Known-Key Distinguishers for Some Block Ciphers*

Lars R. Knudsen[1] and Vincent Rijmen[2,**]

[1] Technical University of Denmark
Department of Mathematics
Building 303S, DK-2800 Lyngby, Denmark
www.ramkilde.com
[2] Graz University of Technology
Institute for Applied Information Processing and Communications
Inffeldgasse 16a, A-8010 Graz, Austria
Vincent.Rijmen@iaik.tugraz.at

Abstract. We present two block cipher distinguishers in a setting where the attacker *knows* the key. One is a distinguisher for AES reduced the seven rounds. The second is a distinguisher for a class of Feistel ciphers with seven rounds. This setting is quite different from traditional settings. We present an open problem: the definition of a new notion of security that covers attacks like the ones we present here, but not more.

Keywords: Block Cipher, Cryptanalysis, Distinguishing algorithms, AES, Feistel ciphers.

1 Introduction

The research leading to this paper was triggered by the following example. Consider an n-bit block cipher and a plaintext/ciphertext pair for which the least significant s bits in both n-bit strings are zeros. With $s < n/2$ such a pair can be found for any reasonable block cipher in time equivalent to approximately 2^s encryptions. Imagine a block cipher where if one is given any key k, one can find such a pair for k in time much less than 2^s, but where no efficient attacks are known in the traditional black-box model. Should we recommend the use of such a cipher? We don't think so!

In the next two sections we present two attacks —or rather distinguishers— for block cipher constructions, where the attacker knows the key. Section 2 presents

* The research described in this paper has been supported by the European Commission under grant number FP6-IST-033563 (Project SMEPP). The information in this document reflects only the authors' views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

** The second author's contribution was made during a stay with the Technical University of Denmark.

K. Kurosawa (Ed.): ASIACRYPT 2007, LNCS 4833, pp. 315–324, 2007.

a distinguisher on AES reduced to seven rounds; Section 3 presents distinguishers for a class of Feistel ciphers, also with seven rounds. At the first glance it might appear strange to consider attacks on a cipher where one is given the secret key. However, by studying this type of attacks, we might learn something about the security margin of a cipher. Intuitively, it seems clear that if one cannot find distinguishers for a block cipher when given the key, then one cannot find a distinguisher where the key is secret. Secondly, in some cases (mainly for block cipher based hashing) block ciphers are used with a key that is known to the attacker, and at least to a certain extent, the key is under the attacker's control. Our attacks are quite relevant to this case.

After introducing our two attacks, we discuss related work in Section 4. We present some thoughts on a new notion of security in Section 5. We conclude in Section 6.

2 Distinguishers for Reduced AES

In this section we present known-key distinguishers for AES [1] reduced to seven (out of ten) rounds. We shall use the so-called integrals [7] to do so.

AES is an iterated cipher where in each iteration the subfunctions SubBytes, ShiftRows, MixColumns, and AddRoundKey are employed, except for the last iteration where the function MixColumns is omitted. The reason for this is that it allows the decryption routine to be implemented in a similar style to the encryption routine.

Consider a collection of 256 texts, which have different values in one byte and equal values in each of the remaining fifteen bytes. It is well-known that after two rounds of encryption the texts take all 256 values in each of the sixteen bytes, and that after three rounds of encryption the sum of the 256 bytes in each position is zero [4]. Such a structure of 256 texts is called a 3-round integral.

2.1 Notation

We introduce some notation for integrals on AES. An integral with the terms $\mathcal{A}^i$ is a collection of 2^{8i} texts. Writing $\mathcal{A}^i_j$ in a byte position means that in the integral the (string) concatenation of all bytes with subscript j take all 2^{8i} $8i$-bit values exactly once. $\mathcal{A}^i$ means that in the integral the particular byte is balanced, that is, it takes all values exactly $2^{8(i-1)}$ times. $\mathcal{C}$ means that the values in the particular byte are constant, and $\mathcal{S}$ means for the particular byte the sum of all texts can be determined. For AES addition is defined by the exclusive-or operation. The special last round of AES in integral attacks has an interesting property, namely that the balance property of an integral is preserved through this round.

2.2 Integrals for AES

It is known that there is a 3-round integral for AES using 2^{32} texts [4,5]. The main observation is that one can choose 2^{32} plaintexts as a collection of 2^{24} 2-round

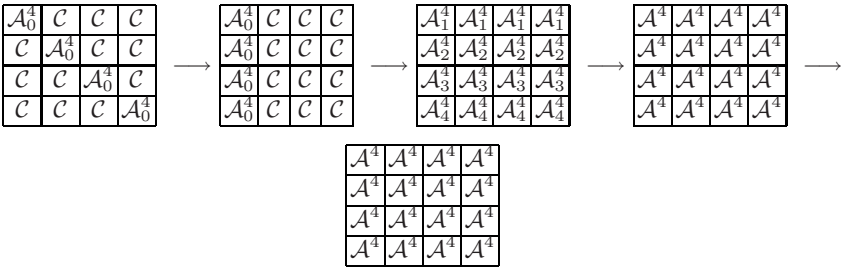

Fig. 1. An integral for 4-round AES with 2^{32} texts. The fourth round is a special round without MixColumns.

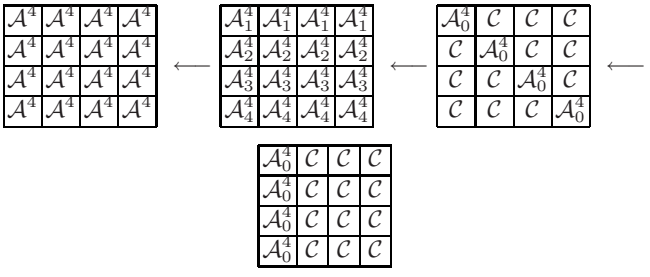

Fig. 2. A backwards integral for three (full) rounds of AES with 2^{32} texts

integrals described above (starting from the second round) each with 2^8 texts. Since the texts in each of these 2-round integrals take all values equally many times in any byte position after the third round, so does the set of all 2^{32} texts.

If we consider AES reduced to four rounds, that is, where the last round is of the special form described above, then one gets that all bytes of the ciphertexts are balanced in the 4-round integral. Figure 1 depicts this 4-round integral. Not surprisingly, one can also define integrals through the inverse cipher of AES. We present a backwards integral for three (full) rounds of AES in Figure 2. (Note the backward integral extended to four rounds does not preserve the balance property nor is it obvious to determine the sum of the texts).

The forward and backward integrals can be combined to integrals over more than four rounds of AES. One chooses a structure of 2^{56} texts which differ in seven bytes and which have constant values in the remaining nine bytes. One can view this as a collection of 2^{24} copies of the forward integral for 4-round AES, but also one can view this as a collection of 2^{24} copies of the backwards 3-round integral. Therefore, when one starts in the middle of the cipher one computes forwards and backwards for the two integrals. Next we show how to employ our findings in known-key distinguishers for AES reduced to seven rounds.

2.3 Known-Key Distinguishers for AES Reduced to Seven Rounds

Consider a variant of AES reduced to seven rounds, where MixColumns is omitted in the last round. Here one can specify the integral of Figure 3, which can

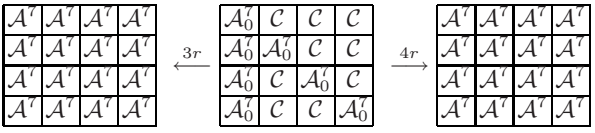

Fig. 3. An integral for 7-round AES with 2^{56} texts. The seventh round is a special round without MixColumns.

be used in an known-key distinguisher. This is constructed from the four-round integral in Figure 1 and the three-round integral of Figure 2.

The known-key distinguisher simply records the frequencies in each byte of the plaintexts and ciphertexts, checks whether the values in each byte of the plaintexts and in each byte of the ciphertexts occur equally often. The time complexity is similar to the time it takes to do 2^{56} 7-round AES encryptions and the memory needed is small.

The big question is of course, what the complexity is to find a similar structure for any 128-bit permutation. The only approach we know of, which comes close to an answer to this is the approach to solve the k-sum problem [10]. Given a function f on n bits, the k-sum problem is to find $x_1, \ldots, x_k$ such that $\sum_{i=1}^{k} f(x_i) = 0$. A solution to this problem is given in [10] with a running time of $\mathcal{O}(k2^{n/(1+\log_2 k)})$. In our case $n = 128$ and $k = 2^{56}$ indicating a running time of 2^{58} operations. However this is a very inaccurate estimation of the complexity we are looking for: the complexity estimate above is in the big $\mathcal{O}$ notation, thus ignoring smaller constants, the approach requires memory (more than for the AES distinguisher), but much more important, the k-sum problem does not give us the structure that we get for reduced AES, merely a collection of texts whose sum through the function f is zero with no conditions of balance on the values of x_i and $f(x_i)$. On the other hand, not much research has gone into finding efficient solutions for this problem. Nevertheless, we feel confident to conjecture that for a randomly chosen 128-bit permutation finding a collection of 2^{56} texts in similar time, using similar (little) memory and with similar properties as in the case of 7-round AES has a probability of succeeding which is very close to zero. Thus, we make the following claim.

Conjecture 1. There is a known-key distinguisher for AES reduced to seven rounds which uses 2^{56} texts.

We note that the above integrals might exist for a randomly chosen permutation but they are hard to find. The point we are making is that for the AES variants one finds the texts in the integrals much faster than for a randomly chosen permutation.

3 Distinguisher for a 7-Round Feistel Cipher

We present here a known-key distinguisher on an n-bit Feistel cipher with 7 rounds. The attack requires that the round function of the Feistel cipher consists

of an XOR of the round key to the round function input, followed by an arbitrary key-independent transformation. An example of a Feistel cipher with such a round function is SEED [8], but note that SEED has 16, rather than 7, rounds.

3.1 Description

The distinguisher computes (in constant time) two plaintexts denoted by $p = (p_L, p_R)$ and $\tilde{p} = (\tilde{p}_L, \tilde{p}_R)$ which have a special property. Let the corresponding ciphertexts be denoted by $c = (c_L, c_R)$ and $\tilde{c} = (\tilde{c}_L, \tilde{c}_R)$, then the following equation will hold with probability 1:

$$p_R \oplus \tilde{p}_R \oplus c_R \oplus \tilde{c}_R. = 0. \tag{1}$$

Figure 4 gives the algorithm to compute the plaintexts p and $\tilde{p}$. Note that the algorithm works only if the round keys of the second and sixth rounds are not equal. For most key schedules, such an equality happens only for a negligible fraction of the keys.

For two randomly chosen plaintexts, (1) will be satisfied with probability only $2^{-n/2}$, so we can build a strong distinguisher in this case. Also, since x can be chosen arbitrarily one can find many such pairs, thereby increasing the advantage of the distinguisher.

3.2 Conditions on the Round Function f

If f is a bijection which is easy to invert, the computations of the pair of plaintexts is straightforward. Also, note that the subkeys can be independent or

Input:
The round function of the Feistel cipher, denoted by f.
The seven subkeys $k_1, \ldots, k_7$, with $k_2 \neq k_6$.

Algorithm:

1. Choose an arbitrary value for x.
2. Define the values γ, α, z as:

$$\begin{aligned} \gamma &= k_2 \oplus k_6 \\ \alpha &= x \oplus f^{-1}(f(x) \oplus \gamma) \\ z &= f^{-1}(k_3 \oplus k_5 \oplus \alpha) \end{aligned}$$

3. Compute

$$\begin{aligned} p &= (z \oplus f(x) \oplus k_4 \oplus f(p_R, k_1), x \oplus k_3 \oplus f(z \oplus f(x) \oplus k_4 \oplus k_2)) \\ \tilde{p} &= (z \oplus f(x) \oplus \gamma \oplus k_4 \oplus f(\tilde{p}_R, k_1), x \oplus \alpha \oplus k_3 \oplus f(z \oplus f(x) \oplus k_6 \oplus k_4)). \end{aligned}$$

It follows that $p_R \oplus c_R = \alpha \oplus k_3 \oplus k_5 = f(z) = \tilde{p}_R \oplus \tilde{c}_R$, see Figure 5. Consequently, $p_R \oplus \tilde{p}_R \oplus c_R \oplus \tilde{c}_R = 0$.

Fig. 4. Algorithm to compute the plaintexts p, $\tilde{p}$ satisfying (1)

$p_L = z \oplus f(x) \oplus k_4 \oplus f(p_R \oplus k_1)$ $\qquad$ $p_R = x \oplus k_3 \oplus f(z \oplus f(x) \oplus k_2 \oplus k_4)$

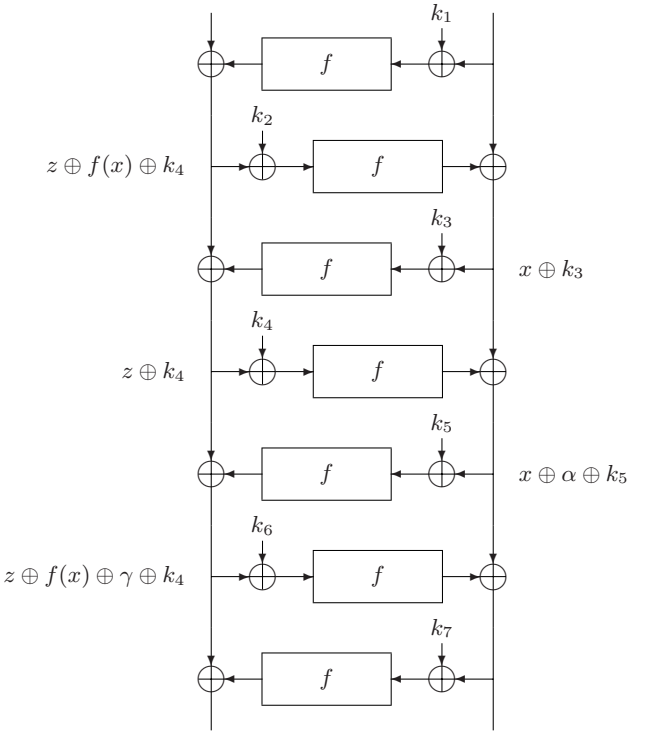

$c_L = z \oplus f(x) \oplus \gamma \oplus k_4 \oplus f(c_R \oplus k_7)$ $\qquad$ $c_R = x \oplus \alpha \oplus k_5 \oplus f(z \oplus f(x) \oplus k_2 \oplus k_4)$

Fig. 5. First encryption in 7-round Feistel cipher distinguisher. The second encryption is where x is replaced by $x \oplus \alpha$ and where $f(x)$ is replaced by $f(x) \oplus \gamma$. Notation: $\gamma = k_2 \oplus k_6$; $\alpha = x \oplus f^{-1}(f(x) \oplus \gamma)$; $z = f^{-1}(k_3 \oplus k_5 \oplus \alpha)$.

computed in a key-schedule, the only requirement we make above is that $k_2 \oplus k_6 \neq 0$. If f is not bijective, the method might still work, if inverting f is not too costly. One example is DES where given $f(w)$ is it relatively easy to find w', such $f(w) = f(w')$.

There is a variant of this attack which works for 7 rounds of Feistel ciphers where f is not bijective and where the following tasks should be "easy":

1. Find $x, y, \alpha \neq 0$ such that $f(x) = f(x \oplus \alpha) = y$,
2. Find z such that $f(z) = k_3 \oplus k_5$.

If one accomplishes these two tasks then one finds a pair of plaintexts such that (1) is satisfied. We omit the details here and refer to Appendix A.

3.3 Impact

To illustrate where the above findings could be exploited in practice consider the Matyas-Meyer-Oseas hashing mode, where the compression function is defined as

$$h(h_{i-1}, m_i) = F_{h_{i-1}}(m_i) \oplus m_i.$$

If F is a 7-round Feistel cipher construction where f is bijective, then one finds a pair of blocks which collide in half of the bits in the outputs of h doing only two encryptions.

4 Related Work

Distinguishing attacks on block ciphers where the key is known were introduced in [3] under the name *correlation intractability*. It was shown that no block cipher can be secure under this notion of security: for every block cipher, there exists a relation such that given the key, it is easy to find plaintext/ciphertext pairs satisfying this relation, but it is difficult to find them for a random permutation. The result is based on the observation that all implementable block ciphers (must) have a description, whereas a random oracle doesn't. The relation is constructed by putting the description of the block cipher in the plaintexts.

It can be argued however, that the relation of [3] is contrived. It is not clear at all how or whether such relation may lead to weaknesses in "reasonable" block-cipher based designs. Secondly, the relation is not interesting from a block cipher designer's point of view, because it applies to all implementable block ciphers. Hence, it gives no guidance on how to construct block ciphers that can be used for instance in block-cipher based hash function constructions, or in any other application where the key is known to the attacker or under her control.

5 Discussion of Known-Key Attacks

The discussion in the previous section suggests there might be a need for a new notion of security, under which the attacks presented in Section 2 and Section 3 count as valid attacks, but the general result of [3] doesn't. Indeed, the foremost idea in our mind, is to evaluate the security of specific, implementable block cipher designs and their suitability for applications which commonly use block ciphers as an underlying component.

However, it appears to be non-trivial to formalize a notion of security and at the same time avoid trivial attacks. A bullet-proof model is likely to be complicated and little transparent. Therefore, we present here some intuitions on what we think are essential elements of such a new notion of security. The introduction of the notion itself remains an open problem.

5.1 Intuitions for the Basic (Known-Key) Scenario

In this scenario, we would measure the security of the cipher against known-key attacks by computing the average advantage over all values of the key k.

A possible way to reduce the number of parasitical attacks in an informal model, would be to make the following thought exercise. Whenever we do a known-key analysis of one specific block cipher, we would rule out attacks which will succeed with approximately the same work effort on any block cipher. Hence such attacks would not change the relative ranking of the block ciphers we would examine.

5.2 Intuitions for Extended Scenarios

In a so-called *weak key* scenario, the attacker would know that the key would come out of a pre-specified subset of the whole key space. Such a scenario could reveal weak keys.

In a *related-key* scenario, we would consider scenarios where the attacker is given several different keys k_i which could have a known relation to one another. By loosening the relation between the k_is, we would eventually measure how well the block cipher would resemble a set of randomly selected permutations.

The above extensions can be illustrated using the block cipher DES. The differential attack on DES [2] uses a 13-round characteristic of average probability 2^{-47}, built from iterating a two-round characteristic of average probability $1/234$. However it is well-known that the exact probability for two rounds is either $1/146$ or $1/585$ depending on the value of one key bit. Thus by restricting ourselves to the subset of keys which provide the highest probabilities better results would be achieved. Also, if $y = DES_k(x)$ then it holds that $DES_{\overline{k}}(\overline{x}) = \overline{y}$ where $\overline{z}$ is the bitwise complemented value of z. This means that for a pair of keys (k_1, k_2), where k_1 is the bitwise complemented value of k_2 it is easy to distinguish the induced encryption functions from two randomly chosen permutations.

6 Conclusion

In this paper we presented two distinguishers for block ciphers, where the attacker is given the key. Although [3] already presented very strong results in this model, we tried to show that our attacks are still interesting from a practical security point of view, in particular when one considers block cipher applications where the key is indeed known to the attacker, e.g. block-cipher based hash functions.

Subsequently we argued that a suitable notion of security is still missing in the cryptographic literature and we presented some intuitions on how such a new notion could look like.

References

1. Specification for the Advanced Encryption Standard (AES), Federal Information Processing Standards Publication (FIPS) 197 (2001)
2. Biham, E., Shamir, A.: Differential cryptanalysis of DES-like cryptosystems. Journal of Cryptology 4(1), 3–72 (1991)
3. Canetti, R., Goldreich, O., Halevi, S.: The random oracle model, revisited. Journal of the ACM 51(4), 557–594 (2004)
4. Daemen, J., Knudsen, L., Rijmen, V.: The block cipher Square. In: Biham, E. (ed.) FSE 1997. LNCS, vol. 1267, pp. 149–165. Springer, Heidelberg (1997)
5. Ferguson, N., Kelsey, J., Lucks, S., Schneier, B., Stay, M., Wagner, D., Whiting, D.: Improved cryptanalysis of Rijndael. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, pp. 213–230. Springer, Heidelberg (2001)
6. Knudsen, L.R.: DEAL - a 128-bit block cipher. Technical Report 151, Department of Informatics, University of Bergen, Norway, February 1998. Submitted as an AES candidate by Richard Outerbridge
7. Knudsen, L.R., Wagner, D.: Integral cryptanalysis. In: Daemen, J., Rijmen, V. (eds.) FSE 2002. LNCS, vol. 2365, pp. 112–127. Springer, Heidelberg (2002)
8. Lee, H.J., Lee, S.J., Yoon, J.H., Cheon, D.H., Lee, J.I.: The SEED encryption algorithm. RFC 4269 (2005)
9. Matyas, S.M., Meyer, C.H., Oseas, J.: Generating strong one-way functions with cryptographic algorithm. IBM Technical Disclosure Bulletin 27, 5658–5659 (1985)
10. Wagner, D.: A generalized birthday problem. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 288–303. Springer, Heidelberg (2002)

A Variant Attack on a 7-Round Feistel Cipher

We present here a variant on the statistical distinguisher presented in Section 3. It works only if the following conditions are met.

1. The round function f must map at least two inputs, denoted by $x, x+\alpha$, to the same output, denoted by y. It must be possible for the attacker to determine x, y and α.
2. For most outputs, it must be easy to construct an input mapping to that output.

The distinguisher can be seen as an extension of the 5-round impossible differential presented in [6]. The transcript consists now of the plaintexts (p_L, p_R), $(\tilde{p}_L, \tilde{p}_R)$ with

$$\begin{aligned}
p_L &= z \oplus y \oplus k_4 \oplus f(x \oplus k_3 \oplus f(z \oplus y \oplus k_4 \oplus k_2) \oplus k_1),\\
p_R &= x \oplus k_3 \oplus f(z \oplus y \oplus k_4 \oplus k_2),\\
\tilde{p}_L &= z \oplus y \oplus k_4 \oplus f(x \oplus \alpha \oplus k_3 \oplus f(z \oplus y \oplus k_4 \oplus k_2) \oplus k_1),\\
\tilde{p}_R &= p_R \oplus \alpha,
\end{aligned}$$

and the corresponding ciphertexts. Here z is defined by $f(z) = k_3 \oplus k_5$. We discuss below what to do if no such z exists. The test is again: verify whether

$$p_R + \tilde{p}_R = c_R + \tilde{c}_R. \tag{2}$$

If it is not possible to find a z such that $f(z) = k_3 \oplus k_5$, then we can search for a z' such that $f(z') = k_3 \oplus k_5 \oplus \alpha$. We can then construct a plaintext pair such that in the first text the inputs to f in round three and five are x, respectively $x \oplus \alpha$, and in the second pair $x \oplus \alpha$, respectively x. This pair will also satisfy (2). Finally, if also this is not possible, there might be another difference α that can be used.