

Inside the Hypercube

Jean-Philippe Aumasson^{1*}, Willi Meier^{1†}, María Naya-Plasencia^{2‡}, and Thomas Peyrin³

¹ FHNW, Windisch, Switzerland

² INRIA project-team SECRET, France

³ Ingenico, France

Some force inside the Hypercube occasionally manifests itself with deadly results.

http://www.staticzombie.com/2003/06/cube_2_hypercube.html

Abstract. Bernstein’s CubeHash is a hash function family that includes four functions submitted to the NIST Hash Competition. A CubeHash function is parametrized by a number of rounds r , a block byte size b , and a digest bit length h (the compression function makes r rounds, while the finalization function makes $10r$ rounds). The 1024-bit internal state of CubeHash is represented as a five-dimensional hypercube. Submissions to NIST have $r = 8$, $b = 1$, and $h \in \{224, 256, 384, 512\}$.

This paper gives the first external analysis of CubeHash, with

- improved standard generic attacks for collisions and preimages
- a multicollision attack that exploits fixed points
- a study of the round function symmetries
- a preimage attack that exploits these symmetries
- a practical collision attack on a weakened version of CubeHash
- high-probability truncated differentials over the 10-round transform

Our results do not seem to contradict the security claims about CubeHash.

1 CubeHash

Bernstein’s CubeHash is a hash function family that includes four functions submitted to the NIST Hash Competition. A CubeHash function is parametrized by a number of rounds r , a block byte size b , and a digest bit length h ; the 1024-bit internal state of CubeHash is viewed as a five dimensional hypercube. Submissions to NIST have $r = 8$, $b = 1$, and $h \in \{224, 256, 384, 512\}$.

CubeHash computes the digest of a message as follows:

- initialize a 1024-bit state as a function of (h, b, r)
- append to the message a 1 bit and enough 0 bits to reach a multiple of $8b$ bits
- for each b -byte message block:
 - xor the block into the first b bytes of the state
 - transform the state through the r -round T function
- xor a 1 bit with the 993th bit of the state
- transform the state through $10r$ -round T
- output the first h bits of the state

*Supported by the Swiss National Science Foundation under project no. 113329.

†Supported by GEBERT RÜF STIFTUNG, project no. GRS-069/07.

‡Supported in part by the French Agence Nationale de la Recherche under contract ANR-06-SETI-013-RAPIDE.

Let $x[0], \dots, x[31]$ represent the 1024-bit state as an array of 32-bit words. The transform function T makes r identical rounds, where each round computes (see also Fig. 1):

```

for  $i = 0, \dots, 15$ :    $x[i + 16] = x[i + 16] + x[i]$ 
for  $i = 0, \dots, 15$ :    $y[i \oplus 8] = x[i]$ 
for  $i = 0, \dots, 15$ :    $x[i] = y[i] \lll 7$ 
for  $i = 0, \dots, 15$ :    $x[i] = x[i] \oplus x[i + 16]$ 
for  $i = 0, \dots, 15$ :    $y[i \oplus 2] = x[i + 16]$ 
for  $i = 0, \dots, 15$ :    $x[i + 16] = y[i]$ 
for  $i = 0, \dots, 15$ :    $x[i + 16] = x[i + 16] + x[i]$ 
for  $i = 0, \dots, 15$ :    $y[i \oplus 4] = x[i]$ 
for  $i = 0, \dots, 15$ :    $x[i] = y[i] \lll 11$ 
for  $i = 0, \dots, 15$ :    $x[i] = x[i] \oplus x[i + 16]$ 
for  $i = 0, \dots, 15$ :    $y[i \oplus 1] = x[i + 16]$ 
for  $i = 0, \dots, 15$ :    $x[i + 16] = y[i]$ 

```

2 Improved standard generic attacks

The author of CubeHash presented [2] the following “standard preimage attack”:

- from (h, b, r) compute the initial state S_0
- from the h -bit image plus some arbitrary $/1024 - h$ bits, invert $10r$ rounds and the “xor 1” to get a state S_f before finalization
- find two n -block sequences that map S_0 (forward) and S_f (backward), respectively, to two states that share the last $(1024 - 8b)$ bits

There are 2^{nb} possible n -block inputs and one looks for a collision over $(1024 - 8b)$ bits. For a success chance $1 - 1/e \approx 0.63$ one thus requires 2^{512-4b} trials in each direction, that is, $2nb > 1024 - 8b$, i.e., $n > 512/b - 4$. In total the number of evaluations of T is approximately

$$2 \times \left(\frac{512}{b} - 4 \right) \times 2^{512-4b} \approx 2^{522-4b-\log b}.$$

Furthermore, [2] estimates that each round of T needs 2^{11} “bit operations”; the above formula gives about $2^{533-4b-\log b+\log r}$ bit operations.

A speed-up of the above attack can be obtained by searching a collision not only in the states resulting of a n -block computation, but in every distinct state reached (i.e. also with the intermediate states). This is made possible by the absence of message length padding. Each call to T gives a new candidate for the collision search; we thus get rid of the $(512/b-4)$ multiplicative factor in the cost estimate. This gives a cost of

$$2 \times 2^{512-4b} = 2^{513-4b}$$

evaluations of T , i.e. $2^{524-4b+\log r}$ bit operations.

The proposed CubeHash-512 has $(h, b, r) = (512, 1, 8)$, our attack thus makes 2^{523} bit operations, against 2^{532} with the original attack. If $r = 8$, our attack needs $b > 3$ to make less than 2^{512} bit operations, against $b > 5$ with the original preimage attack. It is to note that these estimates exclude the nonnegligible communication costs.

One can use the same trick to speed-up the standard collision attack [2]; the cost in T evaluations then drops from $2^{521-4b-\log b}$ to 2^{512-4b} .




Fig. 1. Schematic view of a CubeHash round.

3 Narrow-pipe multicollisions

Based on the “narrow-pipe” attacks in [3], we show a multicollision attack on CubeHash faster than Joux’s [6] or birthday [5,8] methods (for large b ’s). Our attack requires the same amount of computation as narrow-pipe collisions. It exploits the fact that the null state is a fixed point for the compression function T (regardless of r), and that the message padding doesn’t include the message length.

Starting from an initial state S_0 derived from (h, b, r) , one finds two n -block sequences m and m' that map S_0 (forward) and the zero state (backward), respectively, to two states that share the last $(1024 - 8b)$ bits. One finds a connection of the form

$$\begin{aligned} S_0 \oplus m_1 &\xrightarrow{T} S_1 \\ S_1 \oplus m_2 &\xrightarrow{T} \dots \\ &\dots \\ \dots &\xrightarrow{T} S'_1 \\ S'_1 \oplus m'_2 &\xrightarrow{T} 0 \oplus m'_1 \end{aligned}$$

Once a path to the zero state is found, one can add an arbitrary number of zero message blocks to maintain a zero state. Colliding messages are of the form

$$m \| m' \| 0 \| 0 \| \dots \| 0 \| \bar{m},$$

where \bar{m} is an arbitrary sequence of blocks.

Using the technique of §2, this multicollision attack requires approximately 2^{513-4b} evaluations of T . In comparison, a birthday attack finds a k -collision in $(k! \times 2^{n(k-1)})^{1/k}$ trials, and Joux’s attacks in $\log k \times 2^{4(128-b)}$. For example, with $h = 512$ and $b = 112$, our attack finds 2^{64} -collisions within 2^{65} calls to T , against $> 2^{512}$ for a birthday attack and 2^{70} for Joux’s.

4 On state symmetries

The documentation of CubeHash mentions [4, p.3] the existence of symmetries through the round function, and states that the initialization of CubeHash was designed to avoid symmetries. However [4] gives no detail on those symmetries. In this section we present five symmetry classes of 2^{512} states each, and show how to exploit them.

4.1 Symmetry classes

If a 32-word state x satisfies $x[0] = x[1]$, $x[2] = x[3]$, \dots , $x[30] = x[31]$, then this property is preserved through the transformation T , for any number of rounds. One can represent this symmetry with the pattern (each letter stands for a 32-bit word):

AABBCCDD EEFFGGHH IIJJKKLL MMNNOOPP .

In total we found five classes of symmetry:

C_1 : AABBCCDD EEFFGGHH IIJJKKLL MMNNOOPP
 C_2 : ABABCD CD EFEFGH GH IJIKL KL MNNOPOP
 C_3 : ABCD ABCD EFGH EFGH IJKL IJKL MNOP MNOP
 C_4 : ABCDEFGH ABCDEFGH IJKLMN OP IJKLMN OP
 C_5 : ABBACDDC EFFEGHHG IJJKLLK MNNMOPPO

Each class contains 2^{512} states. If a state belongs to several classes, then its image under T also belongs to these classes; for example if $S \in (C_i \cap C_j)$, then $T(S) \in (C_i \cap C_j)$. We have

$$\begin{aligned} |C_1 \cap C_2| &= |C_1 \cap C_5| = |C_2 \cap C_3| = |C_2 \cap C_5| = |C_3 \cap C_4| = 256 \\ |C_1 \cap C_3| &= |C_2 \cap C_4| = |C_3 \cap C_5| = 128 \\ |C_1 \cap C_4| &= |C_4 \cap C_5| = 64 \end{aligned}$$

We thus have $|\cup_{i=1}^5 C_i| \geq 5 \times 2^{512} - 10 \times 2^{256} \approx 2^{514.3}$ distinct symmetric states. Note that symmetry is not preserved by the finalization procedure of CubeHash (the “xor 1” breaks any of the above symmetries).

4.2 Exploiting symmetric states

Preimages. Given a target digest, one can make a preimage attack similar to that in §2, and exploit symmetric states for the connection. The attack goes as follows:

- from the initial state, reach a symmetric state (of any class) by using $2^{1024-514-8} = 2^{502}$ message blocks
- from a state before finalization, reach (backwards) another symmetric state (not necessarily of the same class)
- from these two symmetric states in classes C_i and C_j , use null message blocks in both directions to reach two states in $C_i \cap C_j$
- find a collision by trying $\sqrt{|C_i \cap C_j|}$ messages in each direction

Complexity of steps 1 and 2 is about 2^{503} computations of T . The cost of steps 3 and 4 depends on i and j ; there are three distinct cases (counting in calls to T):

1. $i = j$ (with prob. 5/25): step 3 costs 0 and step 4 costs 2×2^{256}
2. $|C_i \cap C_j| = 2^{256}$ (with prob. 10/25): step 3 costs 2×2^{256} and step 4 costs 2×2^{128}
3. $|C_i \cap C_j| = 2^{128}$ (with prob. 6/25): step 3 costs 2×2^{384} and step 4 costs 2×2^{64}
4. $|C_i \cap C_j| = 2^{64}$ (with prob. 4/25): step 3 costs 2×2^{448} and step 4 costs 2×2^{32}

In any case, the total complexity is about 2^{503} calls to T . This attack, however, finds messages of unauthorized size (more than 2^{257} bytes!).

One can find preimages of reasonable size by using a variant of the above attack: suppose $b > 4$, from the initial state reach a state in a given class C_i , do the same backwards from a state before finalization. For a given b , the complexity of reaching a symmetric state depends on the C_i considered (see below). Then one seeks a collision within C_i by trying messages preserving the symmetry: for example, if $b = 5$ and $C_i = C_1$, then one has to preserve the equality $x[0] = x[1]$ and shall thus pick 5-byte messages of the form X000X (each digit stands for a byte). Since any C_i contains 2^{512} states, the cost of finding a collision within C_i is about 2^{256} trials in each direction.

Below we give the C_i that is the easiest to reach, depending on the value of b :

- $5 \leq b < 9$: the best class is C_1 , which gives $(1024 - 2 \times 4 \times 8)/2 = 480$ equations to verify
- $9 \leq b < 17$: the best class is C_2 or C_5 , which give $(1024 - 2 \times 8 \times 8)/2 = 448$ equations to verify
- $17 \leq b < 33$: the best class is C_3 , which gives $(1024 - 2 \times 16 \times 8)/2 = 384$ equations to verify

- $33 \leq b < 65$: the best class is C_4 , which gives $(1024 - 2 \times 32 \times 8)/2 = 256$ equations to verify

If n equations have to be verified, the cost of reaching a symmetric state is about 2^n evaluations of T . Compared to the preimage attack in §2, the best speed-up obtained from a given C_i is when $b = 4d + 1$, where d is the number of 32-bit words that separate the first repetition of two words.

To illustrate this attack, let's study in more detail the case of C_1 :

- if $b \equiv 0 \pmod 8$, there are $(1024 - 8b)/2 = 512 - 4b$ equations to satisfy, thus about 2^{512-4b} calls to T are necessary
- if $b \equiv 4 \pmod 8$, there are only $(1024 - 8b - 32)/2 = 496 - 4b$ equations to satisfy, because one has no condition on the first state word not xored with the message block
- generalizing, when $b \pmod 8 \leq 4$, about $2^{512-4(b+(b \pmod 4))}$ calls to T are necessary
- when $b \pmod 8 > 4$, there are $(1024 - 8b - 32 + 8(b \pmod 4))/2$ equations to satisfy, which gives a cost $2^{496-4(b-(b \pmod 4))}$

The general formula for the number of equations is

$$512 - 32\lfloor b/8 \rfloor - 32\lfloor (b \pmod 8)/4 \rfloor - [(\lfloor (b \pmod 8)/4 \rfloor + 1) \pmod 2] \times 8(b \pmod 4) .$$

In the best case ($b \equiv 4 \pmod 8$), the attack is 2^{15} times faster than that in §2 (in the worst case, $b \equiv 0 \pmod 8$, it has the same complexity). Note that when $b = 5$, the attack makes about 2^{481} calls to T , against 2^{493} with the attack in §2.

Collisions on a weakened CubeHash. The initialization of CubeHash never leads to a symmetric initial state. Here we present a practical collision attack that would apply if the initial state were symmetric, and in $C_1 \cap C_4$.

Suppose that the initial state of CubeHash $r/b-h$ is in $C_1 \cap C_4$, i.e. is of the form

AAAAAAAA AAAAAAAAAA BBBBBBBBB BBBBBBBB .

If one hashes the $b2^{33}$ -byte message that contain only zeros, then each of the 2^{33} intermediate states is an element of $C_1 \cap C_4$. Assuming that T acts like a random permutation of $C_1 \cap C_4$, one will find two identical states with probability about 0.63, which directly gives a collision.

5 Truncated differentials over T

This section shows how to detect nonrandomness over the 10-round T transform. We start from a weight-64 difference to reach a weight-1 difference after 3 rounds with high probability; this *nonlinear* differential was discovered by simply computing backwards from the weight-1 difference.

We consider the input difference 80000000 in $x[16]$. The word $x[16]$ was chosen because $x[16] \cdots x[31]$ diffuse less in the first rounds than $x[0] \cdots x[15]$. We set a difference 80000000 to minimize the impact of carries.

We consider the following nonlinear differential. Input difference (weight-64):

```
18000000 10000000 08000000 30000000
00000040 00000080 00000000 00000000
00400000 00000000 00400000 01000404
00000003 80802002 00000001 81802004
40000000 08000000 00000000 E8020600
00000000 00000100 00000080 41F001C0
00400008 00000008 00400000 01000404
00000005 80802002 00000001 8080200C
```

Difference after one round (weight-26):

```
000E0000 00000000 00000000 00000000
00000000 00000040 00000080 00000040
01000004 00000000 00000004 00000000
00000000 00000000 00002000 00000000
800E0200 00000000 00000000 00000000
00000000 000000C0 00000080 000001C0
00000000 00000004 00000000 00000004
00000000 00002000 0000C000 00000000
```

Difference after two rounds (weight-9):

```
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 01000000
00000000 00002000 00000000 00002000
00000000 00000000 00000000 80000000
00000000 00000000 00000000 00100000
00000000 00000000 00000000 03000000
00000000 00002000 00000000 00002000
```

Difference after three rounds (weight-1):

```
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
80000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
```

which after another round gives with probability 1 the difference

```
80000000 00000000 80000000 00000000
00000400 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 80000000 00000000 80000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
00000000 00000000 00000000 00000000
```

This differential holds with negligible probability for a random message. But it holds for the message

```
DFB7AA11 7B2872F1 2848B142 64CB0AF9
17DA36E7 320A7AB2 27621CD8 B6E23031
3BCE90DB 0E496C61 AF4156BD 0B4D857F
4379D4C0 D495EAC9 038BD6E5 72A114CC
29065395 824774C3 F0923C34 28F3B2DD
74251DF6 1A562265 BD8EE5E3 DEFDD839
2804D3BE 89417DC3 F001CE4A 6A5328A8
2BEC024E B2306F17 1F2A7C6C 14BC37B6
```

For 32 random bits in $x[25]$ and $x[26]$ (at positions 4, ..., 19 in both), the differential is satisfied with probability approximately 0.985.

Note that in the linear model (i.e. when additions are replaced by xors), a differential path starting from the weight-1 difference cycles over 47 rounds. That is, it comes back to the difference 80000000 in $x[16]$ after 47 rounds. In the original model, however, linear differentials are followed with negligible probability.

Based on the above differential, we empirically looked for high-probability truncated differentials, based on the weight-64 input difference, and applying to each output bit a frequency test similar to that in [7, §2.1], with decision threshold 0.001 and 2^{20} samples. We found 4 output bits with p-value less than 0.001, at positions 579, 778, 841, and 842. Over 11 rounds and more, no bias was detected.

This observation is consistent with the fact that, when starting from the weight-1 differential, we could detect nonrandomness on up to 7 rounds (now this difference is introduced three rounds later). Note that in a previous version of this article [1], we reached 8 rounds by starting one round before the weight-1 difference.

These observations indicate that 10-round T does not act as a random permutation, and that 10 rounds may not be overkill, as suggested in [3]. But note that the settings used don't correspond to a realistic attack scenario. Furthermore, if we restrict ourselves to differences in the first state byte, and put random bits in the rest of the state, then we observe nonrandomness after up to 5 rounds.

References

1. Jean-Philippe Aumasson, Willi Meier, Mara Naya-Plasencia, and Thomas Peyrin. Inside the hypercube. Cryptology ePrint Archive, Report 2008/486, 2008. version 20081124:132635.
2. Daniel J. Bernstein. CubeHash appendix: complexity of generic attacks. Submission to NIST, 2008.
3. Daniel J. Bernstein. CubeHash attack analysis (2.B.5). Submission to NIST, 2008.
4. Daniel J. Bernstein. CubeHash specification (2.B.1). Submission to NIST, 2008.
5. Persi Diaconis and Frederick Mosteller. Methods for studying coincidences. *Journal of the American Statistical Association*, 84(408):853–861, 1989.
6. Antoine Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In *CRYPTO*, 2004.
7. NIST. SP 800-22, a statistical test suite for random and pseudorandom number generators for cryptographic applications, 2001.
8. Kazuhiro Suzuki, Dongvu Tonien, Kaoru Kurosawa, and Koji Toyota. Birthday paradox for multi-collisions. In *ICISC*, 2006.