

Malicious Keccak

Paweł Morawiecki^{1,2}

¹ Institute of Computer Science, Polish Academy of Sciences, Poland

² Section of Informatics, University of Commerce, Kielce, Poland

Abstract. In this paper, we investigate Keccak — the cryptographic hash function adopted as the SHA-3 standard. We propose a malicious variant of the function, where new round constants are introduced. We show that for such the variant, collision and preimage attacks are possible. We also identify a class of weak keys for the malicious Keccak working in the MAC mode. Ideas presented in the paper were verified by implementing the attacks on the function with the 128-bit hash.

Keywords: cryptanalysis, Keccak, SHA-3, malicious hashing

1 Introduction

A malicious variant of a cryptographic primitive is the one with a backdoor. An attacker who knows the backdoor can easily manipulate or even totally compromise the algorithm's security. A holy grail for all intelligence agencies is to have backdoors, which are extremely hard to detect and, at the same time, easy to use, widely applicable. Recent revelations by Edward Snowden have shown that the NSA has deliberately inserted a backdoor in the standardized pseudorandom number generator Dual_EC_DRBG [3]. This backdoor gives the knowledge of the internal state of the generator and consequently the attacker can predict future keystream bits. At the time of Snowden's revelations, NIST actually recommended Dual_EC_DRBG, so there are speculations that the NIST standards are manipulated by NSA, such as the NIST's recommended elliptic curve constants [8].

Malicious cryptography can be also an issue in the commercial cryptography software or in some services relying on supposedly strong cryptographic algorithms. Imagine a multimedia online service, where a user buys and downloads videos and music. Clearly such the service should have a secure and reliable authentication mechanism. If the authentication is backdoored (e.g., a malicious hash function has been inserted), an anonymous worker can blackmail a company by showing an evidence of the system weakness.

In the public domain, very few papers have been published regarding the malicious cryptography. One of the first attempts was cryptovirology project [12]. Another interesting try was to modify Sboxes of CAST and LOKI and hide linear relations inside [11]. Recently, a more formal treatment on malicious hashing was given, along with the malicious variant of SHA-1 [1]. (The constants of SHA-1 were modified in such a way that the collision attack is possible.) In [7], the backdoored version of Streebog is presented. Streebog is a new Russian cryptographic hash standard (GOST R 34.11-2012) and its malicious variant has modified constants, which allows generating collisions, with an aid of differential cryptanalysis.

In this paper we focus on Keccak — the cryptographic hash function adopted as the SHA-3 standard [5]. We propose a new set of constants, which leads to collision and preimage attacks. Additionally, for the malicious Keccak, we identify a class of weak keys. If a key belongs to the class, the forgery attack is possible. An idea behind new, malicious constants is to exploit the symmetric nature of the Keccak permutation. An inspiration comes from the previous attacks on Keccak such as rotational cryptanalysis [9] and internal differential cryptanalysis [6]. Both of these works take advantage of symmetry in the permutation and the low Hamming weights of the constants.

The paper is organised as follows. In Section 2, we describe Keccak and give the malicious constants. In Section 3, first we give the overview of internal differential cryptanalysis in context of the backdoored Keccak. Then, we explain how to attack the malicious Keccak, particularly how to find collisions, preimages and mount the forgery attack on Keccak working in the MAC mode. Before the paper is concluded, we discuss a different set of constants and its impact on the attack complexities.

2 Description of Keccak

Keccak is not a single algorithm but rather a family of algorithms called sponge functions [4]. A sponge function can be treated as a generalisation of the cryptographic hash function with infinite output. It can provide many functionalities, namely a hash function, a stream cipher, the keyed MAC or a pseudorandom bit generator. In this section we give a brief description of Keccak, necessary for understanding the malicious variant we propose and the attacks on the function. An interested reader finds all the details on Keccak in its original specification [5].

Keccak has the b -bit internal state, which is divided into two parts, that is the r -bit bitrate and the c -bit capacity ($r + c = b$). First, the state is filled with all 0's and a message is divided into r -bit blocks. Next, Keccak processes the message in two phases. The first phase is called the absorbing phase, where the r -bit message blocks are absorbed (XORed) into the state, interleaved with calls of the internal permutation Keccak-f. Once all message blocks are processed, the second phase (called squeezing) starts. In this phase, the first r bits of the state are returned (as hash bits for example). If a desired output is longer than r bits, then the internal permutation is called and then another r bits can be squeezed.

Figure 1 shows how the Keccak state is organised. Terms which denote a given part of the state were introduced by the Keccak designers. For the pseudo-code, it is more convenient to represent the state as the two-dimension array $S[x, y]$, where an element of the array is the 64-bit lane. The default variant of Keccak has the 64-bit lane, but smaller variants (such as the 400-bit state with the 16-bit lane) are also defined. A number of rounds is determined by a size of the state. For the default 1600-bit state, a number of rounds is 24.

In the Keccak permutation all rounds are the same except for the ι step (round-dependent constants XORed into the state). Below we provide a pseudo-code of a single round. Steps in the permutation are denoted by Greek letters.

```

Round(A,RC) {
   $\theta$  step
  C[x] = A[x,0] xor A[x,1] xor A[x,2] xor A[x,3] xor A[x,4],           forall x in (0...4)
  D[x] = C[x-1] xor rot(C[x+1],1),                                     forall x in (0...4)
  A[x,y] = A[x,y] xor D[x],                                           forall (x,y) in (0...4,0...4)

   $\rho$  step                                                             forall (x,y) in (0...4,0...4)
  A[x,y] = rot(A[x,y], r[x,y]),

   $\pi$  step                                                             forall (x,y) in (0...4,0...4)
  B[y,2*x+3*y] = A[x,y],

   $\chi$  step                                                             forall (x,y) in (0...4,0...4)
  A[x,y] = B[x,y] xor ((not B[x+1,y]) and B[x+2,y]),

   $\iota$  step

```

```
A[0,0] = A[0,0] xor RC
```

```
return A }
```

In the pseudo-code operations on indices are done modulo 5. The state of the permutation is denoted by A . There are some intermediate variables such as $B[x,y]$, $C[x]$, $D[x]$. The rotation offsets are $r[x,y]$, whereas RC are the round constants. We denote bitwise rotation operation by $\text{rot}(W,m)$. It moves a bit at position i into position $i + m$ in the lane W ($i + m$ are done modulo the lane size).

θ is a linear operation, a main source of diffusion in the algorithm. ρ is a permutation that mixes bits of a lane, while π permutes the whole lanes. χ can be viewed as a layer of the 5-bit Sboxes. The last step is ι , which XORs the round constant into the first lane. The values of the round constants are generated by a simple linear feedback shift register (LFSR) [5].

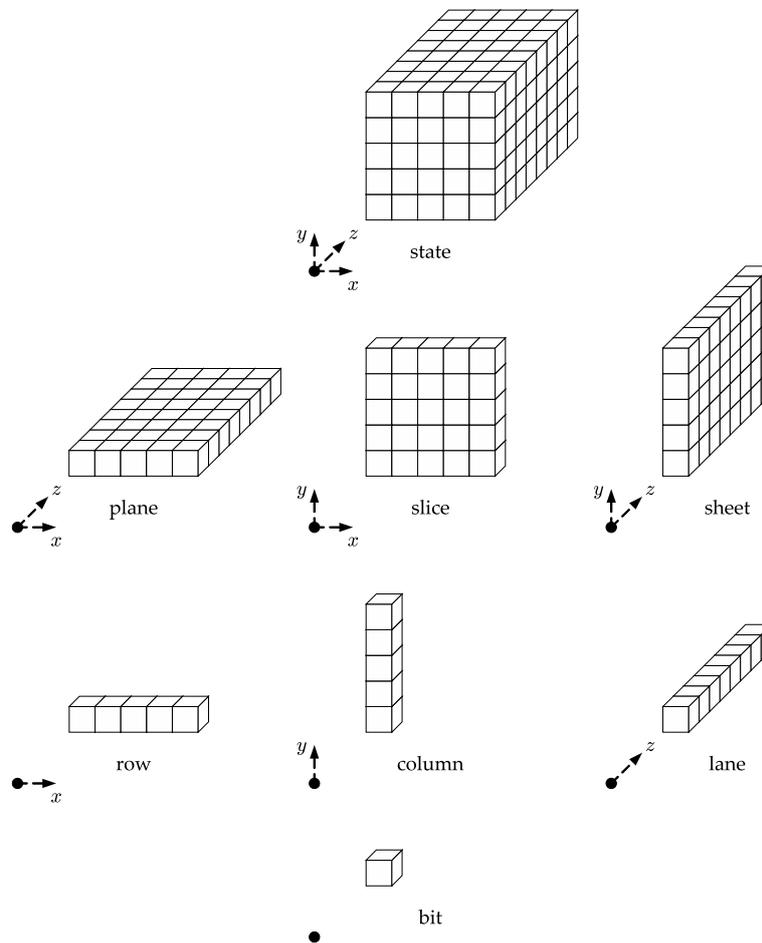


Fig. 1. Pieces of the Keccak state[5]

2.1 Malicious Keccak

For the malicious variant of Keccak, we change the round constants. All other steps and parameters in the algorithm remain the same. Our malicious constants are also generated with

a simple shift register. The register is seeded with "SHA3SHA3" (64-bit constant encoded in ASCII). The subsequent round constants are obtained by rotation of the register by one bit. The constants, given in hexadecimal notation, are as follows.

RC[0] := 5348413353484133	RC[1] := A9A42099A9A42099	RC[2] := D4D2104CD4D2104C
RC[3] := 6A6908266A690826	RC[4] := 3534841335348413	RC[5] := 9A9A42099A9A4209
RC[6] := CD4D2104CD4D2104	RC[7] := 66A6908266A69082	RC[8] := 3353484133534841
RC[9] := 99A9A42099A9A420	RC[10] := 4CD4D2104CD4D210	RC[11] := 266A6908266A6908
RC[12] := 1335348413353484	RC[13] := 099A9A42099A9A42	RC[14] := 04CD4D2104CD4D21
RC[15] := 8266A6908266A690	RC[16] := 4133534841335348	RC[17] := 2099A9A42099A9A4
RC[18] := 104CD4D2104CD4D2	RC[19] := 08266A6908266A69	RC[20] := 8413353484133534
RC[21] := 42099A9A42099A9A	RC[22] := 2104CD4D2104CD4D	RC[23] := 908266A6908266A6

All 24 constants are different, however, a careful reader notices that they are symmetric. This symmetry plays a vital role in attacks on the malicious Keccak.

3 Attacks on Malicious Keccak

In the attacks on the malicious Keccak we use a variant of differential cryptanalysis namely internal differential cryptanalysis. In typical differential attacks, we consider two different plaintexts and follow an evolution of the differences between them. In case of internal differential attacks only one plaintext is considered and we trace a statistical evolution of the differences between its parts. Such analysis was first proposed by Peyrin [10] in the attack on the Grøstl hash function. In [6], Dinur et al. showed that internal differentials could be also used to produce collisions against the round-reduced Keccak.

There is a particular property of Keccak, already noticed by the designers [5], which makes Keccak a promising candidate for internal differential cryptanalysis. That is four out of its five internal mappings (all but ι) are translation invariant in the direction of the z axis. Namely, if one state S is the rotation of another state S' with respect to the z -axis (i.e., $S'[x][y][z] = S[x][y][z+i]$, for some value of i), then applying to them any of the θ, ρ, π, χ operations maintains this property. This leads to the following observation. If we divide the state on two halves along z -axis, where both halves are the same, θ, ρ, π, χ operations do not destroy the symmetry. What does destroy the symmetry is the ι step, namely XORing the round constants to the state. Here comes a natural question, that is, what happens if we craft the constants such as the above-mentioned property works for every step in the algorithm. This is how we come up with an idea of malicious constants in Keccak. If constants are also symmetrical, as the ones proposed in Section 2.1, the symmetry of the state is kept through all the steps and we can exploit it to devise attacks on the Keccak hash function.

3.1 Collision Attack

Here we present the collision attack on the malicious Keccak with the l -bit hash. Applying the standard birthday attack, we can find collisions with $2^{l/2}$ calls. However, for the malicious Keccak, a collision can be found with the $2^{l/4}$ effort. In our attack, instead of calling the algorithm with random messages, we use the messages which have symmetric structure. That is, once the message is absorbed into the state, both halves in each lane are the same. This property, as already explained, would be preserved at every step of the algorithm. In particular, hashes would be also symmetric. It means that, in fact, we search for collisions for the $(l/2)$ -bit hash. Once found, we are sure that the second parts of lanes in the hash also collide. Hence, the cost of the collision attack is $2^{l/4}$.

There are some technicalities worth discussing. First, the attacker has to consider padding. The last 62 bits of the message has to contain all 1's and the length of the message is equal to the bitrate r minus 2. This way the message is padded with two 1's and the last lane in the bitrate part of the state is filled with all 1's (clearly a symmetric lane).

The attack scales naturally for longer or shorter hashes, for example, the attack on Keccak with the 512-bit hash would cost $2^{512/4} = 2^{128}$. If a hash length is not a multiple of the lane size (64), then the attack does not fully exploit the symmetry property. So, for example, the cost of the collision attack for the 256-bit and 224-bit hash is the same. It is because for the 224-bit hash we don't get an extra 32 bits 'for free' in the last lane of the hash string.

We implemented the attack for the Keccak variant with the 1024-bit bitrate and the 128-bit hash. As expected, after about $2^{128/4}$ calls we found a collision. The colliding (padded) messages and the hash are given below.

Table 1. An example of a collision for the malicious Keccak.

m	813e344a813e344a 78d30cf978d30cf9 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 1111111111111111
m'	256a4f71256a4f71 e788dc79e788dc79 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 1111111111111111
hash	a012dcb9a012dcb9 810c79e0810c79e0

3.2 Preimage Attack

In a similar way to the collision attack, we are able to find preimages for the symmetric hashes. Typically, to convince a third party that we can mount the preimage attack, we show a preimage for some non-random hashes, e.g., all 0's or all bytes identical. Such hashes would be covered by our preimage attack as they are clearly symmetric. In the attack against the malicious Keccak with l -bit hash, we search, in fact, the preimage for $l/2$ -bit hash, as the second half of the bits are guaranteed to be the same, due to the symmetric property of the malicious Keccak. Therefore, the preimage attack complexity is $2^{l/2}$, whereas the exhaustive search costs 2^l . As in the collision search, our preimage attack exploits its full potential when a given hash is a multiple of 64.

We found a preimage for the malicious Keccak with the 64-bit hash, where all hash bits are 0.

Table 2. A preimage for a given 64-bit hash.

m	7187c4197187c419 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 0000000000000000 1111111111111111
hash	0000000000000000

3.3 Weak Keys in Keccak-MAC

Keccak can also work in the keyed mode, particularly providing a message authentication code (MAC). The hash-based MAC involves a cryptographic hash function in combination with a secret key and a typical solution is HMAC proposed by Bellare et al. [2]. However, in the case of Keccak, a MAC functionality can be done in a more straightforward way than the nested approach of HMAC. A secret key is simply prepended to a message, so for Keccak with the 128-bit key first two lanes of the state are occupied by key bits. Then, the state is processed and a tag is obtained (squeezed).

A secure MAC algorithm should have two main properties. First, it should be infeasible to mount the key-recovery attack, even if the attacker has many valid message-tag pairs. Second, the attacker should not be able to create a forgery, namely, to show a valid message-tag pair (M, T) for a message M that has not been previously authenticated.

For the malicious Keccak the forgery attack is possible when a key belongs to the weak keys class. The class consists of symmetric keys, that is the first part of a lane is the same as the second one. To attack the Keccak-MAC with the 128-bit key and 128-bit tag, we provide 2^{64} different pairs of (M, T) . Each pair has the same, chosen symmetric tag and different (but also symmetric) message lanes. Because of the symmetric property, one of the 2^{64} chosen pairs will be actually the valid pair, whereas for the secure MAC algorithm we would need 2^{128} tries to provide the valid, previously unseen, message-tag pair. Please note that checking whether a key belongs to the weak keys class requires just a single call of the Keccak-MAC. If a key is in the class, the tag will be symmetric. (There might be false positives but their probability is only 2^{-64} and they can be easily verified with another call of the algorithm, with different message bits.)

3.4 Other Symmetries

The malicious constants we propose are symmetric, that is both 32-bit parts of a lane are the same. We can exploit the symmetry even further and divide a given 64-bit lane on shorter, identical blocks. A size of a block could be 32 (as in our malicious set), 16, 8, 4, 2 — divisors of 64. If a block is smaller, the cost of the attacks is also smaller. This is because we get more bits ‘for free’. So, for example, the malicious constants with a block size 16 (four identical blocks in a lane) decreases the cost of our preimage attack to $2^{l/4}$. A drawback of shorter blocks is that we may not guarantee distinct constants for all 24 rounds and they would look less and less random.

3.5 More ‘nothing up my sleeve’ Constants

To convince a user that the chosen constants are indeed ‘nothing up my sleeve’, we can change the last round constant and set it as, for example, 64 bits taken from π or other well known, mathematical constant. Certainly, such a constant is not a symmetric string of bits. But please note that the constant is XORed into to the state as the last step of the permutation, so it only acts as a ‘mask’. If there were some subsequent steps in the algorithm, then such the constant would be a beginning of breaking the symmetry in the state. For the collision and forgery attack, this new constant does not affect the attack, we can always ‘unmask’ obtained hashes (tags) and observe symmetric hashes (tags). In case of the preimage attack, where we want to show a preimage for a chosen (e.g. all 0’s) hash, the new, non-symmetric constant spoils the attack or at least it makes the attack less convincing. We could still find preimages, but a chosen hash would have been ‘affected’ by the π constant. For example, we could find a preimage for a hash with its first 64-bits equal to π concatenated with the all-zero vector.

4 Conclusion

In this paper, we investigated a malicious version of the Keccak hash function. We propose a new set of constants which exploits a symmetric nature of the Keccak permutation and allows the collision, preimage and forgery attacks, substantially faster than generic ones. Our malicious set of constants has some space for bringing more ‘nothing up my sleeve’ numbers. One can take a well known mathematical constant and set it in the 24th round (or even 23rd round). However, it limits an efficiency of the attack since the symmetry of the state would be disturbed. Our results do not threaten the original SHA-3 standard, however, they clearly show a user should be very careful with the ‘enhanced’ or ‘personalised’ variants of the algorithm, advertised in a commercial cryptography package.

Acknowledgement

Project was financed by Polish National Science Centre, project DEC-2013/09/D/ST6/03918.

References

1. Albertini, A., Aumasson, J., Eichlseder, M., Mendel, F., Schläffer, M.: Malicious hashing: Eve’s variant of SHA-1. In: Selected Areas in Cryptography - SAC 2014 - 21st International Conference, Montreal, QC, Canada, August 14-15, 2014, Revised Selected Papers. pp. 1–19 (2014)
2. Bellare, M., Canetti, R., Krawczyk, H.: Message Authentication Using Hash Functions: the HMAC Construction. *CryptoBytes* 2(1), 12–15 (1996)
3. Bernstein, D.J., Lange, T., Niederhagen, R.: Dual EC: A Standardized Back Door. *Cryptology ePrint Archive*, Report 2015/767 (2015), <http://eprint.iacr.org/>
4. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic Sponges, <http://sponge.noekeon.org/CSF-0.1.pdf>
5. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Keccak Sponge Function Family Main Document, <http://keccak.noekeon.org/Keccak-main-2.1.pdf>
6. Dinur, I., Dunkelman, O., Shamir, A.: Collision Attacks on Up to 5 Rounds of SHA-3 Using Generalized Internal Differentials. In: Fast Software Encryption - 20th International Workshop, FSE 2013, Singapore, March 11-13, 2013. Revised Selected Papers. pp. 219–240 (2013)
7. Federal Agency on Technical Regulation and Metrology (GOST): Gost r 34.11-2012: Streebog hash function. www.streebog.net (2012)
8. Kobitz, N., Menezes, A.: A riddle wrapped in an enigma. *Cryptology ePrint Archive*, Report 2015/1018 (2015), <http://eprint.iacr.org/>
9. Morawiecki, P., Pieprzyk, J., Srebrny, M.: Rotational cryptanalysis of round-reduced Keccak. In: Fast Software Encryption. LNCS, Springer (2013)
10. Peyrin, T.: Improved Differential Attacks for ECHO and Grøstl. In: Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings. pp. 370–392 (2010)
11. Rijmen, V., Preneel, B.: A Family of Trapdoor Ciphers. In: Fast Software Encryption, 4th International Workshop, FSE '97, Haifa, Israel, January 20-22, 1997, Proceedings. pp. 139–148 (1997)
12. Young, A.L., Yung, M.: Malicious cryptography - exposing cryptovirology. Wiley (2004)