference in each word in zero (inactive word) or non-zero (active word).

For a given truncated characteristic D, there exist many ways to instantiate the input/output differences and the intermediate differences. For a given input/output difference (α, β), we consider all the possible intermediate differences following D; this defines a collection of characteristics that all contribute to the same differential. If we can efficiently compute the sum of the probabilities of all those characteristics, this will give a more accurate lower bound of $\Pr[\alpha \rightsquigarrow \beta]$ than by considering a single characteristic.

2.1 Related Work

Recently, a technique was proposed to find differential characteristics using Mixed Integer Linear Programming (MILP) [7,8]. When a good characteristic $\alpha \rightsquigarrow \beta$ is found, this technique can also be used to find a collection of characteristics following the same differential $\alpha \rightsquigarrow \beta$, by adding this constraint to the MILP problem. This has been applied quite successfully, but it inherently requires to enumerate all the considered characteristics.

An analysis of TWINE by Biryukov, Derbez and Perrin is also based on clustering differential characteristics [1], using the same technique as presented here. TWINE is very similar to LBlock (and LAC), but the S-Box has a more uniform differential probability which limits this clustering effect compared to our results on LAC.

2.2 Computation of the Probability of a Truncated Characteristic

In this work we use a technique to compute the probability of a collection of characteristics without having to explicitly list all the characteristics. This allows to take into account a large number of characteristics. The collection of characteristics is defined by a truncated characteristic D, which specifies whether each word is active (*i.e.* with a non-zero difference), or inactive.

We denote by $\Pr[D : \alpha \rightsquigarrow \beta]$ the probability that a pair with input difference α gives an output difference β, in a way that all the intermediate differences follow the truncated characteristic D. We also denote the reduced version of D with only i rounds as D_i. We will compute exactly $\Pr[D : \alpha \rightsquigarrow \beta]$, *i.e.* we consider the collection of *all* characteristics corresponding to the truncated characteristics, with all possible choices of non-zero values.

In order to compute $\Pr[D : \alpha \rightsquigarrow \beta]$ for a given (α, β), we will first compute $\Pr\left[D_1 : \alpha \rightsquigarrow \alpha_j^{(1)}\right]$ for all the differences $\alpha_j^{(1)}$ following D_1. Then we iteratively

build $\Pr\left[D_i : \alpha \rightsquigarrow \alpha_k^{(i)}\right]$ for all $\alpha_k^{(i)}$ following D_i using the results for D_{i-1}:

$$\Pr\left[D_i : \alpha \rightsquigarrow \alpha_k^{(i)}\right] = \sum_j \Pr\left[D_{i-1} : \alpha \rightsquigarrow \alpha_j^{(i-1)}\right] \times \Pr\left[\alpha_j^{(i-1)} \rightsquigarrow \alpha_k^{(i)}\right]$$

This corresponds to a matrix multiplication, where the matrix $M_i = \left[\Pr[\alpha_j^{(i-1)} \rightsquigarrow \alpha_k^{(i)}]\right]$ contains the probability of all the transitions $\alpha_j^{(i-1)} \rightsquigarrow \alpha_k^{(i)}$ corresponding to round i of the truncated characteristic. However, we don't explicitly perform a vector-matrix product, because the matrix is very sparse; instead we deduce the possible $\alpha_j^{(i-1)}$ with a non-zero $\Pr\left[\alpha_j^{(i-1)} \rightsquigarrow \alpha_k^{(i)}\right]$ from each $\alpha_k^{(i)}$.

2.3 Application to LAC

In order to apply this analysis to LAC, we first have to identify a good truncated characteristic. We use an automatic search for truncated characteristics, where we represent the state with a 16-bit vector, with zero and one to represent active and inactive nibbles. After computing all the possible transitions (there are at most $2^8 \times 2^{16}$ allowed transitions because of the Feistel structure), the problem of finding an optimal r-round truncated differential is reduced to the search of a shortest path in graph with $(r+1) \times 2^{16}$ nodes, and at most $(r+1) \times 2^{24}$ edges. Moreover, the graph is structured with edges only from node of round i to nodes of round $i+1$. This allows a very efficient search, round by round, with complexity $(r+1) \times 2^{24}$.

We found several truncated characteristics with 35 active S-Boxes, and we use the one given in Fig. 3. Gray square denote active nibbles. When two active nibbles are xor-ed, the truncated characteristic specifies whether the sum should be zero (slashed square) or non-zero (black square). We note that this characteristic has at most 6 active nibbles at a given round; therefore there are at most 2^{24} possible differences $\alpha_j^{(i)}$ at round i, and the vector $\left[\Pr\left[D_i : \alpha \rightsquigarrow \alpha_j^{(i)}\right]\right]$ has at most 2^{24} entries. Moreover, each step has at most 3 active S-Boxes, therefore we have at most 2^9 possible transitions to consider for any fixed $\alpha_j^{(i)}$. Using this truncated characteristic, the algorithm can compute $\Pr[D : \alpha \rightsquigarrow \beta]$ for a fixed α and for all differences β following D with at most $16 \times 2^9 \times 2^{24} = 2^{37}$ simple operations.

After running this computation with all input differences α allowed by the truncated characteristic, we identified 17512 differentials with probability higher than 2^{-64}; the best differential identified by this algorithm has a probability $\Pr[D : \alpha \rightsquigarrow \beta] \approx 2^{-61.52}$. More precisely, the best differentials found are:

$$\Pr\left[\texttt{0000000000004607} \overset{16}{\rightsquigarrow} \texttt{0000040000004400}\right] \geq 2^{-61.52}$$
$$\Pr\left[\texttt{0000000000004607} \overset{16}{\rightsquigarrow} \texttt{0000060000004400}\right] \geq 2^{-61.52}$$

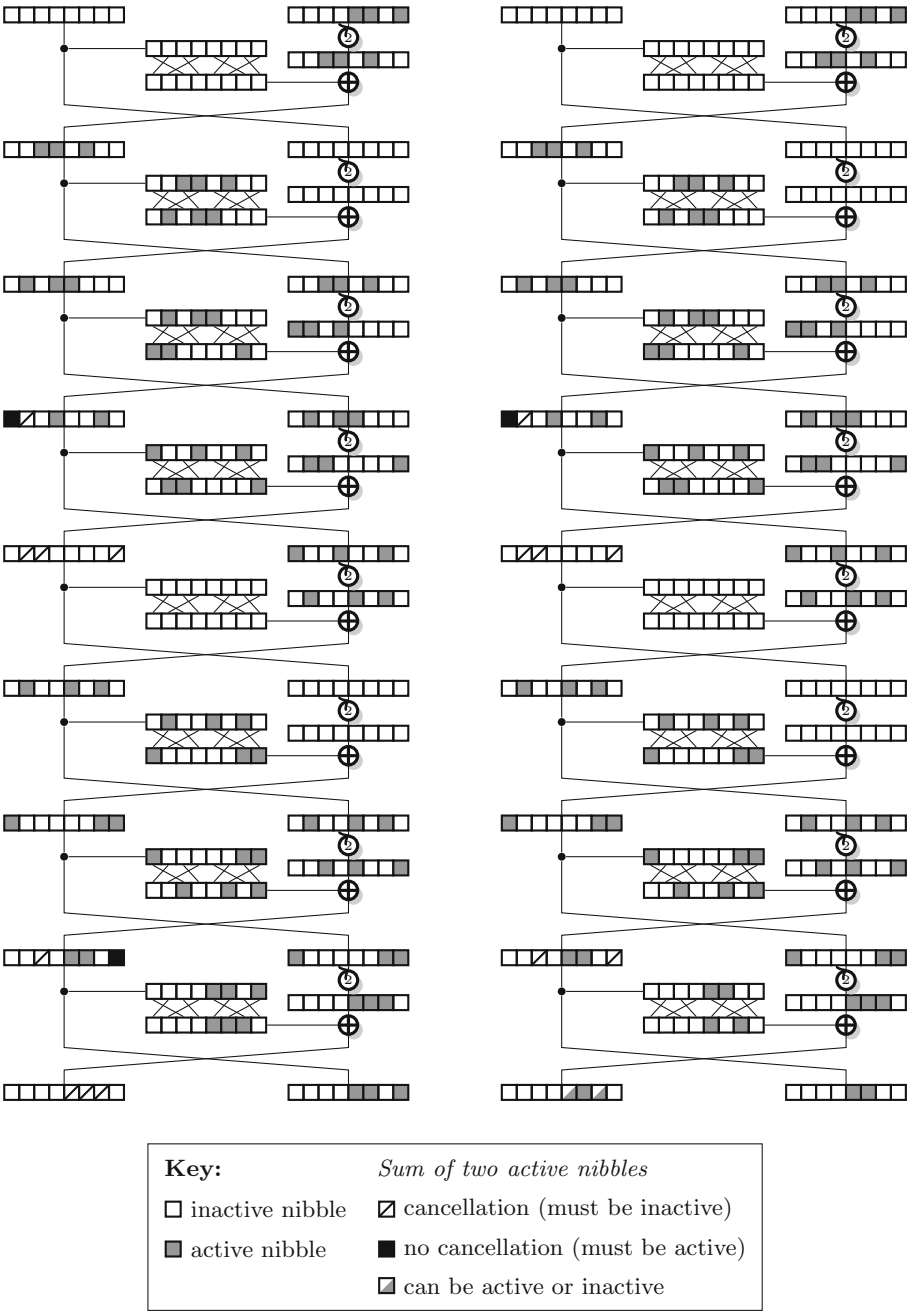

Fig. 3. Truncated characteristic for LAC with 35 active S-boxes.

These probabilities correspond to a collection of 302116704 truncated characteristics. The use of multiple characteristics allows to improve the estimation of the probability of the differential from 2^{-70} to $2^{-61.52}$.

2.4 Experimental Verification

In order to check that the algorithm is correct, we ran it with a reduced version of LAC with 8 rounds. We used the second half of the truncated differentials of Fig. 3, with 17 active S-boxes. We found that this leads to differentials with probability at least $2^{-29.76}$:

$$\Pr\left[0000000000006404 \overset{8}{\rightsquigarrow} 0000040000004400\right] \geq 2^{-29.76}$$
$$\Pr\left[0000000000006404 \overset{8}{\rightsquigarrow} 0000060000004400\right] \geq 2^{-29.76}$$

In this case, the use a multiple characteristics allows to improve the estimation of the probability of the differential from 2^{-34} (17 active S-Boxes) to $2^{-29.76}$.

For this reduced version, we ran experiments with 2^{40} random plaintext pairs following the first differential, and random round keys. We detected 1204 pairs with the expected output difference, which match very closely our prediction ($2^{40} \cdot 2^{-29.76} \approx 1209$). This indicates that our computation is correct, and the lower bound is quite tight in this case.

3 Conclusion

Our analysis shows that there exists differentials for the full G function of LAC with probability higher than 2^{-64}. This allows a simple forgery attack with probability higher than 2^{-64} on the full version of LAC, contradicting the security claims. This shows that the security margin of LAC is insufficient.

Our analysis is based on aggregating a collection of characteristics following the same truncated characteristic. While each characteristic has a probability at most 2^{-70}, a collection of characteristics can have a probability as high as $2^{-61.52}$, giving a lower bound on the probability of the corresponding differential.

Since this technique is relatively simple, we recommend all designers to check whether it can be applied to their designs.

References

1. Biryukov, A., Derbez, P., Perrin, L.: Differential analysis and meet-in-the-middle attack against round-reduced TWINE. In: Leander, G. (ed.) FSE 2015. LNCS, vol. 9054, pp. 3–27. Springer, Heidelberg (2015)
2. Bogdanov, A., Mendel, F., Regazzoni, F., Rijmen, V., Tischhauser, E.: ALE: AES-based lightweight authenticated encryption. In: Moriai, S. (ed.) FSE 2013. LNCS, vol. 8424, pp. 447–466. Springer, Heidelberg (2014)
3. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. http://competitions.cr.yp.to/caesar.html

4. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard. Information Security and Cryptography. Springer, Heidelberg (2002)
5. Keliher, L., Sui, J.: Exact maximum expected differential and linear probability for two-round Advanced Encryption Standard. IET Inf. Secur. **1**(2), 53–57 (2007)
6. Knudsen, L.R.: Truncated and higher order differentials. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1995)
7. Sun, S., Hu, L., Wang, M., Wang, P., Qiao, K., Ma, X., Shi, D., Song, L., Fu, K.: Constructing mixed-integer programming models whose feasible region is exactly the set of all valid differential characteristics of SIMON. IACR Cryptology ePrint Arch. **2015**, 122 (2015). http://eprint.iacr.org/2015/122
8. Sun, S., Hu, L., Wang, P., Qiao, K., Ma, X., Song, L.: Automatic security evaluation and (related-key) differential characteristic search: application to SIMON, PRESENT, LBlock, DES(L) and other bit-oriented block ciphers. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8873, pp. 158–178. Springer, Heidelberg (2014)
9. Wu, W., Zhang, L.: LBlock: a lightweight block cipher. In: Lopez, J., Tsudik, G. (eds.) ACNS 2011. LNCS, vol. 6715, pp. 327–344. Springer, Heidelberg (2011)
10. Zhang, L., Wu, W., Wang, Y., Wu, S., Zhang, J.: LAC: a lightweight authenticated encryption cipher. Submission to CAESAR, March 2014. http://competitions.cr.yp.to/round1/lacv1.pdf (v1)