

On recent higher-order cryptanalysis techniques

Jean-Philippe Aumasson



University of Applied Sciences Northwestern Switzerland
School of Engineering

FHNW, Switzerland

Agenda

Definitions: higher-order cryptanalysis, cube attacks, cube testers

Applications: Grain-128, Grain-v1, KATAN

Most recent developments: zero-sums and k -sums, application to Hamsi, Keccak, Luffa

Conclusion: how to resist higher-order cryptanalysis?

Definitions



=Google('higher-order differential') (???)

Higher-order cryptanalysis (1/2)

Differential cryptanalysis based on **order-1** derivatives:

$$\underline{\text{Ex:}} \quad E_k(m) \oplus E_k(m \oplus \Delta)$$

Higher-order differential cryptanalysis: based on derivatives of **order** ≥ 2

Order- d derivative with respect to d bits is the sum of all 2^d outputs obtained by varying these d bits

$$\frac{\partial^d f}{\partial x_1 \dots \partial x_d} = \sum_{(x_1, \dots, x_d) \in \{0,1\}^d} f(x)$$

for some $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $d \leq n$

Higher-order cryptanalysis (2/2)

Why can it work when classical differential cryptanalysis fails?

Ex: random $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of degree $d \leq n$

- ▶ Its first order derivative looks random
- ▶ Its order- $(d - 1)$ derivative is linear
- ▶ Its order- d derivative is a constant

Previous higher-order attacks called “integral cryptanalysis”, “Square attack”, “saturation attack”, etc.

Most recent and refined version: **cube attacks/testers**

Cube attacks/testers (1/8)

11:10

- 12:10

How to Solve it: New Techniques in Algebraic Cryptanalysis
Adi Shamir

Cube Attacks on Tweakeable Black Box Polynomials

Nad Dinur and Adi Shamir

Computer Science Department
The Weizmann Institute
Rehovot 76100, Israel

Abstract. Almost any cryptographic scheme can be described by tractable polynomials over GF(2), which contain both secret variables (e.g., key bits) and public variables (e.g., plaintext bits or IV bits). The cryptanalyst is allowed to tweak the polynomials by choosing arbitrary values for the public variables, and his goal is to solve the resulting system of polynomial equations in terms of the secret secret variables. In this paper we develop a new technique (called a cube attack) for solving such tractable polynomials, which is a major improvement over several previously published attacks of the same type. For example, on the stream cipher Trivium with a reduced number of initialization rounds, the best previous attack (due to Fauser, Khoo, and Zhou) requires a hardly practical complexity of 2^{27} to attack 472 initialization rounds, whereas a cube attack can find the complete key of the same version in 2^{27} in operations (which side has less a second on a single PC). Trivium with 788 initialization rounds (which could not be attacked by any previous technique) can now be broken with 2^{27} in operations. Trivium with 787 initialization rounds can now be broken with 2^{27} in operations, and the complexity of the attack can almost certainly be further reduced to about 2^{26} in operations. Whereas previous attacks were heuristic, had to be adapted to each cryptosystem, had an general complexity bounds, and were not required to succeed on random looking polynomials, cube attacks are generally successful when applied to random polynomials of degree d over a secret variable whenever the number m of public variables exceeds d^2 inputs. Their complexity is 2^{2d} for $d \leq 4$ operations, which is polynomial in m and asymptotically low when d is small. Cube attacks are applied to many black boxes, stream ciphers, or MAC which are provided as a black box (even when knowing is known about its internal structure) as long as its input-output can be represented by (univariate) polynomial of relatively low degree in the secret and public variables.

Keywords: Cryptanalysis, algebraic attacks, cube attacks, tweakeable black box polynomials, stream ciphers, Trivium.

1 Introduction

Solving large systems of multivariate polynomial equations is considered an overwhelmingly difficult problem, which has been studied extensively over many years.

© 2009 ACM 978-981-06-1222-2/09/0000-0000 \$7.50

Cube Testers and Key Recovery Attacks on Reduced-Round MD6 and Trivium

Jean Philippe Aumasson^{1,*}, Nad Dinur², Willi Meier^{1,*}, and Adi Shamir²

¹ FHNW, Wädswil, Switzerland
² Computer Science Department, The Weizmann Institute, Rehovot, Israel

Abstract. CRYSTO MD6 is a type of algebraic attack applicable to cryptographic functions having a low-degree algebraic normal form over GF(2). This paper applies cube attacks to reduced round MD6, finding the full 128-bit key of a 16-round MD6 with complexity 2^{27} (which takes less than a minute on a single PC). This is the best key recovery attack announced so far for MD6. We then introduce a new class of attack called cube testers, based on efficient pre-processing algorithms, and apply them to MD6 and to the stream cipher Trivium. Unlike the standard cube attacks, cube testers almost never behave better than the pre-processing step mentioned, but they use the attack cryptographic schemes described by non-linear polynomials of relatively high degree. Applied to MD6, cube testers detect non-randomness over 18 rounds in 2^{27} complexity, applied to a slightly modified version of the MD6 compression function, they can distinguish 48 rounds from random in 2^{27} complexity. Cube testers give distinctions on Trivium related to 796 rounds from random with 2^{27} complexity and detect non-randomness over 860 rounds in 2^{27} , improving on the original 788 round cube attack.

1 Introduction

1.1. Cube Attacks

Cube attacks [25] are a new type of algebraic cryptanalysis that exploit implicit low-degree equations in cryptographic algorithms. Cube attacks only require black box access to the target primitive, and were successfully applied to reduced versions of the stream cipher Trivium [8] in [2]. Roughly speaking, a cryptographic function is vulnerable to cube attacks if its algebraic normal form over GF(2) has degree at most d , provided that 2^d computations of the function is feasible. Cube attacks recover a secret key through queries to a black box polynomial with tractable public variables (e.g. stream plaintext or IV bits), followed by solving a linear system of equations in the secret key variables. A one-time preprocessing phase is required to determine which queries should be made to the black box during the on-line phase of the attack. Low-degree implicit equations were previously exploited in [11][22][23] to reconstruct

* Supported by the Swiss National Science Foundation, project no. 131020.

* Supported by CRYPTREC (CRYPTANALYSIS), project no. GR4-040-07.

© 2009 ACM 978-981-06-1222-2/09/0000-0000 \$7.50

© International Association for Cryptologic Research, 2009

[Dinur, Shamir; EUROCRYPT'09] [A., Dinur, Meier, Shamir; FSE'09]

Previous discovery claimed by Vielhaber (“AIDA”)...
ePrint 2007/413, 2009/402

Cube attacks/testers (2/8)

Cube attacks = key-recovery attacks (need a secret)

Offline phase (precomputation)

- ▶ Search for public variables (IV, plaintext) whose derivative is a linear combination of key bits
- ▶ Linearity detected via probabilistic testing
- ▶ Bit-per-bit reconstructions of equations

Online phase

- ▶ Evaluate each linear equation detected during precomputation
- ▶ Solve the linear system obtained

Cube attacks/testers (3/8)

Ex (coefficient of the max-degree monomial):

$$\begin{aligned}f(x_1, x_2, x_3, x_4) &= x_1 + x_3 + x_1x_2x_3 + x_1x_2x_4 \\ &= x_1 + x_3 + x_1x_2x_3 + x_1x_2x_4 + 0 \times x_1x_2x_3x_4\end{aligned}$$

Sum over all values of (x_1, x_2, x_3, x_4) :

$$f(0, 0, 0, 0) + f(0, 0, 0, 1) + f(0, 0, 1, 0) + \dots + f(1, 1, 1, 1) = 0$$

= order-4 derivative

Cube attacks/testers (4/8)

Ex (evaluation of linear combination):

$$\begin{aligned}f(x_1, x_2, x_3, x_4) &= x_1 + x_3 + x_1x_2x_3 + x_1x_2x_4 \\ &= x_1 + x_3 + x_1x_2(x_3 + x_4)\end{aligned}$$

Fix x_3 and x_4 , sum over all values of (x_1, x_2) :

$$\begin{aligned}\sum_{(x_1, x_2) \in \{0,1\}^2} f(x_1, x_2, x_3, x_4) &= 4 \times x_1 + 4 \times x_3 + 1 \times (x_3 + x_4) \\ &= x_3 + x_4\end{aligned}$$

= order-2 derivative

Cube attacks/testers (5/8)

x_3 and x_4 fixed and unknown

$f(\cdot, \cdot, x_3, x_4)$ queried as a **black box**

ANF unknown, except: $x_1 x_2$'s superpoly is $(x_3 + x_4)$

$$f(x_1, x_2, x_3, x_4) = \dots + x_1 x_2 (x_3 + x_4) + \dots$$

Query f to evaluate the superpoly:

$$\sum_{(x_1, x_2) \in \{0,1\}^2} f(x_1, x_2, x_3, x_4) = x_3 + x_4$$

Cube attacks/testers (6/8)

Key recovery attack on a stream cipher

$f : (k, v) \mapsto$ 1st keystream bit:

Offline: find cubes with linear superpolys

$$f(k, v) = \dots + v_1 v_3 v_5 v_7 (k_2 + k_3 + k_5) + \dots$$

$$f(k, v) = \dots + v_1 v_2 v_6 v_8 v_{12} (k_1 + k_2) + \dots$$

$$\dots = \dots$$

$$f(k, v) = \dots + v_3 v_4 v_5 v_6 (k_3 + k_4 + k_5) + \dots$$

Online: evaluate the superpolys, solve the system

Cube attacks/testers (7/8)

Cube testers = distinguishers

Detect a structure in the derivative which is not expected for an ideal algorithm

Ex: linearity, low degree, sparsity, imbalance

Compared to cube attacks

- ▶ At least as powerful (wrt # rounds attacked)
- ▶ Need less precomputation
- ▶ Do not require linear or low-degree derivative

Cube attacks/testers (8/8)

Problem: finding good sets of public variables (bottleneck)

Analytical approach:


- ▶ Analyze internals of the algorithm to determine variables with “lesser” interaction in the computation
- ▶ Ex: study of recurrence relations in Luffa by Hatano and Watanabe

Empirical approach:

- ▶ Use tools such a discrete optimization algorithms
- ▶ Ex: genetic algorithms for attacking Grain-128...

In practice, combine the two approaches

Applications



[D. J. Bernstein](#)
[Hash functions and ciphers](#)

Why haven't cube attacks broken anything?

The talk and the paper

Hundreds of cryptographers were sitting in a dark lecture room at the University of California at Santa Bar "How to solve it: new techniques in algebraic cryptanalysis."

Shamir had already advertised his talk as introducing "cube attacks," a powerful new attack technique that describing a stream cipher with an extremely large key, many S-boxes, etc. David Wagner later wrote that laugh -- since it seemed ridiculous to imagine an attack on the design, yet I knew if he was describing this

Grain-128 (1/2)

State-of-the-art stream cipher developed within


ECRYPT

↓↑↔↻⊕⊕^

's **eSTREAM** Project
(04-08)

- ▶ Designed by Hell, Johansson, Maximov, Meier (2007)
- ▶ 128-bit version of the eSTREAM cowinner Grain-v1
- ▶ 128-bit key, 96-bit IV, 256-bit state
- ▶ Previous DPA and related-key attacks
- ▶ Standard-model attack on 192-round version (of 256)

Grain-128 (2/2)



$\deg f = 1$, $\deg g = 2$, $\deg h = 3$

Initialization: key in NFSR, IV in LFSR, clock 256 times

Then 1 keystream bit per clock

Cube testers on Grain-128 (1/4)

[A., Dinur, Henzen, Meier, Shamir; SHARCS'09]

Method:

1. Select n variables (IV bits)
2. Set the remaining IV bits to zero
3. Set the key bits randomly
4. Run Grain-128 for all 2^n values to evaluate derivative
5. Repeat steps 3-4 N times and make statistics

Try to detect **imbalance**

Ex: if derivatives look like $x_0x_1x_2 + x_1x_2x_3x_4x_5$

Cube testers on Grain-128 (2/4)

Hardware implementation:

- ▶ Xilinx Virtex-5 FPGA
- ▶ 256 instances of $32 \times$ Grain-128 in parallel
- ▶ Efficient VHDL implementation of cube testers
- ▶ Attacks involving more than 2^{54} clocks in ≈ 1 day



Cube testers on Grain-128 (3/4)

Bitsliced C program:

- ▶ Run 64 instances of Grain-128 in parallel
- ▶ Used for parameters optimization (evolutionary algos)

```
u64 grain80_bitsliced64( u64 * key, u64 * tv, int rounds ) {
    u64 l[80+rounds], n[80+rounds], z=0;
    int l, i;

    /* initialize registers */
    for(l=0; l<64; l++){
        n[l]= key[l];
        l[l]= tv[l];
    }
    for(l=64; l<80; l++){
        n[l]= key[l];
        l[l]= 0xFFFFFFFFFFFFFFFFULL;
    }

    for(l=0; l<rounds; l++){
        /* clock */
        l[1+80] = l[1] ^ l[1+13] ^ l[1+23] ^ l[1+38] ^ l[1+51] ^ l[1+62];

        n[1+80] = l[1] ^ n[1] ^ n[1+9] ^ n[1+14] ^ n[1+21] ^ n[1+28] ^ n[1+33] ^ n[1+37] ^ n[1+45] ^ n[1+52] ^ n[1+60] ^
            n[1+62] ^ l[1+63] ^ n[1+68] ^ l[1+68] ^ l[1+68] ^ n[1+37] ^ n[1+33] ^ l[1+15] ^ n[1+9] ^
            ( n[1+60] & n[1+52] & n[1+45] ) ^ ( n[1+33] & n[1+28] & n[1+21] ) ^ ( n[1+63] & n[1+45] & n[1+28] & n[1+9] ) ^
            ( n[1+60] & n[1+52] & n[1+37] & n[1+33] ) ^ ( n[1+63] & n[1+60] & n[1+21] & n[1+15] ) ^
            ( n[1+63] & n[1+60] & n[1+52] & n[1+45] & n[1+37] ) ^ ( n[1+33] & n[1+28] & n[1+21] & n[1+15] & n[1+9] ) ^
            ( n[1+52] & n[1+45] & n[1+37] & n[1+33] & n[1+28] & n[1+21] );

        z = l[1+25] ^ n[1+63] ^ l[1+3] ^ l[1+64] ^ l[1+64] ^ l[1+64] ^ l[1+64] ^ n[1+63] ^ l[1+3] ^ l[1+25] ^ l[1+46] ^
            ( l[1+3] & l[1+46] & l[1+64] ) ^ ( l[1+3] & l[1+46] & n[1+63] ) ^ ( l[1+25] & l[1+46] & n[1+63] ) ^ ( l[1+46] & l[1+64] & n[1+63] );
        z = n[1 + 2] ^ n[1 + 2] ^ n[1 + 4] ^ n[1 + 10] ^ n[1 + 31] ^ n[1 + 45] ^ n[1 + 56] ^ z;

        l[1+80] ^= z; n[1+80] ^= z;
    }


    /* return 1 keystream bit */
    z = (n[1+12] & l[1+81]) ^ (l[1+13] & l[1+20]) ^ (n[1+95] & l[1+42]) ^ (l[1+60] & l[1+79]) ^ (n[1+12] & n[1+95] & l[1+95]);
    z = n[1 + 2] ^ n[1 + 15] ^ n[1 + 36] ^ n[1 + 45] ^ n[1 + 64] ^ n[1 + 73] ^ n[1 + 89] ^ z ^ l[1 + 93];

    return z;
}
```

Cube testers on Grain-128 (4/4)

Distinguisher for 237 rounds (of 256) in 2^{40}

Extrapolation:



Suggests existence of **distinguishers in 2^{77}**

⇒ 128-bit security unlikely

The case of Grain-v1

eSTREAM cowinner, original version of Grain

2×80 -bit state, 160 initialization rounds

Seems to **resist cube testers** (81 rounds in 2^{24}), why?

The case of Grain-v1

eSTREAM cowinner, original version of Grain

2×80 -bit state, 160 initialization rounds

Seems to **resist cube testers** (81 rounds in 2^{24}), why?

- ▶ NFSR feedback of degree 6 (vs. 2 for Grain-128)
- ▶ Filter function of degree 3 (vs. 2)
- ▶ Denser feedback and filter functions
- ▶ Shorter feedback delay (16 vs. 32)
- ▶ Smaller registers (80 vs. 128)

⇒ converges faster towards ideal ANF

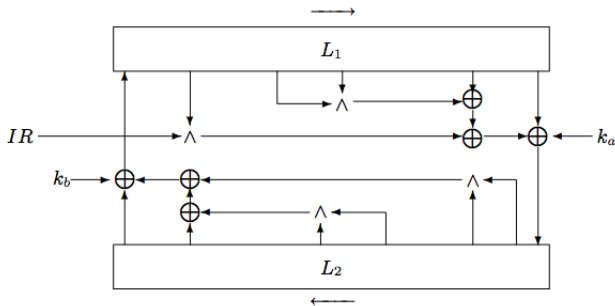
The case of KATAN (1/2)

[De Cannière, Dunkelman, Knezevic; CHES'09]

Lightweight block cipher

80-bit key; 32-, 48-, or 64-bit blocks

Compact in HW: 802 GE, for 6.25 GE/flip-flop (KATAN-32)



The case of KATAN (2/2)

NFSR's properties:

- ▶ Degree-2 and sparse feedback function
- ▶ Short feedback delay (3)

Paper: *“after 160 rounds, the degree of each internal state bit can reach 32”* (for KATAN-32)

Our observations: degree 20 reached after 55 rounds, thus **degree-32 probably reached after 87 rounds only!**

The case of KATAN (2/2)

NFSR's properties:

- ▶ Degree-2 and sparse feedback function
- ▶ Short feedback delay (3)

Paper: *“after 160 rounds, the degree of each internal state bit can reach 32”* (for KATAN-32)

Our observations: degree 20 reached after 55 rounds, thus **degree-32 probably reached after 87 rounds only!**

⇒ sparse and degree-2 function okay when feedback delay is short...

...but combinatorial logic is cheap (a few NAND's), while memory (FSR's) is expensive in hardware...

⇒ better increase degree and density as a safety net?

Most recent developments

[A., Knudsen, Meier]

k -sums

Consider a permutation $F : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$

k -sum: set $\{x_1, \dots, x_k\}$ such that

$$\bigoplus_{i=1}^k F(x_i) = 0$$

Generalized birthday attack in time and space

$$O(k \cdot 2^{\ell/(1+\log k)})$$

Zero-sums (1/2)

Consider a permutation $P : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$

Zero-sum: set $\{x_1, \dots, x_k\}$ such that

$$\bigoplus_{i=1}^k x_i = \bigoplus_{i=1}^k P(x_i) = 0$$


Generic probabilistic methods:

- ▶ Generalized birthday attack
- ▶ XHASH attack (linear algebra)

How to exploit the algebraic structure to find zero-sums?

Zero-sums (2/2)

Cube attack and k -sums need



Zero-sums need



Inside-out strategy

- ▶ Fix state in the middle, vary k bits
- ▶ If degree $< k$ for both halves, 2^k values sum to zero

Need only evaluate at most **half** the algorithm

Application to Keccak (1/3)

- ▶ Second-round SHA-3 candidate
- ▶ 1600-bit state
- ▶ 18 nonidentical rounds

Application to Keccak (1/3)

- ▶ Second-round SHA-3 candidate
- ▶ 1600-bit state
- ▶ 18 nonidentical rounds
- ▶ One round has degree 2
- ▶ One inverse round has degree 3


10 rounds: degree upper bound $2^{10} = 1024$ (suboptimal)
 \Rightarrow higher-order distinguisher in 2^{1024}

13 rounds: degree upper bound $2^{13} \gg 1599$ (optimal)
 \nRightarrow higher-order distinguisher

Application to Keccak (2/3)

Consider 257 variables in the state after 5 rounds

- ▶ Preimage = degree-243 mapping
- ▶ Image = degree-256 mapping
- ▶ Compute order-257 derivative in both directions



Obtain $(x_1, \dots, x_{2^{257}})$ such that $\bigoplus_{i=1}^{2^{257}} x_i$ is the order-257 derivative of a degree-256 polynomial: must be **zero**

\Rightarrow zero-sum on 13 rounds in $2^{257} \times$ first 5 rounds

Application to Keccak (3/3)

Optimizations: exploit structure of the inverse permutation

#rounds	complexity
8	2^{17}
10	2^{60}
12	2^{128}
14	2^{256}
16	2^{1024}

(18 rounds in full version)

Tweak for the second round: #rounds set to 24, rate modified

Application to Luffa

- ▶ Second-round SHA-3 candidate
- ▶ AND/XOR algorithm (like Keccak)
- ▶ Tweaked for the second round

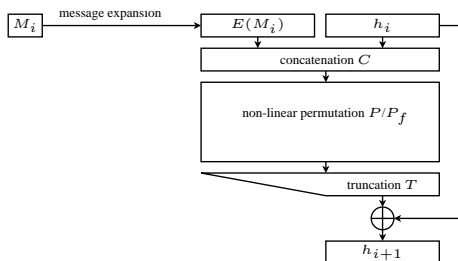
Q permutation of **Luffa** (256-bit)

- ▶ Distinguisher on full version (8 rounds) in 2^{81}
- ▶ Distinguisher on 7 rounds in 2^{27}
- ▶ Not relevant for the hash algorithm

Application to Hamsi (1/3)

- ▶ Second-round SHA-3 candidate
- ▶ Two main instances: **Hamsi-256** and Hamsi-512
- ▶ Serpent-like algorithm (4-bit Sbox + linear layer)

Davies-Meyer compression function



3 rounds (6 for the last compression)

Application to Hamsi (2/3)

Observations:

- ▶ 3 rounds have degree 3 only, instead of ideally 27 (with respect to carefully chosen variables)
- ▶ Distribution of monomials and binomials is sparse

Application to Hamsi (2/3)

Observations:

- ▶ 3 rounds have degree 3 only, instead of ideally 27 (with respect to carefully chosen variables)
- ▶ Distribution of monomials and binomials is sparse

Consequences:

16-, 8-, 4-sums can be found efficiently

Example found for the default IV of Hamsi. . .

Zero-sums can be found efficiently for the permutation

Application to Hamsi (3/3)

Previous near collisions (or 2 sums):

- ▶ (256 – 25)-bit collision from 14 bit differences [Nikolic]
- ▶ (256 – 23)-bit collision from 16 bit differences [Wang et al.]

Application to Hamsi (3/3)

Previous near collisions (or 2 sums):

- ▶ (256 – 25)-bit collision from 14 bit differences [Nikolic]
- ▶ (256 – 23)-bit collision from 16 bit differences [Wang et al.]

We found a differential characteristic of probability 2^{-26}

Consequence:

(256 – 25)-bit collision from 6 bit differences

Easier for the default IV than for a random one. . .

Conclusion



Conclusion

Higher-order methods are diverse, simple, powerful. . .

But only on certain designs, based on

- ▶ AND/XOR
- ▶ Small Sboxes

AXR and AES-based designs immune (even for low #rounds)

Recommendations for new designs

- ▶ If possible, use ADD (or other highly nonlinear op.)
- ▶ Maximum degree achieved with 25% of the #rounds
- ▶ Benchmark with cube testers

On recent higher-order cryptanalysis techniques

Jean-Philippe Aumasson



University of Applied Sciences Northwestern Switzerland
School of Engineering

FHNW, Switzerland