ECOLE NORMALE SUPERIEURE

On the Weak Keys of Blowfish

Serge VAUDENAY

LIENS - 95 - 27

Département de Mathématiques et Informatique

CNRS URA 1327

On the Weak Keys of Blowfish

Serge VAUDENAY

LIENS - 95 - 27

November 1995

Laboratoire d'Informatique de l'Ecole Normale Supérieure
45 rue d'Ulm 75230 PARIS Cedex 05

Tel : (33)(1) 44 32 00 00
Adresse électronique : vaudenay @dmi.ens.fr

On the Weak Keys of Blowfish

Serge Vaudenay*

Ecole Normale Supérieure — DMI
45, rue d'Ulm
75230 Paris Cedex 5 France
Serge.Vaudenay@ens.fr

Abstract

Blowfish is a sixteen-rounds Feistel cipher in which the F function is a part of the private key. In this paper, we show that the disclosure of F allows to perform a differential cryptanalysis which can recover all the rest of the key with 2^{48} chosen plaintexts against a number of rounds reduced to eight. Moreover, for some weak F function, this attack only needs 2^{23} chosen plaintexts against eight rounds, and 3×2^{51} chosen plaintexts against sixteen-rounds. When the F function is safely kept private, one can detect whether it is weak or not with a differential attack using 2^{22} plaintexts against eight rounds.

Blowfish was proposed by Schneier in the Cambridge Security Workshop [6]. It appears to be a very fast encryption function when we always use the same private key. It is based on the Feistel cipher [4]. Differential cryptanalysis [3] is known to be one of the most powerful attack on this kind of cipher. The design of Blowfish includes the new feature that the s-boxes are randomly generated from the private key. Hence, for some particular *weak keys*, a differential cryptanalysis may be successful. This paper shows the first analysis of Blowfish, as an answer to the Dr Dobb's Journal Blowfish Cryptanalysis Contest proposed in [7].

1 Blowfish

Blowfish encrypts a 64-bit plaintext into a 64-bit ciphertext using a variable key length [6]. The encryption proceeds with a suggested number of

Laboratoire d'Informatique de l'Ecole Normale Supérieure, research group affiliated with the CNRS

$t = 16$ rounds. In the following, we may consider a smaller number of rounds t. First, the key is expanded into a 4168-bytes key following a scheduling scheme which works as a pseudo-random generator. As this scheme is very complicated, one has to store definitely the expanded key. Thus, it is reasonable to consider the expanded key as the *real* key in the attack.

The expanded key consists of:

- $t+2$ 32-bit constants $P_1, \ldots, P_{t+2}$;
- four arrays of 256 32-bit values which describe four s-boxes $S_1, \ldots, S_4$ with 8-bit inputs and 32-bit outputs.

The four s-boxes define a 32-bit to 32-bit function F by

$$F([abcd]) = ((S_1(a) + S_2(b)) \oplus S_3(c)) + S_4(d)$$

where $\oplus$ is the bit-wise xor and $+$ is the addition modulo 2^{32} and $[abcd]$ is the concatenated bit string of the four 8-bit strings a, b, c and d.

The plaintext $\mathcal{P} = (L_0, R_0)$ is divided into two 32-bit halves. Each round is defined recursively following the Feistel scheme by

$$R_i = P_i \oplus L_{i-1} \text{ and } L_i = R_{i-1} \oplus F(R_i).$$

the ciphertext is $\mathcal{C} = (R_t \oplus P_{t+2}, L_t \oplus P_{t+1})$ (the left and right registers are not exchanged for the final round).

2 Known F - weak key attack

All through this paper, the term *weak key* means there exists an s-box, say S_1, which has a collision. That is to say, there exist two different bytes a and a' such that $S_1(a) = S_1(a')$. In this section, we assume the opponent knows the part of the private key which describes the F function, that is the four s-boxes. (in fact, we only need to know a and a' to recover seven bits of information on the private key.)

Let δ denote the xor-difference of the collision of S_1 (that is $\delta = a \oplus a'$ with the previous notation) We consider the following iterative characteristic.

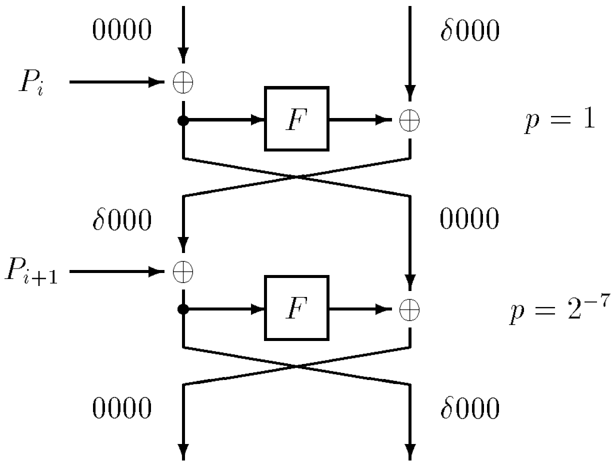

(Throughout this paper, figures represent 8-bit values so that [$\delta 000$] is a 32-bit value.) Assuming there is only one collision for S_1 with difference δ, the probability of this characteristic is 2^{-7}.

For Blowfish reduced to $t = 8$ rounds, we iterate this characteristic three times as shown on figure 1 ($xyzt$ represents an undetermined value). The resulting characteristic has probability 2^{-21}. Hence, trying 2^{21} chosen plaintext pairs with xor [$0000\delta 000$], we easily detect a ciphertext pair $(\mathcal{C}, \mathcal{C}')$ with xor [$\delta 000xyzt$]. Note that for a random pair, the probability of getting such an xor is 2^{-32}. So, a detected pair with the good xor is certainly a good pair. With such a pair, let denote $\mathcal{C} = (L, R)$. Since we have

$$F(L \oplus P_{10}) \oplus F(L \oplus P_{10} \oplus [\delta 000]) = [xyzt]$$

we can try exhautively all the 2^{32} possible P_{10} until this equation holds. It is easy to check that Blowfish with t rounds and a known P_{t+2} is equivalent to Blowfish with $t-1$ rounds. So, this attack allows to reduce the cipher to $t = 7$ rounds.

More generally, for Blowfish with t rounds, the same iterated characteristic has probability $2^{-7 \times \lceil \frac{t-2}{2} \rceil}$. So, we can use $2^{7 \times \lceil \frac{t-2}{2} \rceil}$ chosen plaintext pairs, and, for each ciphertext pair with xor [$\delta 000xyzt$], list the possible values of P_{t+2}. A random pair has this xor with probability 2^{-32}, and each such pair suggests one value of P_{10} on average.

For $t \leq 10$, all pairs which have this xor are good, but for $t \geq 11$, we may get $2^{7 \times \lceil \frac{t-2}{2} \rceil - 32}$ wrong pairs. Each wrong pair suggests one random value for

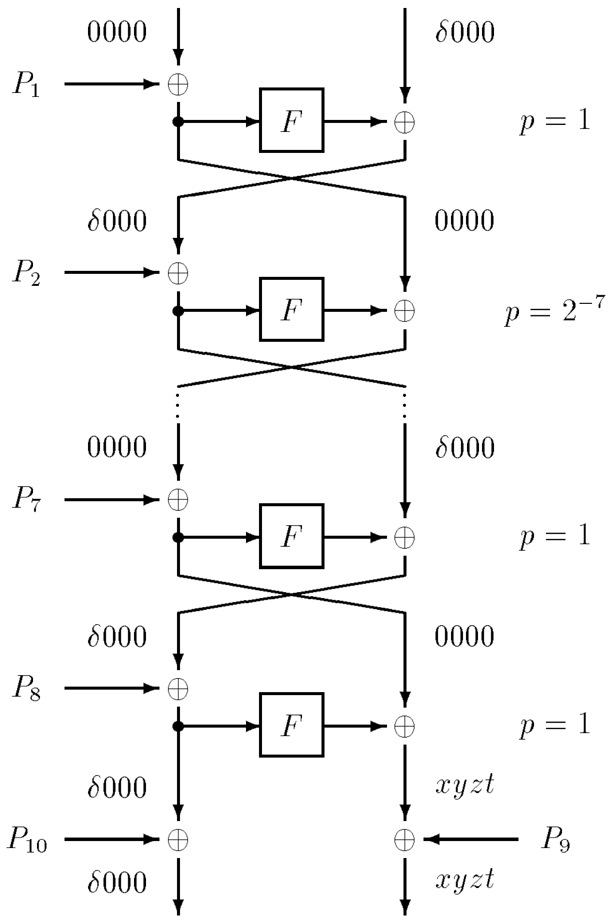

Figure 1: Characteristic for 8 rounds

P_{t+2} on average. So, trying $3 \times 2^{7\times\lceil\frac{t-2}{2}\rceil}$ chosen plaintext pairs, we get three good pairs which suggest the same value with high probability, and no other value may be suggested more than two times for $t \leq 16$.

As the complexity of the attack on $t-1$ rounds is the same as for t rounds if t is even, the number of chosen plaintexts required is $3 \times 2^{2+7\times\lceil\frac{t-2}{2}\rceil}$. For $t = 16$, this is 3×2^{51}.

For $t = 8$, since there is no problem with wrong pairs, the number of plaintexts required is 2^{23}.

3 Known F - random key attack

As in the previous section, we assume the description of the F function has been disclosed, but the private key is not necessarily supposed to be weak. The mapping $(a, b) \mapsto S_1(a) + S_2(b)$ is a 16 to 32 bits function, so it may have a collision $S_1(a) + S_2(b) = S_1(a') + S_2(b')$ with high probability.

Letting $\delta = a \oplus a'$ and $\mu = b \oplus b'$, we consider the following iterative characteristic.

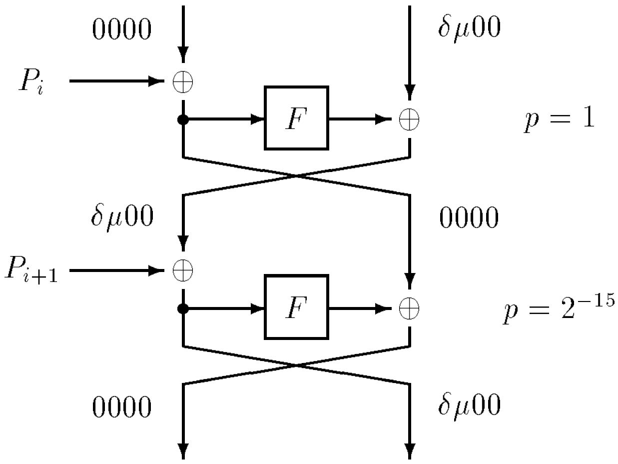

Assuming there is only one collision for $S_1 + S_2$ with difference $\delta\mu$, the probability of this characteristic is 2^{-15}. Hence, for Blowfish with $t = 8$ rounds, this characteristic iterated as on figure 1 has probability 2^{-45}, and 2^{46} chosen plaintext pairs include two good pairs and 2^{14} wrong pairs. So, the good value of P_{10} may be the only value suggested twice.

The attack on $t = 7$ rounds has the same complexity, so the number of plaintexts required is 2^{48}.

4 Weak key detection

What can we do without the description of F? We can try to detect whether a key is weak or not. For a random s-box S_1, the probability that there is no collision is

$$\prod_{i=0}^{2^8-1}\left(1-\frac{i}{2^{32}}\right)=\frac{2^{32}!}{2^{32\times 2^8}(2^{32}-2^8)!}\approx 1-2^{-17.0}.$$

So, for a F function made with random s-boxes, the probability there exists a collision for one s-box is $2^{-15.0}$. Hence, one key out of 2^{15} may be weak. Now let us see how to distinguish which is weak by a chosen plaintext attack.

First one can bet on S_1 and use the characteristic on figure 1. If we pick the bytes B_1, B_2, B_3, B_4, B_6, B_7 and B_8 at random (no B_5 here), in the *structure* of all the 2^8 plaintexts

$$\mathcal{P}=[B_1B_2B_3B_4B_5B_6B_7B_8]$$

there are 2^7 pairs with the good xor. Let

$$\mathcal{C}=[C_1C_2C_3C_4C_5C_6C_7C_8]$$

be the corresponding ciphertexts. In a good pair, we notice that

$$\mathcal{X}=[(B_5\oplus C_5)C_6C_7C_8]$$

must be the same for both messages of the pair. Thus, we can seek for pairs in the structure which makes $\mathcal{X}$ collide. If there are no good pairs, this occurs with probability roughly $2^{-17.0}$. Trying 2^{14} structures, we get one good pair with high probability, and no wrong pair with probability roughly 2^{-3}. Hence, with 2^{22} chosen plaintexts, we can detect a collision on S_1 and get the xor δ of the collision. The same attack holds for S_2, S_3 and S_4.

5 Conclusion

We have shown differential cryptanalysis on Blowfish is possible either against a reduced number of rounds or with the piece of information which describes the F function. This second case appears to be equivalent to an analysis done by Lee, Heys and Tavares [5] against the CAST cipher [1, 2]. In the analysis of CAST, the s-boxes are well design to resist to any attack while they are randomly generated in Blowfish. Compared to CAST, some of the s-boxes generated by Blowfish may be really weak, but it is not sure whether it is sufficient to mount an attack since they are supposed to be private.

We studied weaknesses of the s-boxes based on collisions. This way, we proved there are weak keys in Blowfish that enable to decrease significantly the complexity of the attacks (from 2^{48} to 2^{23} on eight rounds when F is known). We also showed it is possible to detect weak keys using 2^{22} chosen plaintexts (on eight rounds).

Acknowledgments

Many thanks to Lars Knudsen for his help and fruitful discussions.

References

[1] C. M. Adams. *A Formal and Practical Design Procedure for Substitution-Permutation Network Cryptosystems.* PhD thesis, Queen's University, Kingston, Canada, 1990.

[2] C. M. Adams, S. E. Tavares. Designing s-boxes Resistant to Differential Cryptanalysis. In *Proceedings of 3rd Symposium on the State and Progress of Research in Cryptography*, pp. 386–397, Rome, Italy, 1994.

[3] E. Biham, A. Shamir. *Differential Cryptanalysis of the Data Encryption Standard*, Springer-Verlag, 1993.

[4] H. Feistel. Cryptography and computer privacy. In *Scientific American*, vol. 228, pp. 15–23, 1973.

[5] J. Lee, H. M. Heys, S. E. Tavares. On the Resistance of the CAST Encryption Algorithm to Differential Cryptanalysis. Presented at the SAC'95 conference.

[6] B. Schneier. Description of a New Variable-Length Key, 64-Bit Block Cipher (Blowfish). In *Fast Software Encryption – Proceedings of the Cambridge Security Workshop*, Cambridge, United Kingdom, Lectures Notes in Computer Science 809, pp. 191–204, Springer-Verlag, 1994.

[7] B. Schneier. The Blowfish Encryption Algorithm. In *Dr Dobb's Journal*, pp. 38–40, April 1994.