

Mitsuru Matsui (Ed.)

LNCS 2355

Fast Software Encryption

8th International Workshop, FSE 2001
Yokohama, Japan, April 2001
Revised Papers



Springer

Mitsuru Matsui (Ed.)

Fast Software Encryption

8th International Workshop, FSE 2001

Yokohama, Japan, April 2-4, 2001

Revised Papers



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editor

Mitsuru Matsui
Mitsubishi Electric Corporation
5-1-1 Ofuna Kamakura Kanagawa, 247-8501, Japan
E-mail: matsui@iss.isl.melco.co.jp

Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek - CIP-Einheitsaufnahme

Fast software encryption : 8th international workshop ; proceedings / FSE
2001, Yokohama, Japan, April 2 - 4, 2001. Mitsuru Matsui (ed.). - Berlin ;
Heidelberg ; New York ; Barcelona ; Hong Kong ; London ; Milan ; Paris ;
Tokyo : Springer, 2002
(Lecture notes in computer science ; Vol. 2355)
ISBN 3-540-43869-6

CR Subject Classification (1998): E.3, F.2.1, E.4, G.4

ISSN 0302-9743

ISBN 3-540-43869-6 Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2002
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP Berlin, Stefan Sossna e. K.
Printed on acid-free paper SPIN 10870083 06/3142 5 4 3 2 1 0

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Tokyo

Preface

Fast Software Encryption is an eight-year-old workshop on symmetric cryptography, including the design and cryptanalysis of block and stream ciphers, as well as hash functions. The first Fast Software Encryption Workshop was held in Cambridge in 1993, followed by Leuven in 1994, Cambridge in 1996, Haifa in 1997, Paris in 1998, Rome in 1999, and New York in 2000. This Fast Software Encryption Workshop, FSE 2001, was held from 2-4 April 2001 in Yokohama, Japan, in cooperation with the Institute of Industrial Science, of the University of Tokyo.

This year a total of 46 papers were submitted to FSE 2001. After a two-month review process, 27 papers were accepted for presentation at the workshop. In addition, we were fortunate to be able to organize a special talk by Bart Preneel on the NESSIE project, a European initiative to evaluate cryptographic algorithms. The committee of this workshop was:

General Chair

Hideki Imai (The University of Tokyo)

Program Committee

Ross Anderson (Cambridge Univ.)

Cunsheng Ding (Singapore)

Dieter Gollman (Microsoft)

Lars Knudsen (Bergen Univ.)

Mitsuru Matsui (Mitsubishi Electric, Chair)

Bart Preneel (Katholieke Univ. Leuven)

Eli Biham (Technion)

Henri Gilbert (France Telecom)

Thomas Johansson (Lund Univ.)

James Massey (Denmark)

Kaisa Nyberg (Nokia)

Bruce Schneier (Counterpane)

We would like to thank all submitting authors and the committee members for their hard work. We are also appreciative of the financial support provided by Mitsubishi Electric Corporation. Special thanks are due to Toshio Tokita, Junko Nakajima, Yasuyuki Sakai, Seiichi Amada, Toshio Hasegawa, Katsuyuki Takashima, and Toru Sorimachi for their efforts in making the local arrangements for this workshop.

We were very pleased and honored to host the first FSE workshop held in Asia. Finally, we are also happy to announce that the next FSE will be the first FSE workshop sponsored by International Association for Cryptologic Research (IACR).

May 2002

Mitsuru Matsui

Table of Contents

Cryptanalysis of Block Ciphers I

The Saturation Attack – A Bait for Twofish	1
<i>Stefan Lucks</i>	
Linear Cryptanalysis of Reduced Round Serpent	16
<i>Eli Biham, Orr Dunkelman, Nathan Keller</i>	
Cryptanalysis of the Mercy Block Cipher	28
<i>Scott R. Fluhrer</i>	

Hash Functions and Boolean Functions

Producing Collisions for PANAMA	37
<i>Vincent Rijmen, Bart Van Rompay, Bart Preneel, Joos Vandewalle</i>	
The RIPEMD ^L and RIPEMD ^R Improved Variants of MD4 Are Not Collision Free	52
<i>Christophe Debaert, Henri Gilbert</i>	
New Constructions of Resilient Boolean Functions with Maximal Nonlinearity	66
<i>Yuriy Tarannikov</i>	

Modes of Operations

Optimized Self-Synchronizing Mode of Operation	78
<i>Ammar Alkassar, Alexander Geraudy, Birgit Pfitzmann, Ahmad-Reza Sadeghi</i>	
Fast Encryption and Authentication: XCBC Encryption and XECB Authentication Modes	92
<i>Virgil D. Gligor, Pompiliu Donescu</i>	
Incremental Unforgeable Encryption	109
<i>Enrico Buonanno, Jonathan Katz, Moti Yung</i>	

Cryptanalysis of Stream Ciphers I

ZIP Attacks with Reduced Known Plaintext	125
<i>Michael Stay</i>	
Cryptanalysis of the SEAL 3.0 Pseudorandom Function Family	135
<i>Scott R. Fluhrer</i>	

Cryptanalysis of SBLH	144
<i>Goce Jakimovski, Ljupčo Kocarev</i>	

A Practical Attack on Broadcast RC4	152
<i>Itsik Mantin, Adi Shamir</i>	

Cryptanalysis of Block Ciphers II

Improved SQUARE Attacks against Reduced-Round HIEROCRYPT	165
<i>Paulo S.L.M. Barreto, Vincent Rijmen, Jorge Nakahara, Bart Preneel, Joos Vandewalle, Hae Y. Kim</i>	

Differential Cryptanalysis of Q.....	174
<i>Eli Biham, Vladimir Furman, Michal Misztal, Vincent Rijmen</i>	

Differential Cryptanalysis of Nimbus	187
<i>Vladimir Furman</i>	

Cryptanalysis of Stream Ciphers II

Fast Correlation Attack Algorithm with List Decoding and an Application	196
<i>Miodrag J. Mihaljević, Marc P.C. Fossorier, Hideki Imai</i>	

Bias in the LEVIATHAN Stream Cipher	211
<i>Paul Crowley, Stefan Lucks</i>	

Analysis of SSC2.....	219
<i>Daniel Bleichenbacher, Willi Meier</i>	

Pseudo-Randomness

Round Security and Super-Pseudorandomness of MISTY Type Structure	233
<i>Tetsu Iwata, Tomonobu Yoshino, Tomohiro Yuasa, Kaoru Kurosawa</i>	

New Results on the Pseudorandomness of Some Blockcipher Constructions	248
<i>Henri Gilbert, Marine Minier</i>	

FSE 2001 Special Talk

NESSIE: A European Approach to Evaluate Cryptographic Algorithms ...	267
<i>Bart Preneel</i>	

Cryptanalysis of Block Ciphers III

Related Key Attacks on Reduced Round KASUMI	277
<i>Mark Blunden, Adrian Escott</i>	

Security of Camellia against Truncated Differential Cryptanalysis	286
<i>Masayuki Kanda, Tsutomu Matsumoto</i>	
Impossible Differential Cryptanalysis of Zodiac	300
<i>Deukjo Hong, Jaechul Sung, Shiho Moriai, Sangjin Lee, Jongin Lim</i>	
Design and Evaluation	
The Block Cipher SC2000	312
<i>Takeshi Shimoyama, Hitoshi Yanami, Kazuhiro Yokoyama, Masahiko Takenaka, Kouichi Itoh, Jun Yajima, Naoya Torii, Hidema Tanaka</i>	
Flaws in Differential Cryptanalysis of Skipjack	328
<i>Louis Granboulan</i>	
Efficient Algorithms for Computing Differential Properties of Addition	336
<i>Helger Lipmaa, Shiho Moriai</i>	
Author Index	351

The Saturation Attack – A Bait for Twofish

Stefan Lucks*

Theoretische Informatik

University of Mannheim, 68131 Mannheim, Germany

luck@th.informatik.uni-mannheim.de

Abstract. This paper introduces the notion of a “saturation attack”. Consider a permutation p over w -bit words. If p is applied to all 2^w disjoint words, the set of outputs is exactly the same as the set of inputs. A saturation attack exploits this fact. The current paper applies saturation attacks on reduced-round variants of the Twofish block cipher with up to seven rounds with full whitening or eight rounds without whitening at the end (i.e., half of the cipher). The attacks take up to 2^{127} chosen plaintexts (half of the codebook) and are 2–4 times faster than exhaustive search. The attacks are based on key-independent distinguishers for up to six rounds of Twofish, making extensive use of saturation properties.

1 Introduction

Modern b -bit block ciphers often use permutations $p : \{0, 1\}^w \rightarrow \{0, 1\}^w$ with $w < b$ as building blocks. E.g., p may be an S-box, a round function, or a group operation where one of the operands is constant. The constant may be unknown to the cryptanalyst, e.g. as a part of the (round) key. We regard the input for p as a data channel. For the cryptanalyst, p may be known or unknown, and the cryptanalysts may be unable to determine the input for p . A “saturation attack” is based on the idea of choosing a set of $k * 2^w$ plaintexts such that each of the 2^w inputs for p occurs exactly k times. In this case, we say that the data channel into p is “saturated”. A saturation attack exploits the fact that if the input for p is saturated, then the output from p is saturated, too.

The name “saturation attack” is new, but such attacks have been studied before. E.g., the “Square attack” is a saturation attack, developed for the block cipher Square [4]. It works as well for other Square-like ciphers such as the AES candidate Crypton [11,12] and the finalist Rijndael [5], which has recently been chosen as the AES. All these ciphers are 128-bit block ciphers with 8-bit data channels. The attack starts with a set of 2^8 plaintexts with one saturated channel. The other 15 channels are constant. After two rounds, all 16 data channels are saturated. After three rounds, the saturation property is likely to have been lost, but the sum of all values in a data channel is zero. This allows to distinguish the three-round output from random. The best currently known attacks on Crypton [3] and Rijndael/AES [7] are extensions of the Square attack.

* Supported by German Science Foundation (DFG) grant KR 1521/3-2.

“Miss in the middle” attacks [1] are rudimentarily related to saturation attacks, exploiting the fact that given two inputs $x \neq y$ for a permutation p one gets two outputs $p(x)$ and $p(y)$ with $p(x) \neq p(y)$. Also related is the attack on “Ladder DES”, based on choosing $c \cdot 2^{32}$ distinct inputs for a 64-bit data channel and checking if all the outputs are distinct [2].

When using “higher-order differentials” [9], one chooses a certain complete set of plaintexts and, after some rounds of the cipher, predicts a key-independent property with probability one. This resembles the current approach.

This paper shows that saturation attacks are a useful tool for ciphers which are definitely not Square-like. We concentrate on the AES finalist Twofish [15]. So far, the authors of Twofish published some preliminary cryptanalytic results [16,6] themselves, a key separation property has been identified for Twofish [13, 14,8], and some observations on the generation of the Twofish S-Boxes and on differential cryptanalysis have been made [10].

The motivation for this research has been twofold. First, even though Twofish has not been chosen as the AES, it is (and probably will continue to be) used in practice. E.g., recent versions of popular email encryption programs, namely PGP and GnuPG [17], implement Twofish. Second, the study of saturation attacks appears to be of independent interest in cryptanalysis.

1.1 Notation

We will use the notion of a “multiset” to describe a w -bit data channel. A multiset with $k \cdot 2^w$ entries is “saturated” if every value in $\{0,1\}^w$ is found exactly k times in the multiset. If $k = 1$, a saturated multiset is the set $\{0,1\}^w$.

In the context of this paper, a data channel is always 32 bits wide, and we call a value in a data channel a “word”. We interchangeably view a word x as a 32-bit string $x = (x_{31}, \dots, x_0) \in \{0,1\}^{32}$ and as an unsigned integer $x = \sum_i x_i \cdot 2^i$. The addition of values in a data channel is thus addition mod 2^{32} . We write “ $x \ll b$ ” for the rotation of the word x by b bits to the left, and “ $x \gg b$ ” for rotation to the right. E.g. $(x \ll b) \gg b = x$ for all x and b , and $(x_{31}, x_{30}, \dots, x_1, x_0) \ll 1 = (x_{30}, \dots, x_1, x_0, x_{31})$. $\text{LSB}(x) = x \bmod 2$ denotes the “least significant bit (LSB)” of x , and $\text{LSB}^1(x) = \text{LSB}(x \text{ div } 2)$ denotes the 2nd-least significant bit. Similarly, we define the “most significant bit (MSB)”: $\text{MSB}(x) = \text{LSB}(x \ll 1)$. If the multiset M denotes a data channel, the bits at the LSB-position of M are “balanced” if $\bigoplus_{m \in M} \text{LSB}(m) = 0$. It turns out to be useful to also consider “semi-saturated” data channels. The multiset M is semi-saturated if one bit of M is constant and each of the 2^{31} remaining values for M appears exactly $2k$ times in M .

2 A Description of Twofish

In this section, we describe the structure of Twofish. We omit many details, concentrating on the properties of Twofish which are relevant for our attack.

2.1 The Main Operations of Twofish

Twofish is based on the following operations:

Whitening. Decompose a 128-bit text block into words $a_0, \dots, a_3 \in \{0, 1\}^{32}$. The Twofish whitening operation is the XOR of four key words $K_{j+\delta} \in \{0, 1\}^{32}$ to the words a_j : $b_j := a_j \oplus K_{j+\delta}$ for $j \in \{0, \dots, 3\}$, see Figure 1.




Fig. 1. The Twofish Whitening Operation.

Application of the round function. To compute the i -th round function F_i , use a pair $(a, b) \in (\{0, 1\}^{32})^2$ as the input and compute a pair $(a', b') = F_i(a, b) \in (\{0, 1\}^{32})^2$. The round function F_i is defined by two round keys K_{2i+2} and K_{2i+3} and two functions $G_1, G_2 : \{0, 1\}^{32} \rightarrow \{0, 1\}^{32}$, see Fig. 2:

$$a' := G_1(a) + G_2(b) + K_{2i+2}, \quad \text{and} \quad b' := G_1(a) + 2G_2(b) + K_{2i+3},$$

The functions G_1 and G_2 are key-dependent, but do not depend on i . Given the round function's results a' and b' , the remaining two words $c, d \in \{0, 1\}^{32}$ come into play:

$$x := (a' \oplus c) \ggg 1, \quad \text{and} \quad y := b' \oplus (d \lll 1).$$

Except for the rotate operations, Twofish works like a Feistel cipher.




Fig. 2. The Application of the Twofish Round Function.

Our attack greatly depends on the functions G_1 and G_2 to be permutations over $\{0, 1\}^{32}$. Actually, $G_2(x) = G_1(x \lll 8)$. Apart from that, the internal structure of G_1 and G_2 is not relevant for us.

The swap. Replace $(a, b, c, d) \in (\{0, 1\}^{32})^4$ by (c, d, a, b) . See Figure 3.




Fig. 3. The Twofish Swap.

2.2 The Basic Structure of Twofish

Twofish uses a 16-round Feistel structure with two additional one-bit rotates in each round, pre-whitening before the first round and post-whitening after the last round. Twofish works as follows:

1. Generate the key-dependent S-boxes, which define the functions G_1 and G_2 .
2. Generate four subkey words $K_0, \dots, K_3 \in \{0, 1\}^{32}$ for the pre-whitening, two subkey words K_{2i+2}, K_{2i+3} for each round and another four subkey words K_{36}, \dots, K_{39} for the post-whitening.
3. Given a plaintext block, do the pre-whitening.
4. For $i := 1$ to 15 do: (a) Apply the round function F_i .
(b) Do the swap.
5. Apply the last round function F_{16} (no swap in the final round).
6. Do the post-whitening.

The first two of the above steps constitute the “key schedule” described below. Note that we can obviously generalise the Twofish structure to r rounds, where the loop in step 4 is iterated $r - 1$ times.

2.3 The Twofish Key Schedule

A Twofish key consists of 128, 192, or 256 bit¹: $2k$ words $M_0, \dots, M_{2k-1} \in \{0, 1\}^{32}$ with $k \in \{2, 3, 4\}$, organised as two vectors $M_e = (M_0, M_2, \dots, M_{2k-2})$ and $M_o = (M_1, M_3, \dots, M_{2k-1})$. A third vector $S = (S_0, S_1, \dots, S_{k-1})$ is derived from M_e and M_o by using techniques from the theory of Reed-Solomon codes. Given any two of the three vectors M_e , M_o and S , the third one is easy to find.

With these three vectors, the “three halves of a Twofish key”, we can do the first two steps of the structure described above:

1. The vector S determines the internal S-boxes and thus the functions G_1 and G_2 . S is a k -word vector, while the key consists of $2k$ words or $64k$ bit.
2. The 40 subkey words K_0, \dots, K_{39} are defined by using functions h_e and h_o and by doing 20 “subkey generation” steps ($j \in \{0, \dots, 19\}$):

$$\begin{aligned} A_j &:= h_e(j, M_e); & K_{2j} &:= A_j + B_j; \\ B_j &:= h_o(j, M_o); & K_{2j+1} &:= (A_j + 2B_j) \lll 9. \end{aligned}$$

¹ These are the three generic key lengths of Twofish. Other keys of less than 256 bit are padded to the next generic length by appending zeros.

3 Distinguishers for Twofish

Given a well-chosen set of plaintexts, we describe how to distinguish reduced-round versions of Twofish from random permutations.

3.1 A Four-Round Distinguisher

Consider 2^{32} plaintexts $(\alpha_0, \alpha_1, A, \alpha_3)$, where α_0, α_1 , and α_3 are three arbitrary 32-bit constants and A is the set of all 2^{32} words. The pre-whitening changes this set of texts to $(\beta_0, \beta_1, A, \beta_3)$ with new constants β_i .

Given this set of texts as the input for the first round, the input for the round function F_1 is constant: (β_0, β_1) . By (γ_0, γ_1) we denote the output of F_1 , which then generates the texts $(\beta_0, \beta_1, A, \gamma_3)$ with $\gamma_3 = (\beta_3 \lll 1) \oplus \gamma_1$. (Note that $A = \{a_i\} = \{(a_i \oplus \gamma_0) \ggg 1\}$.) The swap changes these texts to $(A, \gamma_3, \beta_0, \beta_1)$.

In the second round, the 2^{32} inputs for the round function are (A, γ_3) . The round function generates the pairs (b_i, c_i) with $b_i = G_1(a_i) + G_2(\gamma_3) + K_6$ and $c_i = G_1(a_i) + 2G_2(\gamma_3) + K_7$ for $a_i \in A$. The sets $B = \{b_i\}$ and $C = \{c_i\}$ are saturated, just like A . Applying the round function here means XORing the constant β_0 to the values of B , followed by a rotation, and XORing $\beta_1 \lll 1$ to C . Neither operation changes the saturated sets B and C . We get 2^{32} texts (A, γ_3, B, C) , where A, B , and C are saturated. By the swap, we get texts (B, C, A, γ_3) .

The 2^{32} inputs for the third round function are of the form (B, C) with saturated B and C . Since both G_1 and G_2 are permutations, $G_1(b_i) \neq G_1(b_j)$ and $G_2(c_i) \neq G_2(c_j)$ for $b_i, b_j \in B, c_i, c_j \in C$, and $i \neq j$. Let $d_i = G_1(b_i) + G_2(c_i) + K_8$ and $e_i = G_1(b_i) + 2G_2(c_i) + K_9$. The 2^{32} outputs of the round function are of the form (D, E) , with the multisets $D = \{d_i | 0 \leq i < 2^{32}\}$ and $E = \{e_i | 0 \leq i < 2^{32}\}$. Neither D nor E is likely to be saturated. However, we are still able to observe a weaker property: Since $\sum_{b_i \in B} b_i = \sum_{c_i \in C} c_i = \sum_{0 \leq i < 2^{32}} i \equiv 2^{31} \pmod{2^{32}}$:

$$\sum_{0 \leq i < 2^{32}} d_i \equiv 2^{31} + 2^{31} + 2^{32} * K_8 \equiv 0 \pmod{2^{32}},$$

$$\sum_{0 \leq i < 2^{32}} e_i \equiv 2^{31} + 2 * 2^{31} + 2^{32} * K_9 \equiv 2^{31} \pmod{2^{32}},$$

thus $\sum d_i \equiv \sum e_i \equiv 0 \pmod{2}$ – i.e., the LSBs of D and E are *balanced*.

Applying the round function means to evaluate 2^{32} pairs (f_i, g_i) with $f_i = f'_i \ggg 1$, $f'_i = a_i \oplus d_i$, and $g_i = (\gamma_3 \lll 1) \oplus e_i$. Define the multisets $F = \{f_i\}$, $F' = \{f'_i\}$, and $G = \{g_i\}$. We observe: The bits at the LSB-positions of both F' and G are balanced, and, due to the rotate, the bits at the MSB-position of F are balanced. Hence, the third round generates 2^{32} texts of the form (B, C, F, G) , which are then swapped to (F, G, B, C) .

The multisets (F, G) of inputs for the fourth round function F_4 are balanced. We write $(?, ?)$ for the outputs. Applying F_4 gives us 2^{32} texts $(F, G, ?, ?)$. After the swap, we get $(?, ?, F, G)$, where one bit in each F and G is balanced.

Figure 4 describes graphically, how the distinguisher works. Having chosen 2^{32} plaintexts, we can check the balancedness of the ciphertext bits at the two positions determined by the MSB of F and the LSB of G . Whatever the keys are, four rounds of Twofish always pass this test – even the post-whitening cannot destroy the balancedness. But a random permutation only passes this test with about a 25% probability.




Fig. 4. The Four-Round Distinguisher from Section 3.1

3.2 Another Four-Round Distinguisher

Our second distinguisher works quite similarly to the first one. We start with 2^{32} plaintexts of the form $(\alpha_0, \alpha_1, \alpha_2, A)$ with arbitrary constants α_i . The pre-whitening changes the constants and we get texts $(\beta_0, \beta_1, \beta_2, A)$. After the first round, including the swap, these are $(\gamma_2, A, \beta_0, \beta_1)$.

In the second round, the inputs to the round function are of the form (γ_2, A) , where $A = \{0, 1\}^{32}$ is a saturated set and γ_2 is constant. The round function generates the pairs (b_i, c_i) with $b_i = G_1(\gamma_2) + G_2(a_i) + K_6$ and $c_i = G_1(\gamma_2) + 2G_2(a_i) + K_7$ for $a_i \in A$. Now the set $B = \{b_i\}$ is saturated like A , but the multiset $C^* = \{c_i\}$ isn't. Instead, it is semi-saturated with a constant $\text{LSB}(c_i) = \gamma^* \in \{0, 1\}$ for all $c_i \in C^*$: $\gamma^* = \text{LSB}(G_1(\gamma_2)) \oplus \text{LSB}(K_7)$. We apply the round function by adding some constants to the elements of B and C^* , and by then rotating the elements of B . The results are a saturated set B and a semi-saturated set C^* , as before. After the swap, we have texts of the form (B, C^*, γ_2, A) .

In the third round the 2^{32} inputs for the round function are of the form (B, C^*) . Consider the round function's outputs (d_i, e_i) with $d_i = G_1(b_i) +$

$G_2(c_i) + K_8$ and $e_i = G_1(b_i) + 2G_2(c_i) + K_9$. Since $B = \{b_i\}$ is saturated, so is $\{G_1(b_i)\}$, and especially

$$\sum_{0 \leq i < 2^{32}} G_1(b_i) \equiv 0 \pmod{2}.$$

Since C^* is semi-saturated, it has 2^{31} different values, each repeated exactly twice. The same holds for the 2^{32} values $G_2(c_i)$ (with $c_i \in C^*$), hence

$$\sum_{0 \leq i < 2^{32}} G_2(c_i) \equiv 0 \pmod{2}.$$

Thus, both multisets $D = \{d_i\}$ and $E = \{e_i\}$ are balanced:

$$\sum_{0 \leq i < 2^{32}} d_i = \sum_i G_1(b_i) + \sum_i G_2(c_i) + 2^{32} * K_8 \equiv 0 \pmod{2}$$

and

$$\sum_{0 \leq i < 2^{32}} e_i = \sum_i G_1(b_i) + 2 * \sum_i G_2(c_i) + 2^{32} * K_9 \equiv 0 \pmod{2}.$$

By applying the round function and swapping, we get 2^{32} texts of the form (F, G, B, C^*) . The bits at the LSB-position of G are balanced, as are the bits at the MSB-position of F (due to the one-bit rotate). The fourth round makes this $(F, G, ?, ?)$, and if we do the swap we get texts of the form $(?, ?, F, G)$.

A random permutation passes the corresponding test only with a probability of about 0.25.

3.3 An Extension to Five Rounds

Next, we show how to extend the distinguisher from Section 3.2 to five rounds. Let α an arbitrary 32-bit constant and c^* an arbitrary 1-bit constant. We choose all 2^{95} plaintexts of the form (α, a_i, b_j, c_k) , with $c_i \text{ div } 2^{31} = c^*$. We write (α, A, B, C^+) for these 2^{95} texts. Note that the multisets A and B are saturated and the multiset C^+ is semi-saturated. The pre-whitening changes the constant α to β , and the constant c^* to γ^* , but leaves A and B saturated and C^+ semi-saturated with a constant MSB. We still have 2^{95} distinct input texts (β, A, B, C^+) for the first round.

Let $(e_i, f_i) = F_1(\beta, a_i)$ with $a_i \in A$. We can write $e_i = \beta_e + G_2(a_i)$ and $f_i = \beta_f + 2G_2(a_i)$, for some constants β_e, β_f . Hence the outputs of F_1 consist of pairs (E, F^*) with saturated E and semi-saturated F^* . Set $\beta^* = f_i \text{ mod } 2$ for the constant LSB of the values $f_i \in F^*$.

For every value $a_i \in A$ there are 2^{63} pairs (b_i, c_i) with a constant bit $\gamma^* = c_i \text{ div } 2^{31} = \text{MSB}(c_i)$. We can fix any constants $\gamma_2, \gamma_3 \in \{0, 1\}^{32}$ with $\gamma_3 \text{ mod } 2 = \gamma^* \oplus \beta^*$ and find pairs (b_i, c_i) in (B, C^+) such that $(e_i \oplus b_i) \ggg 1 = \gamma_2$ and $f_i \oplus (c_i \lll 1) = \gamma_3$ holds for every a_i . (Note that the MSB of c_i is the LSB of $c_i \lll 1$.)

Now the 2^{95} input texts (β, A, B, C^+) can be separated into 2^{63} disjoint groups of 2^{32} texts, determined by the pair (γ_2, γ_3) of constants, such that after applying the first round functions all texts in the same group are of the form $(\beta, A, \gamma_2, \gamma_3)$. The swap changes these to $(\gamma_2, \gamma_3, \beta, A)$.

For each such group, applying the four-round distinguisher from Section 3.2 would result in a set of 2^{32} ciphertexts $(?, ?, F, G)$, where the ciphertext bits at the LSB-position of G and at the MSB-position of F are balanced. Now, we do not know which ciphertexts belong into which group, but if these bits for each group are balanced, then so are all 2^{95} such bits. Five rounds of Twofish always pass this test, while a random permutation passes it with about 25% probability.

The same technique can also be applied to the distinguisher from Section 3.1. Here, we need 2^{96} plaintexts of the form (α, A, B, C) with constant α . A random permutation passes the corresponding test with about 25% probability.

3.4 An Extension to Six Rounds

To attack six rounds, we choose 2^{127} plaintexts (a_i, b_i, c_i, d_i) , (half of the codebook (!)), where $b_i \text{ div } 2^{31} = \text{MSB}(b_i)$ is fixed to an arbitrary constant. Our plaintexts are of the form (A, B^+, C, D) , where A, B , and D are saturated multiset, and B^+ is a semi-saturated one.

Our choice of plaintexts ensures that for each of the 2^{63} left-side pairs (a_i, b_i) , all 2^{64} right-side pairs (c_i, d_i) exist. Neither the pre-whitening nor the application of the first round function change this property. By the swap we get 2^{127} texts (C, D, A, B^+) as the input for the second round. For each 32-bit constant α we get a group of 2^{95} texts (α, D, A, B^+) . These are 2^{32} disjoint groups which are the kind of input we need for the 5-round distinguisher.

After six rounds of Twofish, we get 2^{127} ciphertexts $(?, ?, F, G)$ with balanced bits at two positions. A random permutation does satisfy this with about 25% probability.

3.5 Distinguishers: Summary

In Table 1 we summarise the distinguishers we have found. We describe which section the distinguisher was described in, the number r of Twofish rounds the attack works for, the chosen plaintexts required (how they look like and how many we need), and the probability for a random permutation to pass the test. All tests are one-sided, i.e. r rounds of Twofish pass the test with probability 1.

4 Finding the Key

In modern cryptanalysis, one often uses a distinguisher for some rounds of a product cipher to find the key: Guess some key bits for one or more additional rounds and exploit the distinguishing property to falsify wrong key guesses. This is what we do below, concentrating on using the six-round distinguisher.