Grøstl – a SHA-3 candidate[*]

Praveen Gauravaram[1], Lars R. Knudsen[1], Krystian Matusiewicz[1], Florian Mendel[2], Christian Rechberger[2], Martin Schläffer[2], and Søren S. Thomsen[1]

[1]Department of Mathematics, Technical University of Denmark, Matematiktorvet 303S, DK-2800 Kgs. Lyngby, Denmark
[2]Institute for Applied Information Processing and Communications (IAIK), Graz University of Technology, Inffeldgasse 16a, A-8010 Graz, Austria

January 15, 2009

Summary

Grøstl is a SHA-3 candidate proposal. Grøstl is an iterated hash function with a compression function built from two fixed, large, distinct permutations. The design of Grøstl is transparent and based on principles very different from those used in the SHA-family.

The two permutations are constructed using the wide trail design strategy, which makes it possible to give strong statements about the resistance of Grøstl against large classes of cryptanalytic attacks. Moreover, if these permutations are assumed to be ideal, there is a proof for the security of the hash function.

Grøstl is a byte-oriented SP-network which borrows components from the AES. The S-box used is identical to the one used in the block cipher AES and the diffusion layers are constructed in a similar manner to those of the AES. As a consequence there is a very strong confusion and diffusion in Grøstl.

Grøstl is a so-called wide-pipe construction where the size of the internal state is significantly larger than the size of the output. This has the effect that all known, generic attacks on the hash function are made much more difficult.

Grøstl has good performance on a wide range of platforms, and counter-measures against side-channel attacks are well-understood from similar work on the AES.

---

[*]Document version no. 1.1 (updated January 15, 2009). Most important updates in Section 7.2.2. No changes in the design of Grøstl have been made.

Contents

1 Introduction

In this proposal we present the cryptographic hash function Grøstl as a candidate for the SHA-3 competition initiated by the National Institute of Standards and Technology (NIST).

The paper is organised as follows. In Section 2, we give a high-level summary of the Grøstl proposal, and state the design goals. In Section 3, we present the details of the proposal and in Section 4, we describe the features specific to Grøstl and motivate our design choices. Section 5 introduces some alternative modes of operation of Grøstl for the use as message authentication codes. In Section 6, we present our preliminary cryptanalysis results on Grøstl. Section 7 deals with implementation aspects of Grøstl, including benchmarks results and performance estimates. Finally, we conclude in Section 8.

The name "Grøstl" may cause some problems in terms of pronunciation, and also due to the character 'ø', which has different encodings around the world. Whenever problems with character encodings may arise, we recommend the spelling Groestl. With respect to pronunciation and other information on the name, see Appendix A.

2 Design goals

In this section, we give a brief motivation of the Grøstl proposal. Elegance of the design and simplicity of analysis, as well as proofs of desirable properties are the overall goals. The fact that it iteratively applies a compression function is among the few similarities with commonly used hash functions. Additionally, we aim to have security margins at several layers of abstraction in the design.

2.1 Overall goals for the hash

Here we state overall design goals for Grøstl.

- Simplicity of analysis, hence, Grøstl is based on a small number of permutations instead of a block cipher (with many permutations).

- Provably secure construction (assuming ideal permutations).

- Well-known design principles underlying the permutations (again, allowing simple analysis, provable properties).

- No special preference for a particular platform or word size, and good performance on a very wide range of platforms.

- Side-channel resistance at little additional cost.

- Defining reduced variants for cryptanalysis is made straightforward.

- Prevention of length-extension attacks.

- Allow implementers to exploit parallelisation within the compression function[1].

---

[1]Using Grøstl in a tree-mode, as any other cryptographic hash function for that matter, will also allow to exploit parallelisation at a higher level, but we consider this outside the scope of our submission.

2.2 Failure-tolerant design

Non-random behaviour of the employed permutations do not necessarily lead to non-ideal properties of the compression function. Attacks on the compression function, in turn, may not lead to attacks on the hash function.

- The internal state is significantly larger than the final output – hence, all known generic attacks are thwarted.

- Known techniques that exhibit non-ideal behaviour of the permutations work only for reduced variants.

- Attacks on the compression function do not necessarily translate to attacks on the hash function.

- There are no known attacks on the compression function meeting the proven lower bounds.

2.3 Design considerations for the compression function

Traditional design approaches of hash functions are based on block ciphers, e.g., MD5, SHA-1, SHA-256 [42], Whirlpool [3], Tiger [1], etc. This may seem sound since block cipher designs are well understood. However, the key schedule of the block cipher becomes more important in a setting where the attacker has control over every input and there is little consensus in the community what constitutes a good key schedule. The recent attacks [16, 28, 38, 39, 54] on SHA-1 and Tiger illustrate this issue. For this reason we base our proposal on a few individual permutations rather than a large family of permutations indexed by a key. The advantages of such a design methodology is as follows:

- No threat of attacks via the key schedule (e.g., weak keys).

- Since the key schedule of a block cipher is often rather slow, performance may be improved.

- Simplicity.

3 Specification of Grøstl

Grøstl is a collection of hash functions, capable of returning message digests of any number of bytes from 1 to 64, i.e., from 8 to 512 bits in 8-bit steps. The variant returning n bits is called Grøstl-n. We explicitly state here that this includes the message digest sizes 224, 256, 384, and 512 bits. We now specify the Grøstl hash functions.

3.1 The hash function construction

The Grøstl hash functions iterate the compression function f as follows. The message M is padded and split into ℓ-bit message blocks $m_1, \ldots, m_t$, and each message block is processed sequentially. An initial ℓ-bit value $h_0 = \mathrm{iv}$ is defined, and subsequently the message blocks m_i are processed as

$$h_i \leftarrow f(h_{i-1}, m_i) \quad \text{for } i = 1, \ldots, t.$$

Hence, f maps two inputs of ℓ bits each to an output of ℓ bits. The first input is called the *chaining input*, and the second input is called the *message block*. For Grøstl variants returning up to 256 bits, ℓ is defined to be 512. For larger variants, ℓ is 1024.

After the last message block has been processed, the output $H(M)$ of the hash function is computed as

$$H(M) = \Omega(h_t),$$

where Ω is an *output transformation* which is defined in Section 3.3. The output size of Ω is n bits, and we note that $n < \ell$. See Figure 1.

Figure 1: The Grøstl hash function.

3.2 The compression function construction

The compression function f is based on two underlying ℓ-bit permutations P and Q. It is defined as follows:

$$f(h, m) = P(h \oplus m) \oplus Q(m) \oplus h. \tag{1}$$

The construction of f is illustrated in Figure 2. In Section 3.4, we describe how P and Q are

Figure 2: The compression function f. P and Q are ℓ-bit permutations.

defined.

3.3 The output transformation

Let $\mathrm{trunc}_n(x)$ be the operation that discards all but the trailing n bits of x. Then the output transformation Ω is defined as

$$\Omega(x) = \mathrm{trunc}_n(P(x) \oplus x).$$

See Figure 3.

3.4 The design of P and Q

As mentioned, the compression function f comes in two variants; one is used for short message digests, and one is used for long message digests. Each variant uses its own pair of permutations P and Q. Hence, we define four permutations in total. The permutations will be assigned with subscripts 512 or 1024, whenever it is necessary to distinguish them.

5

Figure 3: The output transformation Ω computes $P(x) \oplus x$ and then truncates the output by returning only the last n bits.

The design of P and Q was inspired by the Rijndael block cipher algorithm [12, 13]. This means that their design consist of a number of *rounds R*, which consists of a number of *round transformations*. Since P and Q are much larger than the 128-bit state size of Rijndael, most round transformations have been redefined. In Grøstl, a total of four round transformations are defined for each permutation. These are

- AddRoundConstant

- SubBytes

- ShiftBytes

- MixBytes.

When a distinction is necessary, the third transformation ShiftBytes will be called ShiftBytesWide when used in the large permutations P_{1024} and Q_{1024}. All other transformations can be described in the same way for all four permutations.

A round R consists of these four round transformations applied in the above order. Hence,

$$R = \text{MixBytes} \circ \text{ShiftBytes} \circ \text{SubBytes} \circ \text{AddRoundConstant}.$$

See Figure 4. We note that all rounds follow this definition. We denote by r the number of rounds. Concrete recommendations for r will be given in Section 3.4.6.

The transformations operate on a state, which is represented as a matrix A of bytes (of 8 bits each). For the short variants, the matrix has 8 rows and 8 columns, and for the large variants, the matrix has 8 rows and 16 columns. In the following, we denote by v the number of columns, and we write constant byte values in sans serif font, e.g., c3. In the following, we describe how to map a byte sequence to a state matrix and back, and then we describe each round transformation.

3.4.1 Mapping from a byte sequence to a state matrix and vice versa

Since Grøstl operates on bytes, it is generally endianness neutral. However, we need to specify how a byte sequence is mapped to the matrix A, and vice versa. This mapping is done in a similar way as in Rijndael. Hence, the 64-byte sequence 00 01 02 ... 3f is mapped to an 8×8
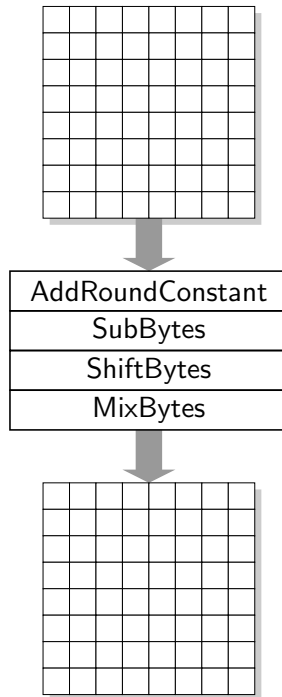
Figure 4: One round of the Grøstl permutations P and Q is a composition of four basic transformations.

matrix as

$$
\begin{bmatrix}
00 & 08 & 10 & 18 & 20 & 28 & 30 & 38 \\
01 & 09 & 11 & 19 & 21 & 29 & 31 & 39 \\
02 & 0a & 12 & 1a & 22 & 2a & 32 & 3a \\
03 & 0b & 13 & 1b & 23 & 2b & 33 & 3b \\
04 & 0c & 14 & 1c & 24 & 2c & 34 & 3c \\
05 & 0d & 15 & 1d & 25 & 2d & 35 & 3d \\
06 & 0e & 16 & 1e & 26 & 2e & 36 & 3e \\
07 & 0f & 17 & 1f & 27 & 2f & 37 & 3f
\end{bmatrix} .
$$

For an 8×16 matrix, this method is extended in the natural way. Mapping from a matrix to a byte sequence is simply the reverse operation. From now on, we do not explicitly mention this mapping.

3.4.2 AddRoundConstant

The AddRoundConstant transformation adds a round-dependent constant to the state matrix A. By addition we mean exclusive-or (XOR). P and Q have different round constants, which is the only difference between the two permutations.

The round constants can be seen as matrices of the same size as the state matrix. All round constants bytes are zero except for a single position. The round constants used for P_{512} and P_{1024} are basically the same: The first 8 columns do not differ for both permutations and the last 8 columns of the round constants used in P_{1024} contain only bytes having the value 00. Likewise for Q_{512} and Q_{1024}.

The byte in the top leftmost corner of the round constant in round i of P has the value i; all other positions in the round constant matrix have the value 00. In Q, the byte in the bottom

leftmost corner has the value $i \oplus \mathsf{ff}$, and all other bytes have the value $\mathsf{00}$. The round number is reduced modulo 256, if necessary.

To be precise, the AddRoundConstant transformation in round i updates the state A as

$$A \leftarrow A \oplus C[i],$$

where $C[i]$ is the round constant used in round i. The round constants $C_P[i]$ and $C_Q[i]$ used in round i of P and Q, respectively, are

$$
C_P[i] = \begin{bmatrix}
i & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00}
\end{bmatrix}
\quad \text{and} \quad
C_Q[i] = \begin{bmatrix}
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
\mathsf{00} & \mathsf{00} & \cdots & \mathsf{00} \\
i \oplus \mathsf{ff} & \mathsf{00} & \cdots & \mathsf{00}
\end{bmatrix}.
$$

See Figure 5.

AddRoundConstant of P AddRoundConstant of Q

Figure 5: AddRoundConstant for permutations P and Q modify a single byte of the state by adding a constant derived from the round number i.

3.4.3 SubBytes

The SubBytes transformation substitutes each byte in the state matrix by another value, taken from the s-box S. This s-box is the same as the one used in Rijndael and its specification can be found in Appendix B. Hence, if $a_{i,j}$ is the element in row i and column j of A, then SubBytes performs the following transformation:

$$a_{i,j} \leftarrow S(a_{i,j}), \quad 0 \leq i < 8, \ 0 \leq j < v.$$

See Figure 6.

3.4.4 ShiftBytes and ShiftBytesWide

ShiftBytes and ShiftBytesWide cyclically shift the bytes within a row to the left by a number of positions. Let $\sigma = [\sigma_0, \sigma_1, \ldots, \sigma_7]$ be a list of distinct integers in the range from 0 to $v - 1$.

Figure 6: SubBytes substitutes each byte of the state by its image under the s-box S.

Then, ShiftBytes moves all bytes in row i of the state matrix σ_i positions to the left, wrapping around as necessary. The vector σ is defined as $\sigma = [0, 1, 2, 3, 4, 5, 6, 7]$ in ShiftBytes, and $\sigma = [0, 1, 2, 3, 4, 5, 6, 11]$ in ShiftBytesWide. See Figure 7.

(a) ShiftBytes
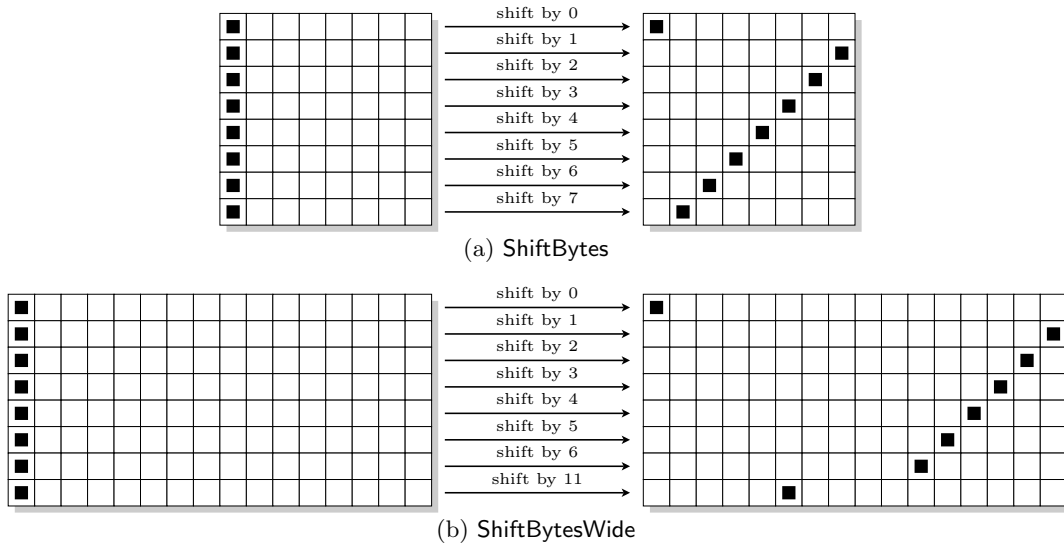
(b) ShiftBytesWide

Figure 7: The ShiftBytes and ShiftBytesWide transformations.

3.4.5   MixBytes

In the MixBytes transformation, each column in the matrix is transformed independently. To describe this transformation we first need to introduce the finite field $\mathbb{F}_{256}$. This finite field is defined in the same way as in Rijndael via the irreducible polynomial $x^8 \oplus x^4 \oplus x^3 \oplus x \oplus 1$ over $\mathbb{F}_2$. The bytes of the state matrix A can be seen as elements of $\mathbb{F}_{256}$, i.e., as polynomials of degree at most 7 with coefficients in $\{0, 1\}$. The least significant bit of each byte determines the coefficient of x^0, etc.

MixBytes multiplies each column of A by a constant 8×8 matrix B in $\mathbb{F}_{256}$. Hence, the transformation on the whole matrix A can be written as the matrix multiplication

$$A \leftarrow B \times A.$$

The matrix B is specified as

$$B = \begin{bmatrix} 02 & 02 & 03 & 04 & 05 & 03 & 05 & 07 \\ 07 & 02 & 02 & 03 & 04 & 05 & 03 & 05 \\ 05 & 07 & 02 & 02 & 03 & 04 & 05 & 03 \\ 03 & 05 & 07 & 02 & 02 & 03 & 04 & 05 \\ 05 & 03 & 05 & 07 & 02 & 02 & 03 & 04 \\ 04 & 05 & 03 & 05 & 07 & 02 & 02 & 03 \\ 03 & 04 & 05 & 03 & 05 & 07 & 02 & 02 \\ 02 & 03 & 04 & 05 & 03 & 05 & 07 & 02 \end{bmatrix}.$$

This matrix is *circulant*, which means that each row is equal to the row above rotated right by one position. In short, we may write $B = \mathrm{circ}(02, 02, 03, 04, 05, 03, 05, 07)$ instead. See also Figure 8.

Figure 8: The MixBytes transformation left-multiplies each column of the state matrix treated as a column vector over $\mathbb{F}_{256}$ by a circulant matrix B.

3.4.6   Number of rounds

The number r of rounds is a tunable security parameter. We recommend the following values of r for the four permutations.

Permutations	Digest sizes	Recommended value of r
P_{512} and Q_{512}	8–256	10
P_{1024} and Q_{1024}	264–512	14

3.5   Initial values

The initial value iv_n of Grøstl-n is the ℓ-bit representation of n. The table below shows the initial values of the required output sizes of 224, 256, 384, and 512 bits.

n	iv_n
224	00 ... 00 00 e0
256	00 ... 00 01 00
384	00 ... 00 01 80
512	00 ... 00 02 00

3.6 Padding

As mentioned, the length of each message block is ℓ. To be able to operate on inputs of varying length, a padding function pad is defined. This padding function takes a string x of length N bits and returns a padded string $x^* = \mathsf{pad}(x)$ of a length which is a multiple of ℓ.

The padding function does the following. First, it appends the bit '1' to x. Then, it appends $w = -N - 65 \bmod \ell$ '0' bits, and finally, it appends a 64-bit representation of $(N + w + 65)/\ell$. This number is an integer due to the choice of w, and it represents the number of message blocks in the final, padded message.

Since it must be possible to encode the number of message blocks in the padded message within 64 bits, the maximum message length is 65 bits short of $2^{64} - 1$ message blocks. For the short variants, the maximum message length in bits is therefore $512 \cdot (2^{64} - 1) - 65 = 2^{73} - 577$, and for the longer variants it is $1024 \cdot (2^{64} - 1) - 65 = 2^{74} - 1089$.

3.7 Summary

First, a message which is to be digested by Grøstl is padded using the padding function pad. The hash function then iterates a compression function $f : \{0, 1\}^\ell \times \{0, 1\}^\ell \to \{0, 1\}^\ell$, which is based on two permutations P and Q. If the output size n of the hash function is at most 256 bits, we set $\ell = 512$. For the longer variants, we set $\ell = 1024$. Hence, we ensure that $\ell \geq 2n$ for all cases. The initial value of Grøstl-n is the ℓ-bit representation of n. At the end, the output of the last call to f is processed by the output transformation Ω, which reduces the output size from ℓ to n bits.

4 Design decisions and design features

In this section, we explain the design decisions made for Grøstl and some features of the Grøstl design. First, we list a number of advantages of Grøstl compared to many other hash functions proposed in the past.

- **Security proof of the construction.** The compression function construction used in Grøstl is provably collision resistant and preimage resistant assuming that the permutations P and Q are ideal. See Section 4.1.

- **Flexibility.** The algorithm can be efficiently implemented on many platforms. The security parameter r, the number of rounds, can be easily changed.

- **Simplicity.** Both the construction and the design of the permutations are simple and easy to understand and remember.

- **Familiarity.** Being based on the well known Rijndael design, most cryptographers and cryptographic software implementors will quickly feel acquainted with Grøstl. Moreover, the design principles behind Rijndael have already proven themselves advantageous.

4.1 The security of the construction

The construction of the compression function f can be proved to be secure assuming that the two permutations P and Q are ideal [20]. The security proof states that at least $2^{\ell/4}$ evaluations of P and/or Q are required to find a collision for the hash function that iterates f, and that at least $2^{\ell/2}$ evaluations are required to find a preimage. Note that these levels are the square root of the security levels for an ideal compression function. However, since $\ell \geq 2n$, internal

collision and preimage attacks on the hash functions have complexities of at least $2^{n/2}$ and 2^n. This analysis assumes that the ℓ output bits of the last call to f are the final output of the hash function. However, in $\mathtt{Grøstl}$, an output transformation is applied. We discuss this output transformation in Section 4.6.

The security proof of the compression function construction assumes that the permutations P and Q are ideal. However, we do not claim that our permutations are ideal. We only use the security proof of the construction to show that the construction is sound. This is similar to using the security proof [8] of one of the PGV constructions [50] to show that this construction is sound, without claiming that the underlying block cipher is ideal. On the other hand, an attack that shows non-ideality of the permutations does not necessarily extend to an attack on the hash function.

4.2   AddRoundConstant

The purpose of adding round constants is to make each round different and at the same time this provides a natural opportunity to differentiate P and Q. If the rounds are all the same, then fixed points x such that $R(x) = x$ for the round function R extend to the entire permutation. For example, if $P = R^{10}$, then fixed points for R^2 and R^5 would also extend to P. Therefore, one can expect several fixed points for P, whereas for an ideal permutation only a single fixed point is expected. By choosing round-dependent constants for AddRoundConstant, we expect the number of fixed points of P and Q to be 1.

In implementations, using two different versions of AddRoundConstant for P and Q does not incur a large penalty, since most of the code implementing the round function can still be shared between P and Q. We decided to use simple round constants in order to reduce the performance penalty of this transformation. Hence, we chose to add a single byte onto a fixed matrix position. Since this is the only transformation in which there is a difference between P and Q, the round constants must be different.

4.3   SubBytes

The SubBytes transformation is the only non-linear transformation in $\mathtt{Grøstl}$. It uses the same s-box as used in Rijndael. For a walk-through of its properties, we refer to one of [12, 13].

The choice for this particular transformation was driven by the following reasoning:

- Size: 8-bit s-boxes are a convenient trade-off between implementation aspects (smallest word size on popular platforms) and cryptanalytic considerations. On the other hand, there are $2^8!$ different permutations to choose from.

- Single s-box rather than many different s-boxes: this is again a trade-off between implementation and cryptanalytic considerations.

- No random s-box: A structured s-box allows for significantly more efficient hardware implementation than a random s-box.

- The particular structure of the chosen s-box was already proposed in 1993 [45] and has therefore undergone a long period of study.

- Since the s-box is inherited from the AES, implementation aspects (especially in hardware) are well studied.

4.4 ShiftBytes and ShiftBytesWide

We had two design criteria for ShiftBytes and ShiftBytesWide. First, we needed shift values which result in optimal diffusion. Let $\nu_{t,c}(a_{i,j})$ be the number of times that a state byte $a_{i,j}$ affects every state byte of column c after t rounds. In detail, $\nu_{t,c}(a_{i,j})$ defines how often (or in how many ways) every state byte of column c depends on $a_{i,j}$. Hence, we have full diffusion after t rounds if $\nu_{t,c}(a_{i,j}) \geq 1$ for all columns c and state bytes $a_{i,j}$. In other words, each state byte is affected by every state byte $a_{i,j}$ at least once. Let t^* be the value of t for which this happens. Then we get optimal diffusion, if $\min(\nu_{t^*,c}(a_{i,j}))$ is maximal for a specific geometry.

Second, as this leaves a large number of set to choose from, we prefer shift values which are similar to the AES shift values to benefit from combined implementations. Combined implementations benefit if we shift into the same direction as AES and the values are of the form $\sigma = [0 + 4k_0, 1 + 4k_1, +4k_2, 3 + 4k_3, 0 + 4k_4, 1 + 4k_5, 2 + 4k_6, 3 + 4k_7]$ with $k_i \in \{0, 1, 2, 3\}$.

The shift values used for P_{512} and Q_{512} are the most obvious ones. They cause optimal diffusion after two rounds. For P_{1024} and Q_{1024} (ShiftBytesWide) we have searched for shift values with optimal diffusion after three rounds (two rounds is not possible) and get optimal diffusion if $\min(\nu_{3,c}(a_{i,j})) = 2$. We have chosen the first set of such values when sorted in lexicographical order, which can additionally be used for combined implementations.

4.5 MixBytes

The main design goal of the MixBytes transformation is to follow the wide trail strategy. Hence, the MixBytes transformation is based on an error-correcting code with the MDS (maximum distance separable) property. This ensures that both the differential and linear branch number is 9. In other words, a difference in $k > 0$ bytes of a column will result in a difference of at least $9 - k$ bytes after one MixBytes application.

Since there exist many MDS codes, we have chosen a code which can be implemented efficiently in many settings. The MixBytes transformation multiplies each column of A with the MDS matrix $B = \text{circ}(02, 02, 03, 04, 05, 03, 05, 07)$ (see Section 3.4.5) over the finite field $\mathbb{F}_{256}$. In most environments, the multiplication with a constant of this matrix is the most expensive part. The implementation costs can be reduced by using constants of low degree. The minimum degree of the constants for an MDS code of size 8 is 2. However, this comes at a higher cost for the additions due to a slightly higher Hamming weight of the elements. Therefore, we have chosen a set of values where we can compensate these costs by the possibility of combining more intermediate results during the matrix multiplication. Especially on 8-bit platforms, this results in more efficient implementations.

4.6 Output transformation

Since the size of the chaining variables is larger than the required output size, an output transformation is needed. Simple truncation would be a possibility. However, since the compression function is not ideal (see Section 6.2), we chose to apply a function which is believed to be one-way and collision resistant, but does not compress before the truncation.

Let $\omega(x) = P(x) \oplus x$. The Matyas-Meyer-Oseas construction [36] for hash functions based on block ciphers provides a compression function g based on the encryption function E_K (with K being the key) as follows:

$$g(h, m) = E_h(m) \oplus m.$$

This function g has been proved to provide a collision resistant and one-way hash function when iterated in the Merkle-Damgård mode [8], under the assumption that E is an ideal block cipher. This implies that g is collision resistant and one-way if h is fixed, since this corresponds

to hashing a one-block message. Hence, $\tilde{g}(m) = E_{h^*}(m) \oplus m$, where h^* is a constant, is one-way and collision resistant as well. Since $\tilde{g} = \omega$ with $P = E_{h^*}$, we believe that ω is one-way and collision resistant. This seems to make it difficult to attack Grøstl via the output transformation.

4.7 Number of rounds

The choice of the (recommended) number of rounds is in part based on the preliminary cryptanalysis results described in Section 6. In particular, the square/integral attack indicates that the permutations might be distinguishable from ideal if the number of rounds is 7 or less in the short variants, and 9 or less in the long variants. The final choice of the number of rounds provides a reasonably large security margin.

4.8 Absence of trap-doors

It should be clear that all constants used in Grøstl, including the s-box, have been selected in a way that does not leave enough freedom to deliberately insert trap-doors in the hash function. In general, we faithfully declare that we have not inserted any hidden weaknesses in Grøstl.

5 Modes of use for Grøstl

Grøstl can be used in a "randomisation mode", e.g., as a message authentication code. Such modes include an additional input, which can be a key, a salt, a randomisation value, etc. We believe that Grøstl is secure when used in existing randomisation modes making use of a hash function, but we also propose a dedicated MAC mode for Grøstl.

5.1 Message authentication

HMAC [5, 43] is a method of constructing a message authentication code (MAC) from a hash function. Given a message M, a key K and a hash function H, the HMAC construction is defined as follows.

$$\text{HMAC}(K, M) = H(\overline{K} \oplus \mathsf{opad} \| H(\overline{K} \oplus \mathsf{ipad} \| M)),$$

where $\overline{K}$ is K padded to a length equal to the block length of the hash function, and ipad and opad are two different constants as defined in [5]. HMAC has been proven to be secure if the compression function of the underlying hash function is a "dual" PRF [4]. A compression function is a dual PRF if it is a PRF when keyed via either the message block or the chaining input. We believe HMAC based on Grøstl is a secure MAC.

The HMAC construction requires two calls to the hash function, which in the case of Grøstl means that the output transformation must be evaluated twice. A more efficient method is the *envelope construction* [51]:

$$\text{MAC}(K, M) = H(\overline{K} \| \overline{M} \| K), \tag{2}$$

where $\overline{M}$ is M padded to a multiple of ℓ bits, and $\overline{K}$ is K padded to ℓ bits. We propose this envelope construction as a dedicated MAC mode using Grøstl. This construction has been proved to be a secure MAC under similar assumptions as HMAC [56]. For the security proof to hold, the key must be processed in blocks that are separate from the blocks of the message M, which explains the additional padding required.

5.2 Randomised hashing

In order to free the security of digital signatures from relying on the collision resistance of a hash function, the input message to the hash function can be randomised using a fresh random value z for every signature following the technique outlined in [15, 24]. The randomised message is then processed using the hash function. This procedure is called randomised hashing. Let the message be M, padded to a multiple of the message block length, and split into message blocks $m_1, \ldots, m_t$. The randomised variant $\tilde{H}$ of the hash function H given randomisation value z is then (roughly) defined as

$$\tilde{H}(z, M) = H(z \| (m_1 \oplus z) \| (m_2 \oplus z) \| \ldots \| (m_t \oplus z)).$$

We believe `Grøstl` to be suitable for use in this randomisation mode.

Being suitable for randomised hashing requires that the following attack [44] has complexity at least 2^{n-k}. The attacker chooses a message M of length at most 2^k bits. The attacker then receives a randomly chosen randomisation value z (not under the control of the attacker). The value $y = \tilde{H}(z, M)$ is computed, and the attacker's task is now to find a pair $(z^*, M^*) \neq (z, M)$ such that $\tilde{H}(z^*, M^*) = y$. When `Grøstl` is used in the mentioned randomisation mode, we restrict the length of the randomisation value to at most n bits.

5.3 Security claims for the mentioned modes of operation

We claim the following security levels for the applications where `Grøstl`-n is deployed. The claimed complexity of the "randomised hashing attack" assumes a first message of at most 2^k blocks.

Attack type	Claimed complexity	Best known attack
Forgery on n-bit HMAC	$2^{n/2}$	2^n
Key recovery on n-bit HMAC	$2^{|K|}$	$2^{|K|}$
Forgery on n-bit envelope MAC	$2^{n/2}$	2^n
Key recovery on n-bit envelope MAC	$2^{|K|}$	$2^{|K|}$
Randomised hashing	2^{n-k}	2^n

6 Cryptanalytic results

In this section, we describe some preliminary cryptanalysis results on `Grøstl`, and we state our security claims

6.1 Attacks exploiting properties of the permutations

We first consider well known attack methods that aim to exploit potential weaknesses in the permutations P and Q.

6.1.1 Differential cryptanalysis

The permutations P and Q have diffusion properties according to the wide trail design strategy. Since the MixBytes transformation has branch number 9, and ShiftBytes is diffusion optimal (moves the bytes in each column to eight different columns), it is guaranteed that for `Grøstl` there are at least $9^2 = 81$ active s-boxes in any four-round differential trail [13, Theorem 9.5.1]. Note that this holds for `Grøstl`-256 as well as for `Grøstl`-512. Hence, there are at least $2 \cdot 81 = 162$ and $3 \cdot 81 = 243$ active s-boxes in any eight-round, respectively twelve-round

differential trail. This, combined with the maximum difference propagation probability of the s-box of 2^{-6}, means that the probabilities of any differential trail (assuming independent rounds) over eight and twelve rounds (for either P or Q) are expected to be at most $2^{-6 \cdot 162} = 2^{-972}$, respectively 2^{-1458}. Therefore, in a classical differential attack where one specifies a differential trail for every round for both P and Q, there is only a very small chance that this would lead to a successful attack for Grøstl-256 and Grøstl-512.

In the collision attack [48] on Grindahl-256 [30], the low probability of any difference propagation through the s-box is circumvented by ignoring the actual values of differences, and instead only considering whether a byte is active or not. Since in Grindahl, a message block overwrites part of the state, the actual values of any differences in this part of the state are irrelevant. This approach means that the probabilistic behaviour of the hash function is now related to the MixColumns/MixBytes transformation, since without knowing the value of an input difference, one cannot predict the output difference. On the other hand, the number of degrees of freedom is essentially doubled, since one does not need to consider a fixed input/output difference. The relatively slow diffusion of Grindahl-256 combined with the continuous ability to influence the state led to the collision attack. In the Grøstl permutations, this approach will result in a complexity well above that of a birthday attack because diffusion is more effective (requiring only two rounds compared to four), and the attacker does not have continuous control over parts of the state. Moreover, since no part of the state is discarded (until the output transformation in the end), the actual value of a difference is significant and therefore, it seems that any input or output difference will have to (probabilistically) match a given difference.

6.1.2 Linear cryptanalysis

Linear and differential trails propagate in a very similar way. Since the MixBytes transformation has linear branch number 9, it is guaranteed that for Grøstl there are at least $9^2 = 81$ active s-boxes in any four-round linear trail [13, Theorem 9.5.1]. Hence, there are at least $2 \cdot 81 = 162$ and $3 \cdot 81 = 243$ active s-boxes in any eight-round, respectively twelve-round linear trail. Since the s-box has maximum correlation of 2^{-3}, the maximum correlation for any four-round linear trial is $2^{-3 \cdot 81} = 2^{-243}$. This means that the correlation of any linear trail over eight and twelve rounds (for either P or Q) are expected to be at most $2^{-3 \cdot 162} = 2^{-486}$, respectively 2^{-729}.

6.1.3 Integrals

Some of the best known attacks on AES are based on so-called integrals [11, 31]. Integrals can be specified also for Grøstl, and although it has not been shown how to utilise integrals in attacks on a hash function, they might say something about the used structure.

Integrals for Grøstl-256 are very similar to integrals for AES. We have identified an integral with 2^{120} texts over 6 rounds of Grøstl-256. The texts in this collection are balanced in every byte of the input and output. Also, we identified an integral with 2^{120} texts over 7 rounds of Grøstl-256. The texts in this collection are balanced in every byte of the input and balanced in every bit of the output. These are similar to the integrals for AES reported in [31]. Note that for AES reduced to 7 rounds, the last round is special. This is not the case for Grøstl.

We have identified integrals for Grøstl-512 for up to 9 rounds. For an 8-round variant the texts are balanced in every byte of the input and output; for an 9-round variant the texts are balanced in every byte of the input and in every bit of the output. For both these integrals, the number of texts is 2^{704}.

With the chosen number of rounds in the Grøstl permutations, 10 respectively 14, we believe it is safe to conclude that integrals cannot be used to show any non-random behaviour of Grøstl.

6.1.4 Algebraic cryptanalysis

It is well-known [10] that one can establish 39 quadratic equations (equations of degree two) over $\mathbb{F}_2$ in the input and output bits of the AES s-box, and there is one additional quadratic equation of probability $\frac{255}{256}$ for the AES s-box. Hence, this is also the case for the s-box in Grøstl. There is a total of 200 s-box applications for one encryption of the AES. Using these 40 equations for AES, it has been shown that from a single AES encryption, one can establish a set of 8000 quadratic equations in 1600 variables (unknowns). The solution of these equations can be used to derive the value of the secret key used in the encryption. The time complexity to solve the above mentioned system of equations for AES is unknown; to the best of our knowledge, it has not been shown that this can lead to an attack faster than an exhaustive search for the key.

For comparison, there is a total of 1280 s-box applications in the compression function of Grøstl-256 and a total of 3584 s-box applications in the compression function of Grøstl-512. It is clear that there are some advantages in an algebraic attack on a hash function compared to a similar attack on a block cipher, since there are no secret keys in the former. However, given that the number of s-box applications is much larger for Grøstl than for AES, we think it is safe to conclude that if an efficient algebraic attack method should be found which exploits the quadratic s-box equations in Grøstl, then a similar attack would be able to break the AES.

6.2 Generic collision attacks

This section deals with collision attacks that do not depend on weaknesses in P and Q. We distinguish between collision attacks on the compression function, and collision attacks on the hash function. Collision attacks on the compression function, where the chaining input is determined by the attack (and is not under the direct control of the attacker), cannot be directly extended to cover the full hash function. The security proof of the construction (1) relates to collision attacks on the hash function. Hence, we cannot rule out the possibility of generic attacks on the compression function below the $2^{\ell/4}$ bound. However, there are good reasons to believe that the bound holds also for the compression function as will be shown next.

6.2.1 Collision attacks on the compression function

Wagner's generalised birthday attack [53] applies to the compression function f: form four lists via the two functions $f_P(x) = P(x) \oplus x$ and $f_Q(x) = Q(x) \oplus x$. Note that $f(h, m) = f_P(h \oplus m) \oplus f_Q(m)$. Find a quadruple (x, x^*, y, y^*) such that $f_P(x) \oplus f_P(x^*) \oplus f_Q(y) \oplus f_Q(y^*) = 0$. Then the two pairs $(x \oplus y, y)$ and $(x^* \oplus y^*, y^*)$ collide.

This attack has complexity $2^{\ell/3}$, and hence is faster than a birthday attack on the compression function. Note that this is still above the proven bound of $2^{\ell/4}$ and above the complexity of a birthday attack on the hash function, since $n \leq \ell/2$. The attack does not provide the attacker with much control over the chaining input, and hence we do not see any methods to extend the attack to the full hash function.

Wagner notes that if f_P and f_Q are considered random functions, then finding a quadruple (x, x^*, y, y^*) such that $f_P(x) \oplus f_P(x^*) \oplus f_Q(y) \oplus f_Q(y^*) = 0$ has complexity at least $2^{\ell/4}$. Assuming this is correct, the complexity extends to the full hash function (where the output transformation is omitted) via the same proof as that of the Merkle-Damgård construction [14, 40].

Wagner's generalised birthday attack is the best attack on the compression function we are aware of. We note that in a Merkle-Damgård hash function, a collision attack on the compression function always extends to a pseudo- or free-start collision attack on the hash function. Hence, Wagner's generalised birthday attack can be used to carry out a free-start collision attack on

Grøstl in time $2^{\ell/3}$. Again, we remind the reader that this complexity is above the complexity of a birthday attack on Grøstl.

6.2.2 Collision attacks on the hash function

The construction (1) is provably collision resistant up to the level of $2^{\ell/4}$ permutation calls. Still, no collision attack of this complexity is known when the permutations are assumed to be ideal. The best known collision attack requires $2^{3\ell/8}$ permutation calls [20], but the true complexity in terms of compression function call equivalents is higher than $2^{\ell/2}$. Hence, a rather large security margin remains.

6.3 Generic attacks on the iteration

The internal state being *at least* twice the size of the hash value for all versions of Grøstl, generic attacks applying to the Merkle-Damgård construction cannot be applied to Grøstl directly via brute force or birthday attacks. However, since the construction used for Grøstl does not achieve security comparable to an ideal iterated hash function with the same internal state size, we do not claim that generic attacks do not apply using some other methods than the standard brute force and birthday attacks.

6.3.1 Multicollision attack

Recall that a d-collision is a set of d messages that all collide pairwise. The multicollision attack of Joux [27] on iterated hash functions applies also to Grøstl; the complexity to find a d-collision is roughly $\log_2(d)2^{\ell/2} \geq \log_2(d)2^n$. This should be compared to a brute-force multicollision attack on the hash function for which the complexity is around $(d!)^{1/d} \cdot 2^{n(d-1)/d}$. For values of d and n of cryptographic relevance, the brute-force attack is always faster than Joux's approach.

6.3.2 Second preimage attack

The second preimage attack of Kelsey and Schneier [29] on the Merkle-Damgård construction also seems to be complicated by the large internal state size. For an n-bit iterated hash function based on an n-bit compression function, given a first preimage of length 2^k message blocks this attack finds a second preimage of the same length in 2^{n-k} evaluations of the compression function. A variant of this attack was published in [2]. Using the techniques of [2, 29], the complexity of carrying out the second preimage attack on Grøstl given a 2^{64}-block first preimage is about $2^{\ell-64}$. For all the message digest sizes of Grøstl, this complexity is well above 2^{n-k}. Hence, our claimed security level for the second preimage resistance is at least 2^{n-k} for any first message of at most 2^k blocks. However, we do not know of an attack with complexity below 2^n.

6.3.3 Length extension attack

The length extension attack on Merkle-Damgård hash functions works as follows. Let (M, M^*) be a collision for the hash function H, with $|M| = |M^*|$. H pads M and M^* to $\overline{M}$ and $\overline{M}^*$ before hashing, and by choosing any message suffix y, we have that $B = \overline{M}\|y$ and $B^* = \overline{M}^*\|y$ also collide. Hence, a single collision gives rise to many new collisions that "come for free".

The length extension method is not trivial to carry out in Grøstl, unless the messages collide before the output transformation. Finding a collision before the output transformation takes time $2^{\ell/2} \geq 2^n$ by the birthday attack. As mentioned several times, there may be collision

attacks on the hash function with the output transformation omitted, that have complexity below the birthday attack, but we do not know of any such attack.

A related weakness of the Merkle-Damgård transformation is the following. Assume the two values $H(M)$ and $|M|$ are known, but M itself is not. Knowing $|M|$, one also knows how M was padded, and hence for any suffix y, one may compute $H(\overline{M}\|y)$, where $\overline{M}$ is the padded version of M, without knowing M. This weakness leads to attacks when a Merkle-Damgård hash function underlies a secret prefix MAC. In Grøstl, this attack does not seem possible due to the output transformation.

6.4 Fixed points

Most existing hash functions, for instance SHA-1 and SHA-2, are based on the Davies-Meyer construction [36], and hence *fixed points* can be easily found for these hash functions [41]. Some applications where this property can be used to attack hash functions have been identified, for instance, in finding an expandable message to carry out the second preimage attack [17, 29]. However, finding an expandable message is only one part of the second preimage attack, and in most cases it is not the most time-consuming task.

Fixed points can also be efficiently found for the compression function f of Grøstl: Choose m arbitrarily, and let $h = P^{-1}(Q(m)) \oplus m$. Then $f(h, m) = h$. Hence, h is computed as a (claimed) one-way function of m, and therefore is not under the direct control of the attacker.

In the case of Grøstl, we note that the internal state is at least twice the size of the hash value, and hence the cost of constructing, e.g., an expandable message using fixed points is expected to be about $2^{\ell/2} \geq 2^n$.

6.5 Security claims and summary of known attacks

With the number of rounds proposed in Section 3.4.6, we claim the following security levels for the Grøstl-n hash function. In the second preimage attack, the first preimage is assumed to be of length at most 2^k blocks.

Attack type	Claimed complexity	Best known attack
Collision	$2^{n/2}$	$2^{n/2}$
d-collision	$\lg(d) \cdot 2^{n/2}$	$(d!)^{1/d} \cdot 2^{n(d-1)/d}$
Preimage	2^n	2^n
Second preimage	2^{n-k}	2^n

Even though compression function attacks do not necessarily translate into attacks on the hash function, we claim the following properties for the compression function:

Attack type	Claimed complexity	Best known attack
Collision	$2^{\ell/4}$	$2^{\ell/3}$
Preimage	$2^{\ell/2}$	$2^{\ell/2}$

7 Implementation aspects

Like Rijndael, Grøstl can be efficiently implemented on a wide variety of processors and allows many trade-offs between resource requirements (memory, registers) and speed. In this section, we describe and estimate performance and resource requirements of implementations on 64-, 32-, and 8-bit architectures, as well as on ASICs and FPGA hardware. As Grøstl is designed to prevent preference for a particular word size, this will also allow efficient implementation of future architectures (like Intel's AVX with 256-bit registers [26]).

7.1 Software implementations

In software, Grøstl is targeted 64-bit processors, but performance is nearly as good on 32-bit processors offering MMX instructions.

7.1.1 64-bit processors

Grøstl can be efficiently implemented on 64-bit processors following a technique very similar to the efficient 32-bit implementation of Rijndael [12]. Consider an implementation of the round function of P_{512} focusing on the effect on column 0. Assume that the AddRoundConstant transformation adds the byte C to $a_{0,0}$. Note that the new column 0 after the round function has been applied depends solely on the 8 bytes $a_{i,i}$, $0 \leq i < 8$, because the ShiftBytes transformation moves these bytes into column 0.

As an example, the round function has the following effect on $a_{0,0}$, the new value of which we denote by $a'_{0,0}$.

$$a'_{0,0} \leftarrow \quad 02 \times S(a_{0,0} \oplus C) \oplus 02 \times S(a_{1,1}) \oplus 03 \times S(a_{2,2}) \oplus 04 \times S(a_{3,3}) \oplus$$
$$05 \times S(a_{4,4}) \oplus 03 \times S(a_{5,5}) \oplus 05 \times S(a_{6,6}) \oplus 07 \times S(a_{7,7}).$$

Similarly, the effect on $a_{1,0}$ is

$$a'_{1,0} \leftarrow \quad 07 \times S(a_{0,0} \oplus C) \oplus 02 \times S(a_{1,1}) \oplus 02 \times S(a_{2,2}) \oplus 03 \times S(a_{3,3}) \oplus$$
$$04 \times S(a_{4,4}) \oplus 05 \times S(a_{5,5}) \oplus 03 \times S(a_{6,6}) \oplus 05 \times S(a_{7,7}).$$

If we continue, we see that, e.g., $a_{0,0}$ affects every byte of the column by the addition of $b \times S(a_{0,0} \oplus C)$, where b is a value from the first column of the matrix B (defined in Section 3.4.5). Hence, when the column is represented by a 64-bit word in an implementation, we may compute the effect of $a_{0,0}$ on *all* bytes in the new column 0 by a single table lookup, the output of which is exactly 8 concatenations of $b \times S(a_{0,0} \oplus C)$, with b varying as defined by the matrix B. Let the table be T_0 containing 256 64-bit words. The value $T_0[i]$ at index i (ignore the addition of C for a moment) will then be

$$02 \times S(i) \parallel 07 \times S(i) \parallel 05 \times S(i) \parallel 03 \times S(i) \parallel 05 \times S(i) \parallel 04 \times S(i) \parallel 03 \times S(i) \parallel 02 \times S(i),$$

interpreted as an 8-byte (64-bit) word. Here, we define the first byte of the word to mean the byte of row 0 in A. In practice, the most convenient ordering depends on the endianness of the processor (the ordering used above is more convenient on big-endian processors, whereas on little-endian processors the byte ordering should be reversed).

A byte in a different row affects the column in a different way, and hence we must define 8 different tables $T_0, \ldots, T_7$. The only difference between them is the ordering of the bytes; they are rotated versions of each other, since the matrix B is circulant. To save space, a single table can be used, and the rotations can be done afterwards.

To sum up, column 0 can be computed as

$$T_0[a_{0,0} \oplus C] \oplus T_1[a_{1,1}] \oplus T_2[a_{2,2}] \oplus T_3[a_{3,3}] \oplus T_4[a_{4,4}] \oplus T_5[a_{5,5}] \oplus T_6[a_{6,6}] \oplus T_7[a_{7,7}],$$

hence using 8 table lookups and 8 XORs (7 for all other columns, since adding C is only needed in column 0).

When the columns are internally represented as 64-bit words, in most programming languages we don't have direct access to the bytes $a_{i,i}$, and hence we must access them by a right-shift and a logical *and*. However, many processors provide instructions for accessing a particular byte of a word.

We note that this technique requires storing 8 tables of 256 64-bit words, taking up 16 kilobytes of memory. As mentioned, a single table of 2 kilobytes can be used instead, but then a rotation is needed for every 7 out of 8 table lookups. This can be generalised; with k tables, $8-k$ rotations are needed for every 8 table lookups. A crude estimate on the performance loss with $0 < k \leq 8$ tables compared to 8 tables is a factor about $\frac{23-k}{15}$. This is based on the estimate that a rotation, a table lookup, and an XOR take about the same time to carry out.

7.1.2 32-bit processors

On a 32-bit processor the above technique cannot be applied directly, but there is a (slower) variant operating with 32-bit words. This method requires half the amount of memory compared to the 64-bit implementation described above, and (roughly) twice the amount of computation. For more details we refer to the Whirlpool specification [3]. The same time/memory trade-offs as mentioned above are possible.

On 32-bit microprocessors with SIMD instruction sets such as MMX, SSE, or SSE2, an implementation like the one described for 64-bit processors is possible. Some overhead is introduced compared to the implementation on a native 64-bit processor, but nevertheless, performance on such 32-bit processors is almost as good as on a 64-bit processor. Most modern 32-bit processors used in personal computers provide these instruction sets. These include virtually all Intel and AMD processors since 1997.

7.1.3 8-bit processors

On 8-bit processors, the round transformations can be applied individually on a byte-by-byte basis. Both the SubBytes and the MixBytes operation can be efficiently realised with lookups in small tables or computed without lookup tables. Various implementation techniques that allow a trade-off between memory usage and performance are possible. Especially in the computation of the MixBytes operation, many intermediate results can be reused depending on the memory requirements. Note that there is no setup time needed for the 8-bit implementation of Grøstl.

As an example for possible trade-offs, preliminary implementation results suggest that Grøstl-256 and Grøstl-224 can be implemented with a performance of (roughly) between 400 and 500 cycles/byte on an 8-bit AVR micro-controller (ATmega163)[2] using between less than 100 bytes and 850 bytes of RAM, and a code size of less than 1KB. The code includes a 256-byte lookup table for SubBytes and up to two 256-byte lookup tables for the multiplication with the constant 02 and/or 04 in the finite field $\mathbb{F}_{256}$ for MixBytes.

Table 1: Examples for Grøstl performance estimates for long messages on 8-bit implementations on the ATmega163 micro-controller.

Hash function	Processor	Memory (bytes)	Speed (cycles/byte)
Grøstl-224/256	ATmega163	<100	475
	ATmega163	850	415
Grøstl-384/512	ATmega163	<200	665
	ATmega163	950	580

---

[2]Running at e.g., 8MHz with no operating system

7.2 Benchmarks on PC platforms

The performance in software of the submitted optimised `Grøstl` implementations in ANSI C has been tested on a number of combinations of processors, operating systems, and compilers. The benchmarks refer to the hash computation of long messages.

The processors used in the tests are shown in the table below (the IDs are meant to serve as mnemonics that will be used in place of the long description in the presentation of the test results).

ID	Processor	Clock speed	Native word size
C2D	Intel Core 2 Duo E4600	2.4 GHz	64 bits
PM	Intel Pentium M 760	2.0 GHz	32 bits

The amount of random access memory available is 2 GB for the C2D processor, and 1 GB for the PM processor. The operating systems used can be found in the table below.

ID	Operating system (OS)
Vis32	Microsoft Windows Vista 32-bit
Vis64	Microsoft Windows Vista 64-bit
XP	Microsoft Windows XP 32-bit
Ubu32	GNU/Linux Ubuntu 8.04 32-bit
Ubu64	GNU/Linux Ubuntu 8.04 64-bit

Finally, the compilers used and the flags set in these can be found in the table below.

ID	Compiler	Version	Flags
gcc64	gcc	4.2.4	`-O3 -funroll-loops`
			`-fno-regmove -fmodulo-sched`
icc64	Intel C compiler	10.1	`-O3 -xP -ipo`
icc32	Intel C compiler	10.0	`-O3 -xN -ipo`
VS64	Visual Studio 2005 (64-bit)	8.0	`/O2 /Ot /GL /LTCG`
VS32	Visual Studio 2008 (32-bit)	9.0	`/O2 /Ot /GL /LTCG`

In all except a single case, the 32-bit optimised implementation was used when compiled in 32-bit mode (the compiler ID ends with 32), and the 64-bit optimised implementation was used when compiled in 64-bit mode. The exception is the compiler icc32, which used the optimised 64-bit implementation, since this turned out to be faster.

Table 2 presents the benchmarking results. The implementations were developed in a Linux environment, and hence do not perform as well in a Windows environment. We expect that the differences would cancel out if the implementations were targeted the Windows/Visual Studio combination. Better performance, in particular on common 32-bit processors, can be achieved; see the benchmarks for implementations using MMX intrinsics or inline assembly in Section 7.2.2.

7.2.1 Algorithm setup time

`Grøstl` spends virtually no setup time before a message can start to be digested. All that is required is to check that the requested output size is a valid one, and if so initialise the state, clear the message block buffer, and reset the block counter. In practice, time for memory allocation will often also be needed.

To demonstrate, we counted the number of cycles required to run the `Init()` function of the optimised implementations. The results are shown in Table 3. We stress that these timings include the time needed for memory allocation.

Table 2: Benchmarks of the optimised `Grøstl` implementations in C.

Hash function	Processor	OS	Compiler	Speed (cycles/byte)
Grøstl-224/256	C2D	Ubu64	gcc64	25.7
	C2D	Ubu64	icc64	25.4
	C2D	Vis64	VS64	27.5
	C2D	Vis32	VS32	77.9
	PM	Ubu32	icc32	65.2
	PM	XP	VS32	79.8
Grøstl-384/512	C2D	Ubu64	gcc64	45.0
	C2D	Ubu64	icc64	36.9
	C2D	Vis64	VS64	42.2
	C2D	Vis32	VS32	123.4
	PM	Ubu32	icc32	85.8
	PM	XP	VS32	126.8

Table 3: The number of cycles required to run the `Init()` function. In all cases, the `gcc` compiler (v. 4.2.4) was used in Ubuntu 8.04 with optimisation flag `-O3`.

Hash function	Processor	OS	Cycles for `Init()`
Grøstl-224/256	C2D	Ubu64	336
	PM	Ubu32	335
Grøstl-384/512	C2D	Ubu64	528
	PM	Ubu32	478

7.2.2 UPDATED v.1.1 Additional implementations

A set of implementations in C using MMX intrinsics, as well as a set using inline assembly with MMX and SSE instructions have also been developed. The inline assembly implementation has been optimised for the Intel Core 2 Duo processor and shows significant improvements in both 32-bit and 64-bit mode, compared to the benchmarks in Table 2. On the Pentium M, the implementation using MMX intrinsics performs better than the inline assembly implementation in the case of `Grøstl`-256, but not in the case of `Grøstl`-512. Table 4 shows the best benchmarks obtained on these two processors in Linux Ubuntu using the `gcc` compiler with the `-O3` optimisation flag (and `-mmmx` or `-msse2` as necessary).

Table 4: Best `Grøstl` benchmarks for implementations using MMX intrinsics or inline assembly with MMX and SSE instructions, as indicated.

Hash function	Processor	Mode	Implementation	Speed (cycles/byte)
Grøstl-224/256	C2D	64-bit	Inline assembly	22.4
	C2D	32-bit	Inline assembly	22.9
	PM	32-bit	MMX intrinsics	28.9
Grøstl-384/512	C2D	64-bit	Inline assembly	30.1
	C2D	32-bit	Inline assembly	37.5
	PM	32-bit	Inline assembly	57.1

7.3 Hardware implementations

Potential settings and scenarios for hardware implementations can be at least as diverse as for software implementations. The many different ways to implement Grøstl allows for a wide range of trade-offs between throughput, latency, gate count, power consumption, etc. Grøstl can be implemented efficiently on architectures with data paths starting from 8-bit up to 1024-bit. In the following, we will give estimates and first implementation results for some extreme examples.

7.3.1 Low-gate count implementations

To illustrate the multitude of different implementation trade-off possibilities the design offers, we also consider implementations where very small area requirements and very low-power requirements are important.

We estimate that Grøstl-256 and Grøstl-224 can be implemented on an ASIC with standard-cell libraries requiring an area of less than 15000 gate equivalents (GE). The dominating factor here is the memory. We use 12390 GE for register-based RAM in our estimate instead of RAM hard-macros. Since low gate count implementations are usually also low-power implementations, the register-based RAM can be used to minimise power consumption by clock gating. We base our estimates on numbers obtained from actual implementations of the AES and other algorithms [18, 19]. This results in 354 GE for the SubBytes and 800 GE for the MixBytes transformation of Grøstl-256. Table 5 gives an overview for all Grøstl variants.

Table 5: Estimates for a low-power architecture with an 8-bit data path implementation of Grøstl, that also has a low gate count.

Part	GE	
	Grøstl-224/256	Grøstl-384/512
RAM	12390	24780
SubBytes	354	354
MixBytes	800	800
Others (conservative)	1400	2000
Sum	< 15000	< 28000

7.3.2 High-throughput implementations

High-throughput implementations of Grøstl can be developed using data paths up to 512 or 1024 bit. Further, the execution of the two permutations can be implemented in parallel or pipelined and interleaved. This results in 1 cycle per round and 10 or 14 cycles per compression function computation. Preliminary implementation results of Grøstl using a $0.18\mu m$ standard cell library are given in Table 6. In these high-throughput implementations the permutations P and Q are computed in parallel and use a 512-bit or 1024-bit data path for each permutation. The s-box is implemented in combinational logic[9].

7.3.3 FPGA implementations

We base our estimates for the 64-bit FPGA implementations of Grøstl on the Whirlpool implementations of [49] on a Xilinx Virtex 2P (xc2vp40-7fg676) device. The gate count and throughput of Grøstl-256 is similar to that of Whirlpool-512 if an 8-bit lookup table is used for the s-box implementation. In [49], the Whirlpool s-box is implemented using 4-bit lookup

Table 6: `Grøstl` implementation results for high-throughput ASIC implementations using a 0.18μm technology.

Hash function	Data path	Size (GE)	Throughput (Mbit/s)	Frequency (MHz)
Grøstl-224/256	512-bit	130640	4379	85.5
Grøstl-384/512	1024-bit	340498	6225	85.1

tables. Since this is not possible for `Grøstl`, an 8-bit lookup table or an implementation in combinational logic has to be used. This may result in about two times the area requirements, which has been taken into account in our estimates. For the `Grøstl`-512 estimation, the area is doubled due to the data path of 1024-bit. The throughput is about 1.5 times higher since more rounds are used but the block size is doubled.

Additionally, two preliminary implementation results of a high-throughput FPGA implementation on a high-end Xilinx Virtex-5 (xc5vlx110-3ff1760) and a standard Xilinx Spartan 3 (xc3s5000-5fg676) are given in Table 7. In both implementations, the s-box is implemented using hardware lookup tables. The permutations P and Q are computed in parallel and use 512-bit and 1024-bit data paths for `Grøstl`-256 and `Grøstl`-512, respectively. Note that the number of CLB slices is a device dependent measure and does not allow direct comparisons.

Table 7: `Grøstl` implementation results and estimates[*] for FPGA implementations using a 64-bit data path and high-throughput implementations for 512-bit or 1024-bit data paths.

Hash function	Data path	Device	Size (CLB slices)	Throughput (Mbit/s)	Frequency (MHz)
Grøstl-224/256	64-bit*	Virtex 2P	3000-4000	400	75-125
	512-bit	Spartan 3	6582	4439	86.7
	512-bit	Virtex 5	1722	10276	200.7
Grøstl-384/512	64-bit*	Virtex 2P	6000-8000	300	35-60
	1024-bit	Spartan 3	20233	5901	80.7
	1024-bit	Virtex 5	5419	15395	210.5

7.4   Implementation attacks

Whenever a key is handled by a machine that implements a cryptographic mechanism in addition to inputs and outputs, various side-channel information may be available to an attacker. Sources for such side-channels can be (but are not limited to) timing information, power consumption and electromagnetic emanations, error messages, etc.

7.4.1   Cache based timing attacks

Cache based timing attacks have been mentioned, discussed and investigated in [6, 7, 32, 47, 52]. Bit-slicing techniques applied to `Grøstl` allow for implementations that are resistant against cache based timing attacks. In the case of AES, implementations using these techniques are in fact the fastest on many modern platforms. Also for `Grøstl`, we do not expect a severe performance loss when using such an implementation technique.

7.4.2 Power- and EM side-channel attacks

Published in 1999 [33], side-channel attacks that exploit information from the power consumption in the form of differential power analysis (DPA) attacks and electromagnetic emanations turn out to be a real threat for many implementations. Many generic (e.g., dual-rail logic) countermeasures and countermeasures specialised for particular algorithms have been proposed since then. Again, the similarity of our proposal to the AES allows to reuse many ideas from previous work.

Popular MAC implementations such as HMAC-SHA-1 and HMAC-SHA-2 have been exposed to DPA attacks [34, 37]. MACs constructed using block cipher based hash functions can be analysed against side channel attacks by assuming that the block cipher or the compression function is side channel resistant. Under this assumption, DPA attacks on several hash function based MACs including HMAC instantiated with the provably secure block cipher based hash functions were demonstrated in [21, 22, 46]. For the MAC modes of Grøstl, we note that these observations do not seem to be directly applicable.

7.4.3 On countermeasures

The in Section 7.4.1 mentioned work on constant-time implementations of AES, and the huge body of work on countermeasures against power- and EM side-channel attacks (see e.g., [35] for a good overview) which is also primarily applied to AES, give a sound basis to counter implementation attacks. On top of that, instruction set extensions that are frequently proposed by CPU manufacturers may be used as well. Preliminary implementation results suggest that the new crypto-related instructions, which Intel is going to introduce in upcoming versions of their CPUs [25], can efficiently be used to implement Grøstl in a constant-time manner and hence resistant against timing attacks.

Grøstl was, on purpose, not designed to use such instructions directly, as diffusion and confusion would be affected and all other platforms penalised. Still, this yet again serves as a powerful illustration for the many ways Grøstl can be implemented.

8 Conclusion

The SHA-3 candidate Grøstl has been proposed. Grøstl is a permutation-based hash function, based on a construction which is provably collision resistant when the permutations are assumed to be ideal. The particular permutations used in Grøstl are based on components of the Rijndael block cipher. As an effect of this, Grøstl has excellent diffusion and confusion properties. The design of Grøstl is very simple and easy to understand. Therefore, it is relatively easy to identify possible attacks and thereby easy to gain confidence in the strength of the construction. We believe that Grøstl is a very strong hash function, yet it can be efficiently implemented on a wide range of platforms. Reference and optimised implementations, test vectors, this document and other information on Grøstl is available at [23].

Acknowledgements

References

[1] R. J. Anderson and E. Biham. TIGER: A Fast New Hash Function. In D. Gollmann, editor, *Fast Software Encryption 1996, Proceedings*, volume 1039 of *Lecture Notes in Computer Science*, pages 89–97. Springer, 1996.

[2] E. Andreeva, C. Bouillaguet, P.-A. Fouque, J. J. Hoch, J. Kelsey, A. Shamir, and S. Zimmer. Second Preimage Attacks on Dithered Hash Functions. In N. Smart, editor, *Advances in Cryptology – EUROCRYPT 2008, Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 270–288. Springer, 2008.

[3] P. S. L. M. Barreto and V. Rijmen. The WHIRLPOOL Hashing Function. Submitted to NESSIE, September 2000. Revised May 2003. Available: `http://paginas.terra.com.br/informatica/paulobarreto/WhirlpoolPage.html` (2008/07/08).

[4] M. Bellare. New Proofs for NMAC and HMAC: Security Without Collision-Resistance. In C. Dwork, editor, *Advances in Cryptology – CRYPTO 2006, Proceedings*, volume 4117 of *Lecture Notes in Computer Science*, pages 602–619. Springer, 2006.

[5] M. Bellare, R. Canetti, and H. Krawczyk. Keying Hash Functions for Message Authentication. In N. Koblitz, editor, *Advances in Cryptology – CRYPTO '96, Proceedings*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 1996.

[6] D. J. Bernstein. Cache-timing attacks on AES. Available: `http://cr.yp.to/antiforgery/cachetiming-20050414.pdf` (2008/10/30).

[7] G. Bertoni, V. Zaccaria, L. Breveglieri, M. Monchiero, and G. Palermo. AES Power Attack Based on Induced Cache Miss and Countermeasure. In *International Conference on Information Technology: Coding and Computing (ITCC 2005), Proceedings*, volume 1, pages 586–591. IEEE Computer Society, April 2005.

[8] J. Black, P. Rogaway, and T. Shrimpton. Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002, Proceedings*, volume 2442 of *Lecture Notes in Computer Science*, pages 320–335. Springer, 2002.

[9] D. Canright. A Very Compact S-Box for AES. In J. R. Rao and B. Sunar, editors, *Cryptographic Hardware and Embedded Systems – CHES 2005, Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 441–455. Springer, 2005.

[10] N. Courtois and J. Pieprzyk. Cryptanalysis of Block Ciphers with Overdefined Systems of Equations. In Y. Zheng, editor, *Advances in Cryptology – ASIACRYPT 2002, Proceedings*, volume 2501 of *Lecture Notes in Computer Science*, pages 267–287. Springer, 2002.

[11] J. Daemen, L. R. Knudsen, and V. Rijmen. The Block Cipher Square. In E. Biham, editor, *Fast Software Encryption 1997, Proceedings*, volume 1267 of *Lecture Notes in Computer Science*, pages 149–165. Springer, 1997.

[12] J. Daemen and V. Rijmen. AES Proposal: Rijndael. AES Algorithm Submission, September 3, 1999. Available: `http://csrc.nist.gov/archive/aes/rijndael/Rijndael-ammended.pdf` (2008/10/29).

[13] J. Daemen and V. Rijmen. *The Design of Rijndael*. Springer, 2002.

[14] I. Damgård. A Design Principle for Hash Functions. In G. Brassard, editor, *Advances in Cryptology – CRYPTO '89, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 416–427. Springer, 1990.

[15] Q. Dang. Draft NIST Special Publication 800-106: Randomized Hashing for Digital Signatures, 2008. Available: `http://csrc.nist.gov/publications/drafts/800-106/2nd-Draft_SP800-106_July2008.pdf` (2008/10/17).

[16] C. De Cannière and C. Rechberger. Finding SHA-1 Characteristics: General Results and Applications. In X. Lai and K. Chen, editors, *Advances in Cryptology – ASIACRYPT 2006, Proceedings*, volume 4284 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2006.

[17] R. D. Dean. *Formal Aspects of Mobile Code Security*. PhD thesis, Princeton University, January 1999.

[18] M. Feldhofer, S. Dominikus, and J. Wolkerstorfer. Strong Authentication for RFID Systems Using the AES Algorithm. In M. Joye and J.-J. Quisquater, editors, *Cryptographic Hardware and Embedded Systems – CHES 2004, Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 357–370. Springer, 2004.

[19] M. Feldhofer and C. Rechberger. A Case Against Currently Used Hash Functions in RFID Protocols. In R. Meersman, Z. Tari, and P. Herrero, editors, *OTM Workshops (1)*, volume 4277 of *Lecture Notes in Computer Science*, pages 372–381. Springer, 2006.

[20] P.-A. Fouque, J. Stern, and S. Zimmer. Cryptanalysis of Tweaked Versions of SMASH and Reparation. In *Selected Areas in Cryptography 2008, Proceedings*, Lecture Notes in Computer Science. Springer. To appear.

[21] P. Gauravaram and K. Okeya. An Update on the Side Channel Cryptanalysis of MACs Based on Cryptographic Hash Functions. In *Progress in Cryptology – INDOCRYPT 2007, Proceedings*, volume 4859 of *Lecture Notes in Computer Science*, pages 393–403. Springer, 2007.

[22] P. Gauravaram and K. Okeya. Side Channel Analysis of Some Hash Function Based MACs: A Response to SHA-3 Requirements. In L. Chen, M. D. Ryan, and G. Wang, editors, *International Conference on Information and Communications Security – ICICS 2008, Proceedings*, volume 5308 of *Lecture Notes in Computer Science*, pages 111–127. Springer, 2008.

[23] The Grøstl web page. `http://www.groestl.info`.

[24] S. Halevi and H. Krawczyk. Strengthening Digital Signatures Via Randomized Hashing. In C. Dwork, editor, *Advances in Cryptology – CRYPTO 2006, Proceedings*, volume 4117 of *Lecture Notes in Computer Science*, pages 41–59. Springer, 2006.

[25] Intel Corporation. Advanced Encryption Standard (AES) Instructions Set. Available: `http://softwarecommunity.intel.com/articles/eng/3788.htm` (2008/10/30).

[26] Intel Corporation. Intel AVX: New Frontiers in Performance Improvements and Energy Efficiency. Available: `http://softwarecommunity.intel.com/articles/eng/3775.htm` (2008/10/30).

[27] A. Joux. Multicollisions in Iterated Hash Functions. Application to Cascaded Constructions. In M. K. Franklin, editor, *Advances in Cryptology – CRYPTO 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 306–316. Springer, 2004.

[28] J. Kelsey and S. Lucks. Collisions and Near-Collisions for Reduced-Round Tiger. In M. J. B. Robshaw, editor, *Fast Software Encryption 2006, Proceedings*, volume 4047 of *Lecture Notes in Computer Science*, pages 111–125. Springer, 2006.

[29] J. Kelsey and B. Schneier. Second Preimages on n-Bit Hash Functions for Much Less than 2^n Work. In R. Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 474–490. Springer, 2005.

[30] L. R. Knudsen, C. Rechberger, and S. S. Thomsen. The Grindahl Hash Functions. In A. Biryukov, editor, *Fast Software Encryption 2007, Proceedings*, volume 4593 of *Lecture Notes in Computer Science*, pages 39–57. Springer, 2007.

[31] L. R. Knudsen and V. Rijmen. Known-Key Distinguishers for Some Block Ciphers. In K. Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007, Proceedings*, volume 4833 of *Lecture Notes in Computer Science*, pages 315–324. Springer, 2007.

[32] P. C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In N. Koblitz, editor, *Advances in Cryptology – CRYPTO '96, Proceedings*, number 1109 in Lecture Notes in Computer Science, pages 104–113. Springer, 1996.

[33] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *Advances in Cryptology – CRYPTO '99, Proceedings*, volume 1666 of *Lecture Notes in Computer Science*, pages 388–397. Springer, 1999.

[34] K. Lemke, K. Schramm, and C. Paar. DPA on n-bit sized boolean and arithmetic operations and its application to IDEA, RC6, and the HMAC-construction. In *Cryptographic Hardware and Embedded Systems – CHES 2004, Proceedings*, volume 3156 of *Lecture Notes in Computer Science*, pages 205–219. Springer, 2004.

[35] S. Mangard, E. Oswald, and T. Popp. *Power Analysis Attacks – Revealing the Secrets of Smart Cards*. Springer, 2007.

[36] S. M. Matyas, C. H. Meyer, and J. Oseas. Generating strong one-way functions with crypographic algorithm. *IBM Technical Disclosure Bulletin*, 27(10A):5658–5659, 1985.

[37] R. P. McEvoy, M. Tunstall, C. C. Murphy, and W. P. Marnane. Differential Power Analysis of HMAC Based on SHA-2, and Countermeasures. In *Workshop on Information Security Applications – WISA 2007, Revised Selected Papers*, volume 4867 of *Lecture Notes in Computer Science*, pages 317–332. Springer, 2007.

[38] F. Mendel, B. Preneel, V. Rijmen, H. Yoshida, and D. Watanabe. Update on Tiger. In R. Barua and T. Lange, editors, *Progress in Cryptology – INDOCRYPT 2006, Proceedings*, volume 4329 of *Lecture Notes in Computer Science*, pages 63–79. Springer, 2006.

[39] F. Mendel and V. Rijmen. Cryptanalysis of the Tiger Hash Function. In K. Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007, Proceedings*, volume 4833 of *Lecture Notes in Computer Science*, pages 536–550. Springer, 2007.

[40] R. C. Merkle. One Way Hash Functions and DES. In G. Brassard, editor, *Advances in Cryptology – CRYPTO '89, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 428–446. Springer, 1990.

[41] S. Miyaguchi, K. Ohta, and M. Iwata. Confirmation that Some Hash Functions Are Not Collision Free. In I. Damgård, editor, *Advances in Cryptology – EUROCRYPT '90, Proceedings*, volume 473 of *Lecture Notes in Computer Science*, pages 326–343. Springer, 1991.

[42] National Institute of Standards and Technology. FIPS PUB 180-2, Secure Hash Standard. Federal Information Processing Standards Publication 180-2, U.S. Department of Commerce, August 2002.

[43] National Institute of Standards and Technology. FIPS PUB 198, The Keyed-Hash Message Authentication Code (HMAC). Federal Information Processing Standards Publication 198, U.S. Department of Commerce, March 2002.

[44] 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: `http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Nov07.pdf` (2008/10/17).

[45] K. Nyberg. Differentially uniform mappings for cryptography. In T. Helleseth, editor, *Advances in Cryptology – EUROCRYPT '93*, volume 765 of *Lecture Notes in Computer Science*, pages 55–64. Springer, 1994.

[46] K. Okeya. Side Channel Attacks Against HMACs Based on Block-Cipher Based Hash Functions. In *Australasian Conference on Information Security and Privacy – ACISP 2006, Proceedings*, volume 4058 of *Lecture Notes in Computer Science*, pages 432–443. Springer, 2006.

[47] D. Page. Theoretical Use of Cache Memory as a Cryptanalytic Side-Channel. Technical Report CSTR-02-003, University of Bristol, Department of Computer Science, June 2002. Available: `http://www.cs.bris.ac.uk/Publications/Papers/1000625.pdf` (2008/10/30).

[48] T. Peyrin. Cryptanalysis of Grindahl. In K. Kurosawa, editor, *Advances in Cryptology – ASIACRYPT 2007, Proceedings*, volume 4833 of *Lecture Notes in Computer Science*, pages 551–567. Springer, 2007.

[49] N. Pramstaller, C. Rechberger, and V. Rijmen. A compact FPGA implementation of the hash function Whirlpool. In *FPGA '06: Proceedings of the 2006 ACM/SIGDA 14th international symposium on Field programmable gate arrays*, pages 159–166, New York, NY, USA, 2006. ACM.

[50] B. Preneel, R. Govaerts, and J. Vandewalle. Hash Functions Based on Block Ciphers: A Synthetic Approach. In D. R. Stinson, editor, *Advances in Cryptology – CRYPTO '93, Proceedings*, volume 773 of *Lecture Notes in Computer Science*, pages 368–378. Springer, 1994.

[51] G. Tsudik. Message Authentication with One-Way Hash Functions. In *INFOCOM '92, Proceedings*, pages 2055–2059, 1992.

[52] Y. Tsunoo, E. Tsujihara, K. Minematsu, and H. Miyauchi. Cryptanalysis of Block Ciphers Implemented on Computers with Cache. In *International Symposium on Information Theory and Its Applications (ISITA 2002), Proceedings*, October 2002.

[53] D. Wagner. A Generalized Birthday Problem. In M. Yung, editor, *Advances in Cryptology – CRYPTO 2002, Proceedings*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–303. Springer, 2002.

[54] X. Wang, Y. L. Yin, and H. Yu. Finding Collisions in the Full SHA-1. In V. Shoup, editor, *Advances in Cryptology – CRYPTO 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 17–36. Springer, 2005.

[55] Wikipedia. Close-mid front rounded vowel. `http://en.wikipedia.org/wiki/Close-mid_front_rounded_vowel` (2008/08/21).

[56] K. Yasuda. "Sandwich" Is Indeed Secure: How to Authenticate a Message with Just One Hashing. In J. Pieprzyk, H. Ghodosi, and E. Dawson, editors, *Australasian Conference on Information Security and Privacy – ACISP 2007, Proceedings*, volume 4586 of *Lecture Notes in Computer Science*, pages 355–369. Springer, 2007.

A   The name

Gröstl is an Austrian dish, usually made of leftover potatoes and pork, cut into slices. These are roasted on a pan together with onions and butterfat. The dish is often seasoned with salt, pepper, marjoram, cumin, and parsley, and served with a fried egg or *kraut* (cabbage). Hence, gröstl is somewhat similar to the American dish called hash.

The letter 'ö' was replaced by 'ø', which is a letter in the Danish alphabet that is pronounced in the same way as 'ö'. This way, the name, like the hash function itself, contains a mix of Austrian and Danish influences.

The pronunciation of `Grøstl` may seem challenging. If you think so, then think of the letter 'ø' as the 'i' in "bird". This letter is a so-called *close-mid front rounded vowel*, and if you need more examples of its pronunciation, or a sound sample, check out [55].

The letter 'ø' may not appear on your keyboard. It can be written in a number of word processing environments as follows:

Environment	Command for 'ø'
LaTeX	`{\o}`
HTML	`ø` or `ø`
Windows	Alt + 0248
Linux	AltGr + o *

(* does not work in all settings.)

B   S-box

The s-box used in `Grøstl` is defined in Table 8.

Table 8: The `Grøstl` s-box (identical to the Rijndael/AES s-box). Given input x, find $x \wedge$ `f0` in the first column ('∧' is logical *and*), and find $x \wedge$ `0f` in the first row. Where the corresponding row and column meet, find the output $S(x)$ of the s-box.

	00	01	02	03	04	05	06	07	08	09	0a	0b	0c	0d	0e	0f
00	63	7c	77	7b	f2	6b	6f	c5	30	01	67	2b	fe	d7	ab	76
10	ca	82	c9	7d	fa	59	47	f0	ad	d4	a2	af	9c	a4	72	c0
20	b7	fd	93	26	36	3f	f7	cc	34	a5	e5	f1	71	d8	31	15
30	04	c7	23	c3	18	96	05	9a	07	12	80	e2	eb	27	b2	75
40	09	83	2c	1a	1b	6e	5a	a0	52	3b	d6	b3	29	e3	2f	84
50	53	d1	00	ed	20	fc	b1	5b	6a	cb	be	39	4a	4c	58	cf
60	d0	ef	aa	fb	43	4d	33	85	45	f9	02	7f	50	3c	9f	a8
70	51	a3	40	8f	92	9d	38	f5	bc	b6	da	21	10	ff	f3	d2
80	cd	0c	13	ec	5f	97	44	17	c4	a7	7e	3d	64	5d	19	73
90	60	81	4f	dc	22	2a	90	88	46	ee	b8	14	de	5e	0b	db
a0	e0	32	3a	0a	49	06	24	5c	c2	d3	ac	62	91	95	e4	79
b0	e7	c8	37	6d	8d	d5	4e	a9	6c	56	f4	ea	65	7a	ae	08
c0	ba	78	25	2e	1c	a6	b4	c6	e8	dd	74	1f	4b	bd	8b	8a
d0	70	3e	b5	66	48	03	f6	0e	61	35	57	b9	86	c1	1d	9e
e0	e1	f8	98	11	69	d9	8e	94	9b	1e	87	e9	ce	55	28	df
f0	8c	a1	89	0d	bf	e6	42	68	41	99	2d	0f	b0	54	bb	16