Updated Differential Analysis of Grøstl

Martin Schläffer

Institute for Applied Information Processing and Communications (IAIK)
Graz University of Technology, Inffeldgasse 16a, A-8010 Graz, Austria

martin.schlaeffer@iaik.tugraz.at

Abstract. Grøstl is a SHA-3 finalist with clear proofs against a large class of differential attacks, similar to those of MD6. Furthermore, in this note we provide an update also regarding more advanced types of differential attacks that have been developed in recent years. We apply the rebound attacks on the initial submission to the tweaked version of Grøstl. We have analyzed the round-reduced hash function and compression function of Grøstl-256 (10 rounds) and Grøstl-512 (14 rounds). For both versions, we get collisions for 3 rounds of the hash function and collisions for 6 rounds of the compression function. We hope that our own efforts on improving the cryptanalysis will continue to motivate and accelerate external cryptanalysis.

1 Introduction

The SHA-3 candidate Grøstl has been tweaked in the final round of the NIST SHA-3 competition. In the following, we call the initial submission Grøstl-0 [1] and the tweaked version Grøstl [2]. In this note we apply the rebound attacks on Grøstl-0 published in [6] to Grøstl.

Since the shift values of ShiftBytes in Grøstl are different for P and Q, it is more difficult to find good truncated differential paths for both permutations which result in a (semi-free-start) collision. In general, most attacks are getting more difficult, since an optimal truncated differential path in P is usually not optimal in Q.

Furthermore, by increasing the size of the round constants in Grøstl, the internal differential attacks [3,7] (which consider differences between P and Q) get infeasible, even for a small number of rounds. In the following, we briefly present the best known attacks for the reduced and tweaked Grøstl compression functions and hash functions.

A summary of the currently best known results is given in Table 1.

Table 1. Summary of results for the round-reduced Grøstl hash and compression functions.

Target	Hash Size	Rounds	Time	Memory	Type	Reference
hash	224,256	3/10	2^{64}	-	collision	Sect. 2
function	512	3/14	2^{192}	-	collision	Sect. 3
compression	256	6/10	2^{120}	2^{64}	semi-free-start collision	Sect. 4
function	384,512	6/14	2^{180}	2^{64}	semi-free-start collision	Sect. 5

2 Collisions for 3 Rounds of Grøstl-256

When analyzing the hash function, we need to ensure that the pattern of active bytes prior to the last round is the same in each permutation. Furthermore, the different shift constants of SubBytes make the SuperBox match over the full first round (see [6]) more difficult. In the following we present two truncated differential paths which lead to a collision attack for 3 rounds of the hash function.

2.1 Path 1

The most simple case is to consider only one active byte prior to MixBytes in the last round of each permutation. Then we immediately get the minimum 3-round truncated differential path given in Fig. 1, with full active states at the input of each permutation.

Next, we need to verify if the truncated differential path is valid, i.e. if we have enough freedom such that the expected number of pairs is at least 1. This expected number of pairs can be computed by multiplying the total number of input pairs by the probability that the truncated differential path is followed for each input pair. Usually, we call the log_2 of the expected number of pairs the degrees of freedom we have in an attack.

For the truncated differential path of Fig. 1, the total number of input pairs depends on the number of pairs for the message M_i and for the chaining input H_{i-1} or initial value (IV), and we get approximately:

$$\underbrace{2^{8\cdot(64+64)}}_{M_1}\cdot\underbrace{1}_{IV}=2^{1024}$$

The probability of the given truncated differential path is determined by the probabilistic propagation in the MixBytes transformations of round r_1 and r_2 and in the final XOR at the output. For example, in the MixBytes transformation of round r_2 in permutation Q, the path reduces from 8 → 1 active bytes which happens with a probability of about 2^{-56}. In total, the approximate probability of the truncated differential path can be computed as follows:

$$\underbrace{2^{-8\cdot 56}\cdot 2^{-8\cdot 56}}_{\mathsf{MB(r_1)}}\cdot\underbrace{2^{-8\cdot 7}\cdot 2^{-8\cdot 7}}_{\mathsf{MB(r_2)}}\cdot\underbrace{2^{-8\cdot 1}}_{\mathsf{XOR}}=2^{-1016}$$

Hence, the expected number of pairs of the truncated differential path given in Fig. 1 can be computed as follows:

$$\underbrace{2^{8\cdot(64+64)}}_{M_1}\cdot\underbrace{1}_{IV}\cdot\underbrace{2^{-8\cdot 56}\cdot 2^{-8\cdot 56}}_{\mathsf{MB(r_1)}}\cdot\underbrace{2^{-8\cdot 7}\cdot 2^{-8\cdot 7}}_{\mathsf{MB(r_2)}}\cdot\underbrace{2^{-8\cdot 1}}_{\mathsf{XOR}}=2^{8}$$

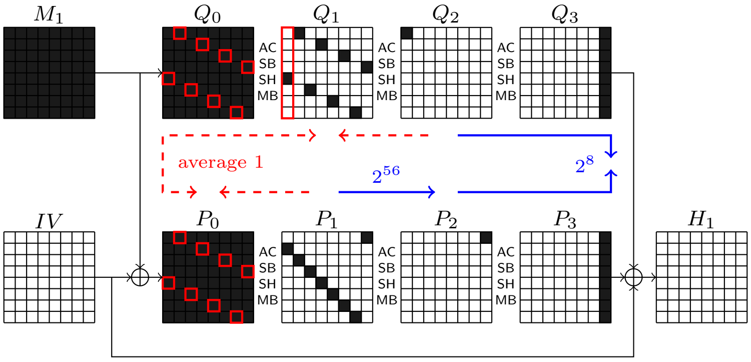

Fig. 1. The truncated differential path to get a collision attack on 3 out of 10 rounds for the hash function of Grøstl-256. The inbound phase (red) can be solved with average complexity 1, the outbound phase (blue) with a total complexity of about $2^{56}\cdot 2^{8}=2^{64}$. The first SuperBox in the inbound phase is shown by red rectangles.

We use the rebound attack to find pairs for the truncated differential paths in P and Q. First, we compute pairs for the inbound phase between rounds r_1 and r_2 in Q and round r_1 in P. Note that in this path, the SuperBoxes are not fully active. In this case, the memory complexity of the attack can be reduced significantly. In fact we can even apply the techniques published in [5] with memory complexities of at most 2^{16}. Also the techniques published in [8] and [4] could be applied. In any case, the complexity to find a conforming input pair according to the truncated differential path until state Q_2 in permutation Q, and until state P_1 in permutation P is 1 on average. We compute 2^{64} such pairs and propagate them outwards. With a probability of 2^{-56} we get one active byte in P_2 and with a probability of 2^{-8} also the 1-byte differences in the last round prior to MixBytes are equal. Hence, we get a collision for 3 rounds of the hash function with a total complexity of 2^{64} in time and negligible memory requirements.

2.2 Path 2

Again we can use a second truncated differential path which has the same time complexity, but higher memory complexities. We still mention this path here since it could be interesting in future analysis of the Grøstl-256 hash function. The path is constructed in a similar way as the second path of the compression function attacks on Grøstl-256 and given in Fig. 2. Note that the pattern of active bytes in Q_2 can be determined from the pattern in P_2 by the relation

$$Q_2 \leftarrow \mathsf{ShiftBytes}_{\mathsf{Q}}^{-1} \circ \mathsf{ShiftBytes}_{\mathsf{P}} \circ \mathsf{P}_2$$

which results in the following left-shift values (mod 8):

$$\{0,1,2,3,4,5,6,7\} - \{1,3,5,7,0,2,4,6\} = \{7,6,5,4,4,3,2,1\}$$

Again, we verify if the truncated differential path is valid and compute the expected number of solutions. The path is probabilistic in the MixBytes transformations of round r_1 and r_2 in Q, in the MixBytes transformations of round r_1 in P, and in the XOR at the output. Hence, the expected number of pairs is given as follows:

$$\underbrace{2^{8\cdot(64+64)}}_{M_1} \cdot \underbrace{1}_{IV} \cdot \underbrace{2^{-8\cdot 8} \cdot 2^{-8\cdot 56}}_{\mathsf{MB(r_1)}} \cdot \underbrace{2^{-8\cdot 49}}_{\mathsf{MB(r_2)}} \cdot \underbrace{2^{-8\cdot 8}}_{\mathsf{XOR}} = 2^{56}$$

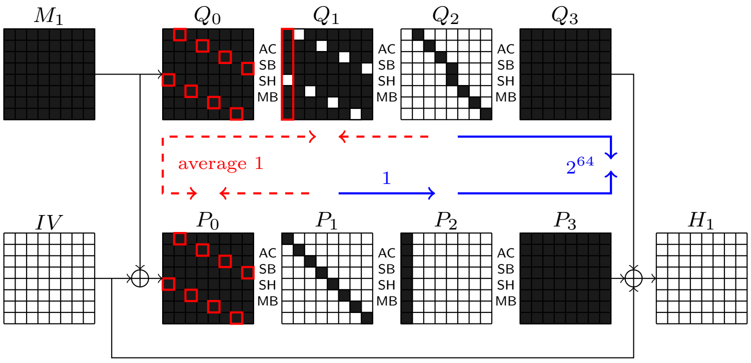

Fig. 2. Another truncated differential path to get a collision attack on 3 out of 10 rounds for the hash function of Grøstl-256. The inbound phase (red) can be solved with average complexity 1, the outbound phase (blue) with a total complexity of about $1 \cdot 2^{64} = 2^{64}$. The first SuperBox in the inbound phase is shown by red rectangles.

In the inbound phase, we can do a standard SuperBox match with memory complexity 2^{64}, or use the non-full active SuperBox techniques of [8] with a memory complexity of 2^{56} since only 7 bytes are active. We get one pair on average for the inbound phase, such that the truncated differential path until states Q_2 and P_1 is fulfilled. Furthermore, each of these pairs also follows the truncated differential path until the end of each permutation with a probability of almost 1. We get a collision if the 8-byte differences prior to the last MixBytes transformation are equal which happens with a probability of 2^{-64}. Hence, the total complexity to get a collision for 3 rounds of the Grøstl-256 hash function using this path is 2^{64} with memory requirements of 2^{56}.

3 Collisions for 3 Rounds of Grøstl-512

When analyzing the hash function of Grøstl-512, we first need to construct a colliding and valid truncated differential path. Similar as for the compression function, we need an (almost) full active state at least once in the path due to the wide-trail design strategy.

The used truncated differential path is shown in Fig. 3. Due to the slower diffusion in Grøstl-512, we cannot use a single active byte in the last round. However, we can use a path with 3 active bytes in both P_2 and Q_2. This path is similar as Path 1 in the hash function attack on Grøstl-256. Note that we need at least 3 active bytes in the last round such that we get full active states in both Q_0 and P_0. Otherwise, the pattern of active bytes would not match at the input of the hash function. Note that the pattern of active bytes in Q_2 can be determined from the pattern in P_2 by the relation

$$Q_2 \leftarrow \mathsf{ShiftBytes}_\mathsf{Q}^{-1} \circ \mathsf{ShiftBytes}_\mathsf{P} \circ \mathsf{P}_2$$

which results in the following left-shift values (mod 16):

$$\{0, 1, 2, 3, 4, 5, 6, 11\} - \{1, 3, 5, 11, 0, 2, 4, 6\} = \{15, 14, 13, 8, 4, 3, 2, 5\}.$$

Finally, we also verify if this truncated differential path is valid and compute the expected number of solutions. The path is probabilistic in the MixBytes transformations of round r_1 and r_2 in both P and Q, and in the XOR at the output. Hence, the expected number of pairs is given as follows:

$$\underbrace{2^{8\cdot(128+128)}}_{M_1} \cdot \underbrace{1}_{IV} \cdot \underbrace{2^{-8\cdot 104} \cdot 2^{-8\cdot 104}}_{\mathsf{MB(r_1)}} \cdot \underbrace{2^{-8\cdot 21} \cdot 2^{-8\cdot 21}}_{\mathsf{MB(r_2)}} \cdot \underbrace{2^{-8\cdot 3}}_{\mathsf{XOR}} = 2^{24}$$

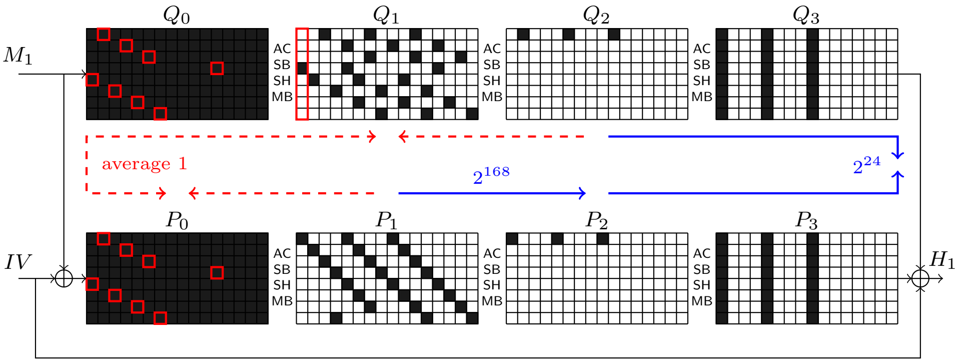

Fig. 3. The truncated differential path to get a collision attack on 3 out of 14 rounds for the hash function of Grøstl-512. The inbound phase (red) can be solved with average complexity 1, the outbound phase (blue) with a total complexity of about $2^{3\cdot 56} \cdot 2^{3\cdot 8} = 2^{192}$. The first SuperBox in the inbound phase is shown by red rectangles.

Again we use the rebound attack to find pairs according to this truncated differential path. First, we compute pairs for the inbound phase between rounds r_1 and r_2 in Q and round r_1 in P. The complexity to find a solution for the truncated differential path until state Q_2 in permutation Q, and until state P_1 in permutation P is 1 on average with memory requirements of 2^{64} for a standard SuperBox match. Using non-full active SuperBox matches or by solving linearly for pairs, we can significantly reduce the memory requirements to at most 2^{16}. We compute 2^{192} such pairs in the inbound phase and propagate them outwards. With a probability of 2^{-168} we get 3 active bytes after the MixBytes transformation in round r_2 of permutation P. The 3-byte differences in the last round prior to MixBytes are equal with a probability of 2^{-24}. Hence, we get a collision for 3 rounds of the hash function with a total complexity of 2^{192} in time with negligible memory requirements.

4 Semi-Free-Start Collisions for 6 Rounds of Grøstl-256

To get a semi-free-start collision for the compression function of Grøstl we first need to construct a colliding truncated differential path for the two permutations. Due to the different shift value in P and Q this gets more difficult than for Grøstl-0. In the following, we show two truncated differential paths which lead to collision attacks on the compression function with the same attack complexity.

4.1 Path 1

The most straight-forward approach is to consider only one active byte prior to the first and after the last SubBytes layer. This way, the ShiftBytes transformations do not change the pattern of active bytes in the first and last round and we can get a collision at the output of the compression function. The number of active bytes for each round in both P and Q is then given as follows:

$$1 \xrightarrow{r_1} 8 \xrightarrow{r_2} 64 \xrightarrow{r_3} 64 \xrightarrow{r_4} 8 \xrightarrow{r_5} 1 \xrightarrow{r_6} 8$$

The truncated differential path is shown in Fig. 4. Note that the path and also the pattern of active bytes is still similar in P and Q.

Next, we verify if the truncated differential path is valid, i.e. if the expected number of solutions is at least 1. This number can be computed by multiplying the total number of input pairs by the probability that the truncated differential path is followed for each input pair. Note that the path is only probabilistic in the MixBytes transformations of round r_4 and r_5, and in the XOR at the output. Hence, the expected number of pairs is given as follows:

$$\underbrace{2^{8\cdot(64+1)}}_{M_i} \cdot \underbrace{2^{8\cdot 64}}_{H_{i-1}} \cdot \underbrace{2^{-8\cdot 56} \cdot 2^{-8\cdot 56}}_{\mathsf{MB}(r_4)} \cdot \underbrace{2^{-8\cdot 7} \cdot 2^{-8\cdot 7}}_{\mathsf{MB}(r_5)} \cdot \underbrace{2^{-8\cdot 1}}_{\mathsf{XOR}} = 2^{16}$$

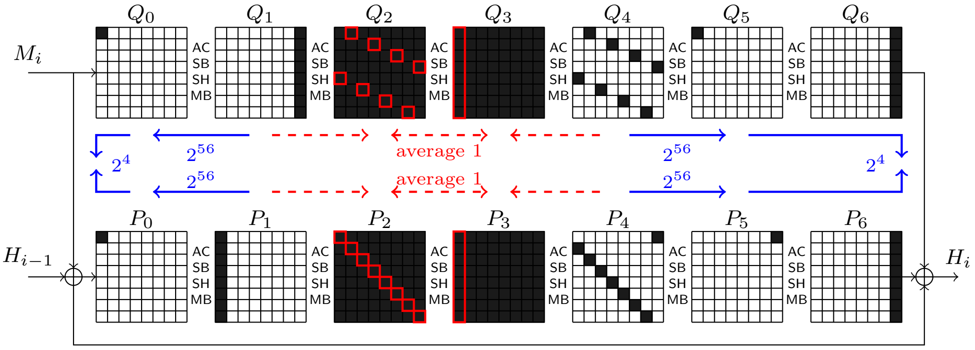

Fig. 4. The truncated differential path to get a semi-free-start collision attack for 6 out of 10 rounds of the compression function of Grøstl-256. The inbound phase (red) can be solved with average complexity 1, and the outbound phase (blue) with a total complexity of about $2^4 \cdot 2^{56} \cdot 2^{56} \cdot 2^4 = 2^{120}$ compression function evaluations. The first SuperBox in the inbound phase is shown by red rectangles.

In the compression function attack, we first compute pairs for each permutation independently and then, match the input and output differences using a birthday attack. By computing the inbound phase with SuperBox matches, we can find pairs for the three middle rounds r_2, r_3 and r_4 with an average complexity of 1 and memory requirements of 2^{64}. For each permutation, we independently propagate the resulting pairs outwards and get one active byte at the input (P_0, Q_0) and one active byte after round r_5 (P_5, Q_5) with a complexity of $2^{2\cdot 56} = 2^{112}$. To get a semi-free-start collision, the 1-byte differences at the input, and the 1-byte differences prior to the last MixBytes transformation need to be equal. This 16-bit condition can be fulfilled with a complexity of 2^8 using the birthday effect. In total, the complexity to get a semi-free-start collision for 6 rounds of Grøstl-256 is $2^{112} \cdot 2^8 = 2^{120}$ in time with memory requirements of 2^{64}.

4.2 Path 2

Note that also another path with more active bytes can be used to get a semi-free-start collisions for 6 rounds and with the same complexity. This time, we use two truncated differential paths in P and Q where the full active state does not occur in the same round. Hence, the number of active bytes in P and Q are different and given as follows:

$$Q : 8 \xrightarrow{r_1} 1 \xrightarrow{r_2} 8 \xrightarrow{r_3} 64 \xrightarrow{r_4} 56 \xrightarrow{r_5} 8 \xrightarrow{r_6} 64$$
$$P : 8 \xrightarrow{r_1} 56 \xrightarrow{r_2} 64 \xrightarrow{r_3} 8 \xrightarrow{r_4} 1 \xrightarrow{r_5} 8 \xrightarrow{r_6} 64$$

The respective truncated differential path is shown in Fig. 5. Remember that we need the same pattern of active bytes at the input and prior to the last MixBytes transformation to get a semi-free-start collision. For the given truncated differential path, we have 8 active bytes in these states. Due to the different shift values in P and Q, we can get a single active byte only in one of P or Q. For the other permutation, we immediately get an almost full active state. In the given path, we get an (almost) full active state in permutation P since we require a single active byte in Q after the first round in states P_1 and Q_1. We get the opposite behavior in backward direction in states P_4 (single active byte) and Q_4 (almost full active state). Nevertheless, this truncated differential path can be used to efficiently find collisions for the compression function of Grøstl-256.

Again, we verify if the truncated differential path is valid and compute the expected number of solutions. This time, the path is probabilistic in the MixBytes transformations of round r_1, r_4 and r_5 of Q, in the MixBytes transformations of round r_3 and r_4 of P, and in the XOR at the output. Hence, the expected number of pairs is given as follows:

$$\underbrace{2^{8\cdot(64+8)}}_{M_i} \cdot \underbrace{2^{8\cdot 64}}_{H_{i-1}} \cdot \underbrace{2^{-8\cdot 7}}_{\mathsf{MB(r_1)}} \cdot \underbrace{2^{-8\cdot 56}}_{\mathsf{MB(r_3)}} \cdot \underbrace{2^{-8\cdot 8} \cdot 2^{-8\cdot 7}}_{\mathsf{MB(r_4)}} \cdot \underbrace{2^{-8\cdot 49}}_{\mathsf{MB(r_5)}} \cdot \underbrace{2^{-8\cdot 8}}_{\mathsf{XOR}} = 2^8$$

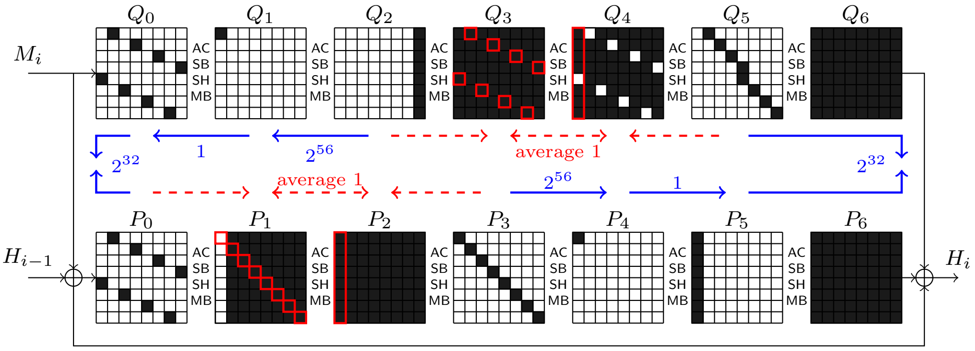

Fig. 5. Another truncated differential path to get a semi-free-start collision attack for 6 out of 10 rounds of the compression function of Grøstl-256. The inbound phase (red) can be solved with average complexity 1, and the outbound phase (blue) with a total complexity of about $2^{32} \cdot 2^{56} \cdot 2^{32} = 2^{120}$ compression function evaluations. The first SuperBox in the inbound phase is shown by red rectangles.

When applying the rebound attack to this path, we can solve the inbound phase for rounds r_1, r_2 and r_3 in permutation P, and for rounds r_3, r_4 and r_5 in permutation Q independently and with average complexity 1. The memory requirements are 2^{64} again. In each permutation, we have one propagation through MixBytes from 8 to 1 active byte which has a complexity of 2^{56} in each case. This time, we get a 128-bit condition such that the differences of the 8 active bytes at the input and output (prior to MixBytes) cancel each other. Using a birthday attack we can match the differences with a complexity of 2^{64} in time and memory. In total, the complexity for this semi-free-start collision attack on 6 rounds is again $2^{56} \cdot 2^{64} = 2^{120}$ in time with memory requirements of 2^{64}.

5 Semi-Free-Start Collisions for 6 Rounds of Grøstl-512

To get a semi-free-start collision for the compression function of Grøstl-512, we use a similar path as in [6]. Due to the different shift values in P and Q we need to reduce this path by one round to get a colliding truncated difference pattern at the input and output. This gets much easier if the number of active bytes at the input and output is very low. The truncated differential path is shown in Fig. 6.

Next, we verify if the truncated differential path is valid and compute the expected number of solutions. The path is probabilistic in the MixBytes transformations of round r_3, r_4 and r_5, and in the XOR at the output. The expected number of pairs is given as follows:

$$\underbrace{2^{8\cdot(128+1)}}_{M_i}\cdot\underbrace{2^{8\cdot(128)}}_{H_{i-1}}\cdot\underbrace{2^{-8\cdot 16}\cdot 2^{-8\cdot 16}}_{\mathsf{MB(r_3)}}\cdot\underbrace{2^{-8\cdot 96}\cdot 2^{-8\cdot 96}}_{\mathsf{MB(r_4)}}\cdot\underbrace{2^{-8\cdot 14}\cdot 2^{-8\cdot 14}}_{\mathsf{MB(r_5)}}\cdot\underbrace{2^{-8\cdot 2}}_{\mathsf{XOR}}=2^{24}$$

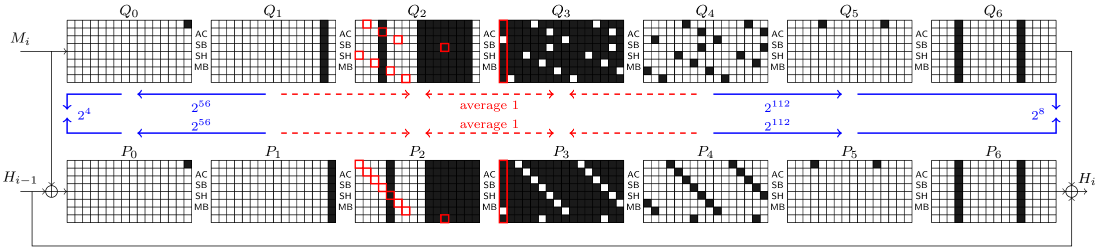

Fig. 6. The truncated differential path to get a semi-free-start collision attack for 6 out of 14 rounds of the compression function of Grøstl-512. The inbound phase (red) can be solved with average complexity 1, and the outbound phase (blue) with a total complexity of about $2^4 \cdot 2^{56} \cdot 2^{112} \cdot 2^8 = 2^{180}$ compression function evaluations. The first SuperBox in the inbound phase is shown by red rectangles.

Again, we use the rebound attack to find pairs for each truncated differential path in P and Q. We compute pairs for each permutation independently and match the input and output differences using a birthday attack. By computing the inbound phase with SuperBox matches, we can find pairs for the three middle rounds r_2, r_3 and r_4 with an average complexity of 1 and memory requirements of 2^{64}. For each permutation, we independently propagate the resulting pairs outwards and get one active byte at the input (P_0, Q_0) and one active byte after round r_5 (P_5, Q_5) with a complexity of $2^{3\cdot 56} = 2^{168}$. To get a semi-free-start collision, the 1-byte differences at the input, and the 2-byte differences prior to the last MixBytes transformation need to be equal. This 24-bit condition can be fulfilled with a complexity of 2^{12} using the birthday effect. In total, the complexity to get a semi-free-start collision for 6 rounds of Grøstl-512 is $2^{168} \cdot 2^{12} = 2^{180}$ in time with memory requirements of 2^{64}.

6 Conclusion

In this note we have updated the cryptanalysis results on Grøstl to the tweaked version. The given results and truncated differential paths provide a starting point for future independent analysis of Grøstl. Furthermore, we encourage the analysis of Grøstl-0, the initial submission to the SHA-3 competition.

References

1. Praveen Gauravaram, Lars R. Knudsen, Krystian Matusiewicz, Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. Grøstl – a SHA-3 candidate. Submission to NIST, 2008. Available online: http://groestl.info.

2. Praveen Gauravaram, Lars R. Knudsen, Krystian Matusiewicz, Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. Grøstl – a SHA-3 candidate. Submission to NIST (Round 3), 2011. Available online: `http://groestl.info`.
3. Kota Ideguchi, Elmar Tischhauser, and Bart Preneel. Improved collision attacks on the reduced-round Grøstl hash function. In *Information Security Conference*, 2011. To appear.
4. Jérémy Jean and Pierre-Alain Fouque. Practical Near-Collisions and Collisions on Round-Reduced ECHO-256 Compression Function. In *Fast Software Encryption*, 2011. To appear.
5. Florian Mendel, Thomas Peyrin, Christian Rechberger, and Martin Schläffer. Improved Cryptanalysis of the Reduced Grøstl Compression Function, ECHO Permutation and AES Block Cipher. In Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini, editors, *Selected Areas in Cryptography*, volume 5867 of *LNCS*, pages 16–35. Springer, 2009.
6. Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. Rebound Attacks on the Reduced Grøstl Hash Function. In Josef Pieprzyk, editor, *CT-RSA*, volume 5985 of *LNCS*, pages 350–365. Springer, 2010.
7. Thomas Peyrin. Improved Differential Attacks for ECHO and Grøstl. In Tal Rabin, editor, *CRYPTO*, volume 6223 of *LNCS*, pages 370–392. Springer, 2010.
8. Yu Sasaki, Yang Li, Lei Wang, Kazuo Sakiyama, and Kazuo Ohta. Non-full-active Super-Sbox Analysis: Applications to ECHO and Grøstl. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *LNCS*, pages 38–55. Springer, 2010.