

Inside the Hypercube

Jean-Philippe Aumasson^{1,*}, Eric Brier³, 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. The submissions to NIST recommends $r = 8$, $b = 1$, and $h \in \{224, 256, 384, 512\}$. This paper presents 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
- a study of fixed points and an example of nontrivial fixed point
- high-probability truncated differentials over 10 rounds

Since the first publication of these results, several collision attacks for reduced versions of CubeHash were published by Dai, Peyrin, et al. Our results are more general, since they apply to any choice of the parameters, and show intrinsic properties of the CubeHash design, rather than attacks on specific versions.

1 CubeHash

Bernstein’s CubeHash is a hash function family 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. The submissions to NIST recommends $r = 8$, $b = 1$, and $h \in \{224, 256, 384, 512\}$.

CubeHash computes a message digest as follows:

*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.

- 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 993rd bit of the state
- transform the state through $10r$ -round T
- output the first h bits of the state

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  $i = 0, \dots, 15$ :  $x[i + 16] = y[i]$ 

```

See [5] for a more detailed description of CubeHash.

This paper presents 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
- a study of fixed points and an example of nontrivial fixed point
- high-probability truncated differentials over the 10-round transform

After the first publication of this article [2], Dai, Peyrin, et al. presented a series of collision attacks [1, 6–8] on reduced versions of CubeHash. Their best results (as of Feb. 6) are an example of collision on CubeHash3/64 and a collision attack on CubeHash4/3 in about 2^{207} simple operations [6]. Our results, however, are more general, since they apply to any choice of the parameters, and show intrinsic properties of the CubeHash design, rather than attacks on specific versions.

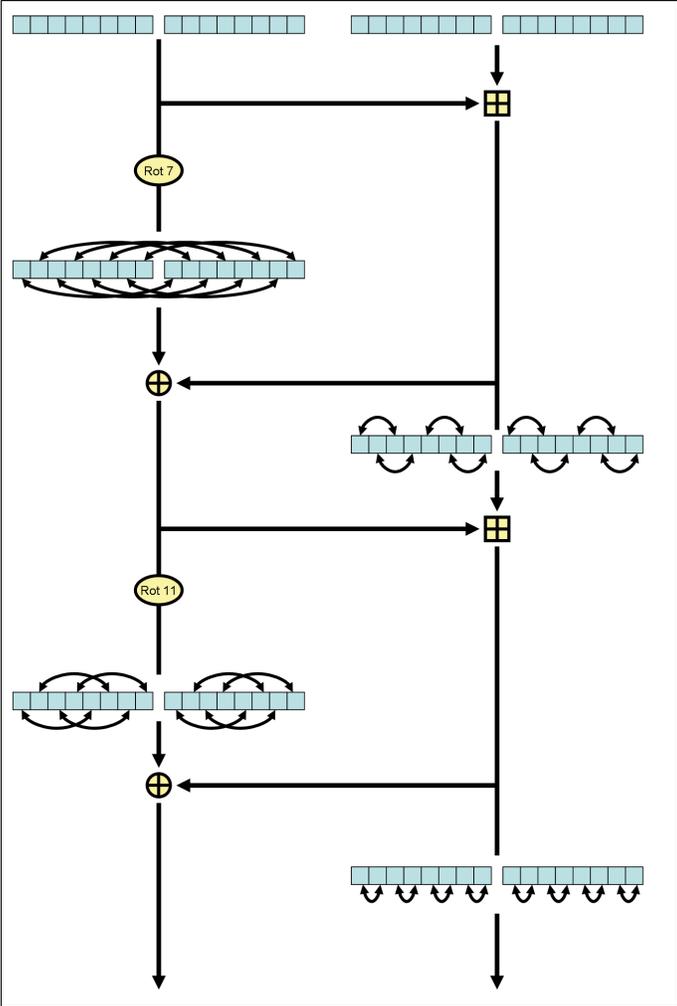


Fig. 1. Schematic view of a CubeHash round.

2 Improved standard generic attacks

The author of CubeHash presented [3] 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, [3] 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 non-negligible communication costs.

One can use the same trick to speed-up the standard collision attack [3]; the cost in T evaluations then drops from $2^{521-4b-\log b}$ to 2^{512-4b} .

3 Narrow-pipe multicollisions

Based on the “narrow-pipe” attacks in [4], we show a multicollision attack on CubeHash faster than Joux’s [10] or birthday [9,12] 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 does not 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 State symmetries

The documentation of CubeHash mentions [5, p.3] the existence of symmetries through the round function, and states that the initialization of CubeHash was designed to avoid them. However [5] gives no detail on those symmetries. In this section, we provide a reasoning that finds all symmetries inherent in the transformation T . In total we are able to show 15 symmetry classes of 2^{512} states each, and show how to exploit these.

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 , with probability equal to 1, for any number of rounds. One can represent this symmetry with the pattern (each letter stands for a 32-bit word):

$$\text{AABBCDD EEFGGHH I I JKKLL MMNNOOPP .}$$

In total we found 15 classes of symmetry:

C_1 : AABBCDD EEFGGHH IJJKKLL MMNNOOP
 C_2 : ABABCD CD EFEFGHG IJIKLKL MNNOPOP
 C_3 : ABBACDDC EFFEFGHG IJJIKLLK MNMOPPO
 C_4 : ABCDABCD EFGHEFGH IJKLIJKL MNOPMNOP
 C_5 : ABCDBADC EFGHFEHG IJKLJILK MNOPNMPO
 C_6 : ABCDCDAB EFGHGHEF IJKLKLIJ MNOPOPMM
 C_7 : ABCDDCBA EFGHHGFE IJKLLKJI MNOPPONM
 C_8 : ABCDEFGH ABCDEFGH IJKLMNOP IJKLMNOP
 C_9 : ABCDEFGH BADCFEFGH IJKLMNOP JILKNMPO
 C_{10} : ABCDEFGH CDABGHEF IJKLMNOP KLIJOPMN
 C_{11} : ABCDEFGH DCBAHGFE IJKLMNOP LKJIPONM
 C_{12} : ABCDEFGH EFGHABCD IJKLMNOP MNOPIJKL
 C_{13} : ABCDEFGH FEHGBADC IJKLMNOP NMPOJILK
 C_{14} : ABCDEFGH GHEFCDAB IJKLMNOP OPMNKLIJ
 C_{15} : ABCDEFGH HGFEDCBA IJKLMNOP PONMLKJI

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

$$|C_i \cap C_j| \leq 2^{256} .$$

By the inclusion-exclusion principle, the number of distinct symmetric states is

$$|\cup_{i=1}^{15} C_i| = 15 \times 2^{512} - 70 \times 2^{256} + 120 \times 2^{128} - 64 \times 2^{64} \approx 2^{516} .$$

Note that symmetry is not preserved by the finalization procedure of CubeHash (the “xor 1” breaks any of the above symmetries).

4.2 Finding all symmetry classes

Now we prove that the classes C_1, \dots, C_{15} capture all the possible symmetries of CubeHash’s transform T . A symmetry class can be represented as a set of pairs (i, j) , where each (i, j) means $x[i] = x[j]$. For example, C_1 can be described by the set

$$\begin{aligned}
& (0,1) \quad (2,3) \quad (4,5) \quad (6,7) \quad (8,9) \quad (10,11) \quad (12,13) \quad (14,15) \\
& (16,17) \quad (18,19) \quad (20,21) \quad (22,23) \quad (24,25) \quad (26,27) \quad (28,29) \quad (30,31)
\end{aligned}$$

We want a symmetry class to propagate through one round of the scheme with probability equal to one. It is easy to see that this condition imposes that the equality constraints at the left and at the right branch of the scheme must be the same (because of the intra-word rotations that are only present in the left branch of the scheme). That is, for any relation (i, j) with $0 \leq i, j \leq 15$, we must also have the relation $(i + 16, j + 16)$. In other words, a symmetry pattern is the same for the left and for the right branch. We thus only need to consider 16-word symmetry patterns.

To describe all possible symmetries, we start by fixing $(0, k)$, for a fixed k in $\{1, \dots, 15\}$. We then compute T backwards to identify the relations implied by $(0, k)$: the first substitution and xor encountered force us to have

$$(0, k) (4, k \oplus 4).$$

Then, the second substitution and the modular addition force to have (note that the intra-word rotations can be omitted since they leave the symmetry pattern unchanged)

$$(0, k) (4, k \oplus 4) (1, k \oplus 1) (5, k \oplus 5).$$

The third substitution and xor yield

$$\begin{aligned} &(0, k) \quad (4, k \oplus 4) \quad (1, k \oplus 1) \quad (5, k \oplus 5) \\ &(2, k \oplus 2) \quad (6, k \oplus 6) \quad (3, k \oplus 3) \quad (7, k \oplus 7). \end{aligned}$$

Finally, the last substitution and the modular addition imply

$$\begin{aligned} &(0, k) \quad (4, k \oplus 4) \quad (1, k \oplus 1) \quad (5, k \oplus 5) \\ &(2, k \oplus 2) \quad (6, k \oplus 6) \quad (3, k \oplus 3) \quad (7, k \oplus 7) \\ &(8, k \oplus 8) \quad (12, k \oplus 12) \quad (9, k \oplus 9) \quad (13, k \oplus 13) \\ &(10, k \oplus 10) \quad (14, k \oplus 14) \quad (11, k \oplus 11) \quad (15, k \oplus 15). \end{aligned}$$

Eventually, each symmetry that contains the relation $(0, k)$ —i.e., $x[0] = x[k]$ —also has the relations $(i, k \oplus i)$, $1, \dots, 15$. Therefore, we have 15 distinct wordwise symmetry classes, of the form

$$(i, k \oplus i), i = 0, \dots, 15$$

for $k \in \{1, \dots, 15\}$. Each class contains 2^{512} states. For example, the case $k = 1$ provides directly C_1 , and more generally $k = i$ corresponds to C_i .

4.3 Exploiting symmetric states for finding 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 $2^{1024-516-8} = 2^{500}$ 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^{501} computations of T . The cost of steps 3 and 4 depends on i and j ; but it is upper bounded by 2×2^{256} operations.

Thus, in any case, the total complexity is about 2^{501} calls to T . This attack, however, finds messages of unauthorized size (more than 2^{256} 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. 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 a class example C_i that is the easiest to reach, depending on the value of b :

- $5 \leq b < 9$: one of the best classes is C_1 , which gives $(1024 - 2 \times 4 \times 8)/2 = 480$ equations to verify
- $9 \leq b < 17$: one of the best classes is C_2 , which give $(1024 - 2 \times 8 \times 8)/2 = 448$ equations to verify
- $17 \leq b < 33$: one of the best classes is C_4 , which gives $(1024 - 2 \times 16 \times 8)/2 = 384$ equations to verify
- $33 \leq b < 65$: one of the best classes is C_8 , 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.

4.4 Exploiting symmetric states for finding collisions

We present a technique to find collisions for a weakened version of CubeHash, in which we modify the IV (initial state). 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 $C_1 \cap C_2 \cap C_4 \cap C_8$.

Suppose that the initial state of CubeHash $r/b-h$ is in $C_1 \cap C_8$, i.e. is of the form

AAAAAAAAA AAAAAAAAAA BBBBBBBBBB BBBBBBBBB .

If one hashes the $b2^{33}$ -byte message that contain only zeros, then each of the 2^{33} intermediate states is an element of $C_1 \cap C_2 \cap C_4 \cap C_8$. Assuming that T acts like a random permutation over this set, one will find two identical states with probability about 0.63, which directly gives a collision.

5 On the fixed points of T

In this section we let T be the 1-round transform of CubeHash. A fixed point for T^r , $r > 0$, is a state x that is left unchanged by T^r , i.e., $T^r(x) = x$. Recall that the average number of cycles of length k is $1/k$ for a random permutation. If T were a random permutation, T would thus have one fixed point. Noting that a cycle of length r gives r fixed points for T^r , we have that T^2 would have two fixed points (the one of T and one due to an average of $2 \times 1/2$ fixed points from cycles of length two); T^4 would have three fixed points (one for each cycle length in 1, 2, 4), etc. More generally, the average number of fixed points for T^n would be the number of divisors of n , if T were a random permutation.

Note that each symmetry class represents a class of cycles of T , and that the 15 symmetry classes give 67 distinct subsets. Modeling T as a random permutation over each of those subsets, one expects 67 fixed points. This gives for T^8 $1 + 4 \times 67 = 269$ fixed points, where 4 is the number of divisors of 8, i.e., of cycles length that give fixed points for T^8 . Note that this results assumes a random behavior of T with respect to fixed points over the 67 subsets considered.

Finding examples of fixed points seems difficult, however: the zero state $x[0] = \dots = x[31] = 0$ is a trivial fixed point for T , and thus also for T^n , $n \geq 0$. Among the states of the form $x[0] = \dots = x[15]$, $x[16] = \dots = x[31]$, the only nontrivial fixed point is the state with $x[0] = 54E5FC8A$ and $x[16] = 84FE49D2$.

6 Truncated differentials over T

This section shows how to detect non-randomness 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 $x[16] \dots x[31]$ diffuse less in the first rounds than $x[0] \dots x[15]$. 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 input. But it holds for the input

```
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.

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 [11, §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 difference, we could detect non-randomness on up to 7 rounds (now this difference is introduced three rounds later). Note that in a previous version of this article [2], 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 [4]. 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 non-randomness after up to 5 rounds.

References

1. Jean-Philippe Aumasson. Collision for CubeHash2/120-512. NIST mailing list, 4 Dec 2008, 2008. <http://ehash.iaik.tugraz.at/uploads/a/a9/Cubehash.txt>.
2. Jean-Philippe Aumasson, Willi Meier, Mara Naya-Plasencia, and Thomas Peyrin. Inside the hypercube. Cryptology ePrint Archive, Report 2008/486, 2008. version 20081124:132635.
3. Daniel J. Bernstein. CubeHash appendix: complexity of generic attacks. Submission to NIST, 2008.
4. Daniel J. Bernstein. CubeHash attack analysis (2.B.5). Submission to NIST, 2008.
5. Daniel J. Bernstein. CubeHash specification (2.B.1). Submission to NIST, 2008.
6. Eric Brier, Shahram Khazaei, Willi Meier, and Thomas Peyrin. Attack for CubeHash-2/2 and collision for CubeHash-3/64. NIST mailing list (local link), 2009. http://ehash.iaik.tugraz.at/uploads/3/3a/Peyrin_ch22_ch364.txt.
7. Eric Brier and Thomas Peyrin. Cryptanalysis of CubeHash. Available online, 2009. <http://thomas.peyrin.googlepages.com/BrierPeyrinCubehash.pdf>.
8. Wei Dai. Collisions for CubeHash1/45 and CubeHash2/89. Available online, 2008. <http://www.cryptopp.com/sha3/cubehash.pdf>.
9. Persi Diaconis and Frederick Mosteller. Methods for studying coincidences. *Journal of the American Statistical Association*, 84(408):853–861, 1989.
10. Antoine Joux. Multicollisions in iterated hash functions. application to cascaded constructions. In *CRYPTO*, 2004.
11. NIST. SP 800-22, a statistical test suite for random and pseudorandom number generators for cryptographic applications, 2001.
12. Kazuhiro Suzuki, Dongyu Tonien, Kaoru Kurosawa, and Koji Toyota. Birthday paradox for multi-collisions. In *ICISC*, 2006.