

Collision Attack on 5 Rounds of Grøstl [★]

Florian Mendel¹, Vincent Rijmen², and Martin Schl affer¹

¹ IAIK, Graz University of Technology, Austria

² Dept. ESAT/COSIC, KU Leuven and Security Dept., iMinds, Belgium

Abstract. In this article, we describe a novel collision attack for up to 5 rounds of the Grøstl hash function. This significantly improves upon the best previously published results on 3 rounds. By using a new type of differential trail spanning over more than one message block we are able to construct collisions for Grøstl-256 on 4 and 5 rounds with complexity of 2^{67} and 2^{120} , respectively. Both attacks need 2^{64} memory. Due to the generic nature of our attack we can even construct meaningful collisions in the chosen-prefix setting with the same attack complexity.

Keywords: hash functions, SHA-3 candidate, Grøstl, collision attack

1 Introduction

In the last few years the cryptanalysis of hash functions has become an important topic within the cryptographic community. Especially the collision attacks on the MD4 family of hash functions (MD5, SHA-1) have weakened the security assumptions of these commonly used hash functions [26–28]. As a consequence NIST has decided to organize a public competition in order to design a new hash function, leading to the selection of Keccak as SHA-3 [19].

During the SHA-3 competition, the three classical security requirements (collision-, preimage- and second-preimage resistance) were not the main target of cryptanalytic attacks. Most results were published on building blocks such as the compression function, block cipher or permutation used in a hash function. Additionally, many distinguishers on these building blocks with minor relevance in practice were considered. Although these results are important from a theoretical point of view, vulnerabilities that can be exploited for the hash function are certainly more important.

In this work, we present new results on the collision resistance of the SHA-3 finalist Grøstl. Grøstl is an iterated hash function based on design principles very different from those used in the MD4 family. The compression function of Grøstl is built from two different permutations that follow the design strategy of the Advanced Encryption Standard (AES) [3, 17]. The simple construction of the compression function and the byte-oriented design of Grøstl facilitates the security analysis. In the last years Grøstl has received a large amount of cryptanalysis. However, most of the analysis focus on the building blocks of Grøstl and only a few results have been published for the hash function so far.

[★] © IACR 2014. This article is the final version submitted by the authors to the IACR and to Springer-Verlag on 2014-04-30, which appears in the proceedings of FSE 2014.

Related Work. `Grøst1` is one of the SHA-3 candidates that has probably received the largest amount of cryptanalysis during the competition. Security analysis of `Grøst1` was initiated by the design team itself which led to the rebound attack [15]. Since then, several improvements to the rebound attack technique have been made, leading to new results on both the hash function [16] and its underlying components [6, 14, 22]. For the final round, `Grøst1` was been tweaked to thwart internal differential attacks [7, 21] and to reduce the impact of the rebound attack and its extensions.

The best published attacks on the final version of both `Grøst1-256` and `Grøst1-512` are collision attacks on 3 rounds of the hash function and on 6 rounds of the compression function [9, 23]. Preimage attacks for the compression function of `Grøst1-256` and `Grøst1-512` have been shown in [29] for 5 and 8 rounds, respectively. Additionally, non-random properties of the `Grøst1` permutation have been discussed in [1, 8]. For a detailed overview of the existing attacks on `Grøst1` we refer to the ECRYPT II SHA-3 Zoo [4].

Our Contribution. By using a new type of differential trail we are able to show collision attacks on `Grøst1` for up to 5 rounds. The extension becomes possible by considering differential trails spanning over more than one message block to iteratively cancel differences in the chaining variable. Our new attack combines ideas of the attack on SMASH [13] with the rebound attack [15] on `Grøst1`. A similar approach has also been used in the attack on Grindahl [20]. The results are collision attacks on the `Grøst1-256` hash function reduced to 4 and 5 rounds with a complexity of 2^{67} and 2^{120} , respectively. Both attacks have memory requirements of 2^{64} . Note that the best previously known collision attack on the `Grøst1` hash function was on 3 rounds with a complexity of 2^{64} [23]. We want to note that the same attack also applies to 5 rounds of `Grøst1-512`.

Additionally, we show that due to the generic nature of our attack we can construct collisions in the chosen-prefix setting with the same complexity. It has been demonstrated in [24, 25] that chosen-prefix collisions can be exploited to construct colliding X.509 certificates and a rogue CA certificate for MD5. Note that in most cases constructing such collisions is more complicated than constructing (random) collisions. Our results and related work for the `Grøst1` hash function are shown in Table 1.

Table 1: Collision Attacks on the `Grøst1-256` hash function.

rounds	complexity	memory	reference
3	2^{64}	-	[23]
4	2^{67}	2^{64}	this work
5	2^{120}	2^{64}	this work

Outline. The paper is structured as follows. In Section 2, we give a short description of the `Grøstl` hash function. The basic attack strategy and the collision attack for 4 rounds of the hash function is presented in Section 3. In Section 4, we describe the extension of the attack to 5 rounds, and in Section 5 the construction of meaningful collisions is discussed. Finally, we conclude in Section 6.

2 Short Description of `Grøstl`

The hash function `Grøstl` [5] was one of the 5 finalists in the SHA-3 competition [18]. `Grøstl` is an iterated hash function with a compression function built from two distinct permutations P and Q , which are based on the AES design principles. In the following, we describe the components of the `Grøstl` hash function in more detail.

2.1 The Hash Function

The two main variants, `Grøstl-256` and `Grøstl-512` are used for hash output sizes of $n = 256$ and $n = 512$ bits. The hash function first pads the input message M and splits the message into blocks m_1, m_2, \dots, m_t of ℓ bits with $\ell = 512$ for `Grøstl-256`, and $\ell = 1024$ for `Grøstl-512`. The message blocks are processed via the compression function $f(h_{i-1}, m_i)$ and output transformation $\Omega(h_t)$. The size of the chaining value h_i is ℓ bits as well.

$$\begin{aligned} h_0 &= IV \\ h_i &= f(h_{i-1}, m_i) \quad \text{for } 1 \leq i \leq t \\ h &= \Omega(h_t). \end{aligned}$$

The compression function f is based on two ℓ -bit permutations P and Q (sometimes denoted by P_ℓ and Q_ℓ) and is defined as follows:

$$f(h_{i-1}, m_i) = P(h_{i-1} \oplus m_i) \oplus Q(m_i) \oplus h_{i-1}.$$

The output transformation Ω is applied to h_t to give the final hash value h of size n , where $\text{trunc}_n(x)$ discards all but the least significant n bits of x :

$$\Omega(h_t) = \text{trunc}_n(P(h_t) \oplus h_t).$$

2.2 The Permutations P and Q

The two permutations P and Q are designed according to the wide trail strategy [2] and their structure is very similar to the AES. In `Grøstl-256` each permutation updates an 8×8 state of 64 bytes in 10 rounds. In one round, the round transformation updates the state by means of the sequence of transformations

$$\text{MB} \circ \text{SH} \circ \text{SB} \circ \text{AC} .$$

In the following, we briefly describe the round transformations of P and Q used in the compression function f in more detail.

AddRoundConstant (AC). In this transformation, the state is modified by combining it with a round constant with a bitwise `xor` operation. Different constants are used for the permutations P and Q .

SubBytes (SB). The `SubBytes` transformation is the same for P and Q and is the only non-linear transformation of the permutations. It is a permutation consisting of an S-box applied to each byte of the state. The 8-bit S-box is the same as in the AES with good cryptographic properties against differential and linear attacks. For a detailed description of the S-box, we refer to [17].

ShiftBytes (SH). The `ShiftBytes` transformation is a byte transposition that cyclically shifts the rows of the state over different offsets. The `ShiftBytes` transformation is different for the two permutations P and Q .

MixBytes (MB). The `MixBytes` transformation is a permutation operating on the state column by column. To be more precise, it is a left-multiplication by an 8×8 MDS matrix over \mathbb{F}_{2^8} . The coefficients of the matrix are determined in such a way that the *branch number* of `MixBytes` (the smallest nonzero sum of active input and output bytes of each column) is 9, which is the maximum possible for a transformation with these dimensions. This transformation is the same for both permutations P and Q .

2.3 Alternative Description of `Grøst1`

To simplify the description of attack in the following sections, we use an equivalent alternative description of `Grøst1`. Let P' and Q' denote the permutation P and Q without the last application of `MixBytes`. Then, by setting

$$\begin{aligned} h'_0 &= \text{MB}^{-1}(\text{IV}) \\ h'_i &= P'(\text{MB}(h'_{i-1}) \oplus m_i) \oplus Q'(m_i) \oplus h'_{i-1} \quad \text{for } 1 \leq i \leq t \\ h &= \Omega(\text{MB}(h'_t)) \end{aligned}$$

with $h_i = \text{MB}(h'_i)$, we get an equivalent description of `Grøst1`, where the last `MixBytes` transformation of the permutations has been swapped with the XOR operation of the feed-forward.

3 Collision Attack for 4 Rounds of `Grøst1`

To get improved attacks on the `Grøst1` hash function, we view `Grøst1` as a strengthened variant of SMASH [10]. The essential difference between the two designs is that `Grøst1` employs a second nonlinear permutation Q , where SMASH employs a scaling by the constant θ , i.e. a linear map (see Fig. 1). The hash function SMASH has been broken by subsequently controlling the output difference of the compression function using the linearity of θ . After the application

of 257 respectively 513 message blocks, a colliding output difference can be constructed [13]. In this section, we show how to achieve the same for 4 rounds of `Grøst1` by having differences in only one permutation.




Fig. 1: The compression function of (a) SMASH and (b) `Grøst1`.

3.1 The Second-Preimage Attack on SMASH

The second-preimage attack on SMASH presented in [13] is based on the technique of *controllable output differences* for the compression function. By carefully selecting consecutive message blocks, an attacker can step-by-step convert an arbitrary starting difference in the chaining variable into an arbitrary output difference. The attack is deterministic and the number of consecutive controllable message blocks is equal to the length of the chaining variable. The nonlinearity of f is made ineffective by strictly controlling its input differences and values. Controlling the input values of f implies that the input values of the linear map are determined. Fortunately, for a linear map it suffices to know the input *difference* to compute the output difference. The output difference of the linear map is controlled by the number of message blocks.

3.2 Application to Reduced `Grøst1`

At the first sight, the attack on SMASH does not apply to `Grøst1`, because the strong nonlinearity of P and Q makes it difficult to control the output differences of both permutations. However, by having no differences in Q , we can use the whole freedom of the message block to control the differential propagation in P . Since we cannot control the differences completely, we need to apply a variation of the technique on SMASH, to get a zero output difference at the compression function.

Our attack will start from an arbitrary difference in the chaining variable and convert it into an output difference equal to zero after 9 steps. The first message block can be selected arbitrarily. The only requirement is a difference in the message. The next 8 message blocks are fully controlled by the attacker and must not contain any differences. Then, each of the 8 message blocks is used to cancel one eighth of the differences at the output of the compression function to result in a collision at the end (see Fig. 2).

3.3 Details of the Attack

To simplify the description of the attack we use the alternative description of `Grøst1` given in Section 2.3. Since the last `MixBytes` transformation is moved out of the compression function, the limited set of differences at the output are more clearly visible.

The core of our collision attack on the reduced hash function are truncated differential trails with only 8 active bytes at the output of P' . Two full active states are placed at the beginning and the number of active bytes for the 4-round trail are given as follows:

$$64 \xrightarrow{r_1} 64 \xrightarrow{r_2} 8 \xrightarrow{r_3} 8 \xrightarrow{r_4} 8. \quad (1)$$

For such a truncated trail, we can construct a pair following the trail with an amortized complexity of 1 (even for a given input differences). We postpone the detailed explanation how to do so until Section 3.4.

The high-level overview of the 4-round attack is shown in Fig. 2. In each iteration, the differences in 8 bytes are canceled. Since this has a probability 2^{-64} , we need to compute 2^{64} pairs for P' (for the given input differences) to find a *right* pair that result in the desired output difference. The attack can then be summarized as follows:

1. Choose arbitrary message blocks m_1, m_1^* and compute h'_1 . Repeat this until one gets a full active state in h'_1 . Note that randomly selected m_1, m_1^* produce a full active state in h'_1 with probability at least $3/4$.
2. Use a *right* pair for P' following the trail of (1) to cancel 8 bytes of the difference in the state h'_2 , cf. Fig. 2.
3. Use a *right* pair for P' for a rotated variant of the trail of (1) to cancel another 8 bytes of the difference in the state h'_3 .
4. Repeat steps 3-4 in total 8 times until a collision has been found in h'_9 .

The complexity of the attack is 8 times finding a *right* pair for P' to iteratively cancel the difference in the state h'_2, \dots, h'_9 . We will show in the following section that such a *right* pair can be constructed with complexity of 2^{64} , resulting in a total attack complexity of $8 \cdot 2^{64} = 2^{67}$.




Fig. 2: Overview of the attack on 4 rounds.

3.4 Finding a Right Pair for P'

In this section, we show how to find a *right* pair for P' reduced to 4 rounds following the truncated differential trail in (1) using the rebound attack [15]. Note that the input difference is fixed by $\text{MB}(h'_{i-1})$ and we target an output difference such that 8 bytes of the difference in h'_i can be canceled. Unlike the classical rebound attack, the inbound phase is placed at the beginning and covers the first two rounds, while the outbound phase covers the last two rounds.

Using super-box matches [6, 11, 12], we can find 2^{64} pairs (solutions) for the inbound phase with a complexity of 2^{64} in time and memory. In the outbound phase, all these pairs will follow the 4 round truncated differential trail with a probability of 1 and one of these 2^{64} pairs will match the desired output difference (condition on 64 bits). In other words, using the rebound attack we can find a *right* pair for P' with a complexity of 2^{64} in time and memory.




Fig. 3: Truncated differential trail for P' used in the attack on 4 rounds.

In order to make the subsequent description of the rebound attack easier, we swap the SubBytes and ShiftBytes transformation in the first round of permutation P' (see Fig. 3). Note that this can always be done without affecting the output of the round. Then, the attack can be summarized as follows:

1. Compute the input difference of the permutation (P'_0) forward to state T_1^* .
2. Compute all 2^{64} differences of state P'_2 backward to state T_2^* and store them in a list L .
3. Connect the single difference of state T_1^* with the 2^{64} differences of state T_2^* using independent super-box matches. For each column $c = \{0, 1, \dots, 7\}$ we proceed as follows:
 - (a) Take all 2^{64} values for column c of state T_1^* and compute both values and differences forward to column c of state T_2^* .
 - (b) Check for matching 8-byte column differences in list L . Since we compute 2^{64} differences forward and have 2^{64} entries in L , we get 2^{64} solutions (differences and values) for the match. We update L to contain these 2^{64} solutions.
4. For each column and thus, for the whole inbound phase the number of resulting solutions is 2^{64} . The total complexity is 2^{64} in time and memory.

Since the truncated differential trail in the outbound part (the last 2 rounds) has probability 1, we get in total 2^{64} pairs following the truncated differential trail and one of these pairs is expected to be a right pair, i.e. result in the desired output difference (condition on 64 bits).

4 Extending the Attack to 5 Rounds

In this section, we present a collision attack for the `Grøst1-256` hash function reduced to 5 rounds with a complexity of about 2^{120} and memory requirements of 2^{64} . The attack is an extension of the attack on 4 rounds. However, since the freedom in finding right pairs for the 5-round trail is limited, we need more message blocks for the attack to succeed. In the attack, we use the following sequence of active bytes in P' (cf. Fig. 4):

$$64 \xrightarrow{r_1} 64 \xrightarrow{r_2} 8 \xrightarrow{r_3} 1 \xrightarrow{r_4} 8 \xrightarrow{r_5} 8. \quad (2)$$

However, it is important to note that for this truncated differential trail (with a fixed input difference) only 2^8 pairs exist, in contrast to 2^{64} for the 4 round trail. This complicates the application of the attack. The complexity of finding these 2^8 pairs is 2^{64} using the rebound attack, as described in Section 3.4.




Fig. 4: Truncated differential trail for P' used in the attack on 5 rounds.

4.1 Details of the Attack

Since we can construct only 2^8 pairs following the truncated differential trail, but need to cancel a 64-bit difference at the output of the compression function, each step of the attack succeeds only with a probability of 2^{-56} . However, this can be compensated for by using more message blocks in each step of the attack. Then, the attack can be summarized as follows:

1. Use the rebound attack (cf. Section 3.4) to find 2^8 pairs following the truncated differential trail. This has a complexity of 2^{64} in time and memory.
2. For each of these 2^8 pairs check if it can be used to cancel the corresponding 8 bytes of differences in state h'_i . This has a probability of 2^{-56} .

3. If no such *right* pair exists, then choose arbitrarily one of the 2^8 pairs and compute the state h'_i . This generates a new starting point with new differences for the next iteration, while keeping the same bytes inactive.
4. After 2^{56} new starting points, we expect to find a *right* pair with the desired output difference.

Since we need 2^{56} new starting points to cancel the differences in 8 bytes of the state, the complexity of the attack is equivalent to $8 \cdot 2^{64+56} = 2^{123}$ compression function evaluations. Note that the length of such a colliding message is about $8 \cdot 2^{56} = 2^{59}$ blocks.

4.2 Reducing the Length of the Colliding Message Pair

Beside the large time and memory complexity of the attack, one might also see the length of the colliding message pair as a limiting factor of the attack. However, the length of the colliding message pair can be significantly reduced by using a tree-based approach. Instead of choosing only one of the 2^8 pairs to generate a new starting point, we can continue with all pairs in parallel. By using a huge tree with 8 levels and 2^8 branches at each level, we get $(2^8)^8 = 2^{64}$ nodes at level 8. One of these 2^{64} nodes will have the desired output difference. This way, the length of the colliding message pair can be reduced to $8 \cdot 8 + 1 = 65$ message blocks.

4.3 Improving the Complexity of the Attack

The complexity of the attack can be slightly improved by using denser characteristics except when canceling the last 8 bytes. Instead of using a truncated differential trail with a $8 \rightarrow 1$ transition in round 3 of the trail, we can use truncated differential trails with $8 \rightarrow 8$, $8 \rightarrow 7, \dots, 8 \rightarrow 2$ which have a probability greater than 2^{-48} . The complexity of the attack is then dominated by the last iteration where we still need an $8 \rightarrow 1$ transition. This will improve the attack complexity by a factor of 8 resulting in a total complexity of 2^{120} compression function evaluations and 2^{64} memory.

5 Collisions in the Chosen-Prefix Setting

In a collision attack on a hash function an attacker has to find two arbitrary messages M and M^* such that $H(M) = H(M^*)$. However, in practice it might be required that the two messages contain some meaningful information, such that it can be used to practically compromise a cryptographic system. Such an example are, for instance, collisions in the chosen-prefix setting, where an attacker searches for a pair (M, M^*) such that

$$H(M_{pre} \| M) = H(M_{pre}^* \| M^*) \quad (3)$$

for a chosen-prefix (M_{pre}, M_{pre}^*) . In [24, 25], it was shown that such a more powerful attack exists for MD5. Moreover, the application of the attack to construct colliding X.509 certificates and the creation of a rogue certification authority certificate has been shown.

However, in most cases constructing such collisions is more complicated than constructing (random) collisions. In the case of MD5 the collision attack in the chosen-prefix setting has a complexity of 2^{49} , while the currently best collision attack on MD5 has a complexity of 2^{16} . However, in `Grøst1` the collision attack in the chosen-prefix setting has the same complexity as the collision attack. Due to the generic nature of the collision attack, differences in the chaining variables can be canceled efficiently (cf. Section 3).

6 Conclusion

In this work, we have provided new and improved cryptanalysis results for the `Grøst1` hash function, which significantly improving on previously known results. To be more precise, by using a new type of differential trail we were able to show collision attacks on `Grøst1` for up to 5 rounds. The extension becomes possible by considering differential trails spanning over more than one message block.

Moreover, due to the generic nature of our attack we can also construct meaningful collisions, i.e. collisions in the chosen-prefix setting with the same complexity. It has been shown in the past that such collisions might be exploited for instance to construct colliding X.509 certificates.

Although our results do not threaten the security of `Grøst1`, we believe that they will lead to a better understanding of the security margin of the hash function.

Acknowledgments. The work has been supported in part by the Austrian Government through the research program COMET (Project SeCoS, Project Number 836628) and through the research program FIT-IT Trust in IT Systems (Project SePAG, Project Number 835919), by the Secure Information Technology Center-Austria (A-SIT), and by the Research Fund KU Leuven, OT/13/071.

References

1. C. Boura, A. Canteaut, and C. De Cannière. Higher-Order Differential Properties of Keccak and Luffa. In A. Joux, editor, *FSE*, volume 6733 of *LNCS*, pages 252–269. Springer, 2011.
2. J. Daemen and V. Rijmen. The Wide Trail Design Strategy. In B. Honary, editor, *IMA Int. Conf.*, volume 2260 of *LNCS*, pages 222–238. Springer, 2001.
3. J. Daemen and V. Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer, 2002.
4. European Network of Excellence in Cryptology. ECRYPT II SHA-3 Zoo. http://ehash.iaik.tugraz.at/wiki/The_SHA-3_Zoo.

5. P. Gauravaram, L. R. Knudsen, K. Matusiewicz, F. Mendel, C. Rechberger, M. Schl affer, and S. S. Thomsen. Gr ostl – a SHA-3 candidate. Submission to NIST (Round 3), January 2011. Available online: <http://www.groestl.info>.
6. H. Gilbert and T. Peyrin. Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In S. Hong and T. Iwata, editors, *FSE*, volume 6147 of *LNCS*, pages 365–383. Springer, 2010.
7. K. Ideguchi, E. Tischhauser, and B. Preneel. Improved Collision Attacks on the Reduced-Round Gr ostl Hash Function. In M. Burmester, G. Tsudik, S. S. Magliveras, and I. Ilic, editors, *ISC*, volume 6531 of *LNCS*, pages 1–16. Springer, 2010.
8. J. Jean, M. Naya-Plasencia, and T. Peyrin. Improved Rebound Attack on the Finalist Gr ostl. In A. Canteaut, editor, *FSE*, volume 7549 of *LNCS*, pages 110–126. Springer, 2012.
9. J. Jean, M. Naya-Plasencia, and T. Peyrin. Multiple Limited-Birthday Distinguishers and Applications. In T. Lange, K. Lauter, and P. Lisonek, editors, *Selected Areas in Cryptography*, LNCS. Springer, 2013. (in press).
10. L. R. Knudsen. SMASH - A Cryptographic Hash Function. In H. Gilbert and H. Handschuh, editors, *FSE*, volume 3557 of *LNCS*, pages 228–242. Springer, 2005.
11. M. Lamberger, F. Mendel, C. Rechberger, V. Rijmen, and M. Schl affer. Rebound Distinguishers: Results on the Full Whirlpool Compression Function. In M. Matsui, editor, *ASIACRYPT*, volume 5912 of *LNCS*, pages 126–143. Springer, 2009.
12. M. Lamberger, F. Mendel, C. Rechberger, V. Rijmen, and M. Schl affer. The Rebound Attack and Subspace Distinguishers: Application to Whirlpool. Cryptology ePrint Archive, Report 2010/198, 2010. <http://eprint.iacr.org/>.
13. M. Lamberger, N. Pramstaller, C. Rechberger, and V. Rijmen. Second Preimages for SMASH. In M. Abe, editor, *CT-RSA*, volume 4377 of *LNCS*, pages 101–111. Springer, 2007.
14. F. Mendel, T. Peyrin, C. Rechberger, and M. Schl affer. Improved Cryptanalysis of the Reduced Gr ostl Compression Function, ECHO Permutation and AES Block Cipher. In M. J. Jacobson Jr., V. Rijmen, and R. Safavi-Naini, editors, *Selected Areas in Cryptography*, volume 5867 of *LNCS*, pages 16–35. Springer, 2009.
15. F. Mendel, C. Rechberger, M. Schl affer, and S. S. Thomsen. The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Gr ostl. In O. Dunkelman, editor, *FSE*, volume 5665 of *LNCS*, pages 260–276. Springer, 2009.
16. F. Mendel, C. Rechberger, M. Schl affer, and S. S. Thomsen. Rebound Attacks on the Reduced Gr ostl Hash Function. In J. Pieprzyk, editor, *CT-RSA*, volume 5985 of *LNCS*, pages 350–365. Springer, 2010.
17. National Institute of Standards and Technology. FIPS PUB 197: Advanced Encryption Standard. Federal Information Processing Standards Publication 197, U.S. Department of Commerce, November 2001. Available online: <http://www.itl.nist.gov/fipspubs>.
18. National Institute of Standards and Technology. Announcing Request for Candidate Algorithm Nominations for a New Cryptographic Hash Algorithm (SHA-3) Family. *Federal Register*, 27(212):62212–62220, November 2007. Available online: http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Nov07.pdf.
19. National Institute of Standards and Technology. SHA-3 Selection Announcement, October 2012. Available online: http://csrc.nist.gov/groups/ST/hash/sha-3/sha-3_selection_announcement.pdf.
20. T. Peyrin. Cryptanalysis of Grindahl. In K. Kurosawa, editor, *ASIACRYPT*, volume 4833 of *LNCS*, pages 551–567. Springer, 2007.
21. T. Peyrin. Improved Differential Attacks for ECHO and Gr ostl. In T. Rabin, editor, *CRYPTO*, volume 6223 of *LNCS*, pages 370–392. Springer, 2010.

22. Y. Sasaki, Y. Li, L. Wang, K. Sakiyama, and K. Ohta. Non-full-active Super-Sbox Analysis: Applications to ECHO and Grøstl. In M. Abe, editor, *ASIACRYPT*, volume 6477 of *LNCS*, pages 38–55. Springer, 2010.
23. M. Schl affer. Updated Differential Analysis of Grøstl, 2011. Available online: <http://www.groestl.info/>.
24. M. Stevens, A. K. Lenstra, and B. de Weger. Chosen-Prefix Collisions for MD5 and Colliding X.509 Certificates for Different Identities. In M. Naor, editor, *EUROCRYPT*, volume 4515 of *LNCS*, pages 1–22. Springer, 2007.
25. M. Stevens, A. Sotirov, J. Appelbaum, A. K. Lenstra, D. Molnar, D. A. Osvik, and B. de Weger. Short Chosen-Prefix Collisions for MD5 and the Creation of a Rogue CA Certificate. In S. Halevi, editor, *CRYPTO*, volume 5677 of *LNCS*, pages 55–69. Springer, 2009.
26. X. Wang, X. Lai, D. Feng, H. Chen, and X. Yu. Cryptanalysis of the Hash Functions MD4 and RIPEMD. In R. Cramer, editor, *EUROCRYPT*, volume 3494 of *LNCS*, pages 1–18. Springer, 2005.
27. X. Wang, Y. L. Yin, and H. Yu. Finding Collisions in the Full SHA-1. In V. Shoup, editor, *CRYPTO*, volume 3621 of *LNCS*, pages 17–36. Springer, 2005.
28. X. Wang and H. Yu. How to Break MD5 and Other Hash Functions. In R. Cramer, editor, *EUROCRYPT*, volume 3494 of *LNCS*, pages 19–35. Springer, 2005.
29. S. Wu, D. Feng, W. Wu, J. Guo, L. Dong, and J. Zou. (Pseudo) Preimage Attack on Round-Reduced Grøstl Hash Function and Others. In A. Canteaut, editor, *FSE*, volume 7549 of *LNCS*, pages 127–145. Springer, 2012.