

Hash Function *Luffa*

Supporting Document

Christophe De Cannière

ESAT-COSIC, Katholieke Universiteit Leuven

Hisayoshi Sato, Dai Watanabe

Systems Development Laboratory, Hitachi, Ltd.

31 October 2008

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 3 |
| 2 | Design Rationale | 3 |
| 2.1 | Chaining | 3 |
| 2.2 | Non-Linear Permutation | 4 |
| 2.2.1 | Sbox in SubCrumb | 5 |
| 2.2.2 | MixWord | 6 |
| 2.2.3 | Constants | 7 |
| 2.2.4 | Number of Steps | 7 |
| 2.2.5 | Tweaks | 8 |
| 3 | Security Analysis of Permutation | 8 |
| 3.1 | Basic Properties | 8 |
| 3.1.1 | Sbox S | 8 |
| 3.1.2 | Differential Propagation | 10 |
| 3.2 | Collision Attack Based on A Differential Path | 11 |
| 3.2.1 | General Discussion | 11 |
| 3.2.2 | How to Find A Collision for 5 Steps without Tweaks | 13 |
| 3.3 | Birthday Problem on The Unbalanced Function | 13 |
| 4 | Security Analysis of Chaining | 14 |
| 4.1 | The Basic Properties of The Message Injection Functions | 15 |
| 4.2 | The First Naive Attack | 15 |
| 4.2.1 | Long Message Attack to Find An Inner Collision | 16 |
| 4.2.2 | How to Reduce The Message Length | 16 |
| 4.2.3 | The Efficiency of The Attack | 17 |
| 4.3 | Meet-In-The-Middle Attack on <i>Luffa</i> | 17 |
| 4.3.1 | Attack Description | 18 |
| 4.3.2 | The Complexity of The Attack | 18 |
| 4.4 | An Application of The Multicollision Attack on A Weakened MI | 19 |
| 4.4.1 | Attack Description for $w = 3$ | 20 |
| 4.4.2 | Extension to $w > 3$ | 20 |
| 4.5 | Collision, Second Preimage, And Preimage | 21 |
| 4.6 | Security Analysis of Finalization | 21 |
| 4.6.1 | Saturation Property of <i>Luffa</i> without A Blank Round | 21 |
| 4.6.2 | Slide Attack | 22 |
| 5 | Implementation Aspects | 22 |
| 5.1 | Performance Figures for Software Implementations | 22 |
| 5.1.1 | 8-bit Processors | 23 |
| 5.1.2 | 32-bit Processors | 24 |
| 5.1.3 | 64-bit Processor | 26 |
| 5.2 | Performance Figures for Hardware Implementations | 28 |
| 5.3 | Security against Side Channel Attacks | 30 |

1 Introduction

This document provides the currently known results on the security and the performances of a family of hash function *Luffa*. We refer to [22] for the specification and the notations used throughout this document.

The rest of this document is organized as follows: Firstly the design rationale of *Luffa* is presented in Section 2 to clarify that there is no backdoor in our design. Secondly the security issues of *Luffa* are discussed. The security of some randomness properties, especially the differential characteristics, of the non-linear permutation is discussed in Section 3. The security of the chaining is discussed in Section 4 by assuming that the underlying permutations are ideal function. After that the implementation issues are discussed in Section 5. Also the performances on some platforms are given in Section 5.

2 Design Rationale

Luffa is a family of cryptographic hash functions suitable for multipurpose. Besides NIST's requirements, we would like to make the design as simple as possible.

2.1 Chaining

The chaining method of *Luffa* hash function is a variant of a sponge function [5, 6] whose security is based only on the randomness of the underlying permutation. Although the simple construction is very attractive, a sponge function costs more than a traditional block cipher based hash construction from the viewpoint of implementations. In the case of PGV construction, the block length n_b of the underlying block cipher is equal to the hash length n_h (A key scheduling function should be taken into account in practice). On the other hand, a sponge function requires a permutation of $n_b \geq 2 \cdot n_h + m$ bits length, where m is the message block length. Generally speaking, the calculation cost of a function increases if the input (and the output) length get larger. In addition, constructing proper large permutations for each hash length is far from a scalable design.

These undesirable properties of a naive sponge lead us to its variant such

that the underlying function is a family of sub-permutations with shorter input/output. The resulting chaining method *Luffa* employs the message injection function and one blank round.

The message injection functions of *Luffa* are linear functions. These functions are designed to avoid the serious attack described in Section 4.4 caused by the simple property of keeping the independency of input/output of each sub-permutations. More precisely, the message injection functions are designed to have the maximum branch number, thus the input of each round is sufficiently mixed and the independency is kept to the minimum. Consequently, it is conceivable that these components make the construction *Luffa* sufficiently intractable to find a collision. Moreover we believe that these components also provide sufficient resistance against the preimage/second preimage attacks.

The finalization of *Luffa* consists of one blank round and an output function. One blank round is appended in order to be resistant against unknown future attacks. Note that if the number of padded message blocks is one, then the blank round is omitted for the efficiency.

The block length of the sub-permutations is chosen to be suitable for 256, 384 and 512 bits hash length. It is possible to choose the shorter block length, however all sub-permutations must be different each other so that too small block length is not practical. In addition, a bit slice permutation has more flexibility in the design if the block length is large. Therefore 256 bits is chosen as a block length of a sub-permutation.

2.2 Non-Linear Permutation

The sub-function of *Luffa* adopts a bit slice substitution permutation network (SPN). The reasons to choose an SPN more than a look up table (LUT) based function are as follows:

- If the CPU design is evolved, the throughput increases.
- A cipher consisting of logical operations is believed to be secure against cash timing attacks.
- It seems easier to implement compared to LUTs.

The first is the main reason to adopt a bit slice permutation. For example, four **SubCrumb** can be executed at once with the SSE instructions of Intel[®] 686 processors and the throughput is much faster than the code with 32-bit instructions only.

2.2.1 Sbox in SubCrumb

Serpent [2, 7] and Noekeon [10] are typical examples of a bit slice block cipher and they adopt Sboxes of 4 bits input/output.

In the proposal of Serpent, the designers explained the eight Sboxes are almost randomly generated and ones satisfying good differential/linear properties are chosen. In addition, those representation as Boolean functions should have the highest degree, namely three. The efficiency of the implementations is not well considered. Osvik proposed an efficient method to find a good sequence for a given Sbox on Intel[®] 586 processors [18]. In his experimental results, S_2 of Serpent is the fastest, the number of instructions is 16 and it can be executed in 8 cycles. Note that his method is not exhaustive, so that resultant sequence might not be optimal. In addition, how to choose the optimal Sbox in software implementation is not clarified in his method.

On the other hand, the Sbox of Noekeon is defined as a sequence of instructions by nature. Noekeon is intended to have symmetric property, i.e., the encryption and the decryption can be done by the same function. The reason to design the Sbox in this manner might be that the Sbox and the inverse must be equal. Unfortunately, the given set of instructions of Noekeon's Sbox is not suitable for software implementations.

Our approach to design the Sbox is similar to Noekeon. In other words, it is defined by the sequence of instructions to have a desirable property such that the implementation on Intel[®]Core[™]2 Duo processors is optimal. By simple thought experiment, we believe that at least five cycles are needed in order to achieve optimal differential/linear probability. In fact, we found a few Sboxes satisfying those properties. The smallest Sbox consists of 9 instructions and it is executable in 5 cycles.

The Sbox chosen for *Luffa* consists of 16 instructions and it is executable in 6 cycles. It seems to have some good properties as follows:

- The maximum differential probability is $1/4$.
- The maximum linear probability is $1/4$.
- It has no fixed point.
- The degree of the Boolean representation is 3 for all output bits.

2.2.2 MixWord

The linear diffusion layer of the sub-permutation of *Luffa* consists of XORings and rotations in the same way as Serpent and Noekeon. Those Sboxes do not mix bits at a different bit position in a word, but mix bits at a same bit position in different words, so that the linear diffusions are intended to mix the bits at different bit positions.

Different from those 128-bit block ciphers, the linear diffusion of *Luffa* is required to mix the outputs of different **SubCrumb**, i.e., x_k and x_{k+4} in addition. In order to achieve these requirements, a traditional Feistel ladder with rotations is chosen, where the inputs are x_k and x_{k+4} . The rotations are removed from the runcles of the ladder to execute an XORing and a rotation in parallel in software implementations. The number of iterations is chosen to be four because we believe that the weight of the non-linear mixing should be nearly equal to that of the linear diffusion.

Now we are going to explain how to choose the rotations σ_1 , σ_2 , σ_3 and σ_4 . Firstly, the linear code given by the iteration of **MixWord** is considered. Let A be the representation matrix of **MixWord** and $G_n = I||A||\cdots||A^n$ be the generator matrix of a code. We searched for the lowest weight code word generated from low weight inputs. Most of the candidates seems to have the same diffusion property for $n \leq 4$, and the chosen parameter shows the best property.

Next, the polynomial representation of **MixWord** is considered. An XORing and a rotation can be represented as the operations on a polynomial ring $\text{GF}(2)[t]$. Let the two input words be $a(t)$ and $b(t)$, then the output is given by

$$\begin{aligned} a'(t) &= (1 + t^{\sigma_1} + t^{\sigma_2} + t^{\sigma_3} + t^{\sigma_1+\sigma_3})a(t) + (1 + t^{\sigma_2} + t^{\sigma_3})b(t), \\ b'(t) &= t^{\sigma_4}(1 + t^{\sigma_1} + t^{\sigma_2})a(t) + t^{\sigma_4}(1 + t^{\sigma_2})b(t). \end{aligned}$$

It is clear from the equations that the branch number of `MixWord` is upper-bounded by six (In practice, it is five).

The notable input is $(0, b)$ and the corresponding output becomes $(a', 0)$ if $b(x)(1 + x^{\sigma_2}) = 0$, i.e. b should have a zero divisor. Such a differential path with small number of active Sboxes is not desirable from the security viewpoint. The lowest Hamming weight of such inputs b is given by $32/p$, where $\sigma_2 = 2^p(2\tau+1)$. In addition, the Hamming weight of the corresponding output $(a', 0)$ is also $32/p$. This fact indicates that the factor of 2 in σ_2 should be small.

On the other hand, we restricted the parameters σ_1 , σ_2 and σ_3 to be even in order to make it feasible to search for the best (truncated) differential path for 4 steps. Then a step of `MixWord` does not mix odd bits and even bits. In other words, the transformation can be separated into two independent functions. Note that the above choice might reduce the diffusion property in general. σ_4 should be odd in order to mix odd bits and even bits at the next step.

Throughout the above process, a set of parameters $\sigma_1 = 2, \sigma_2 = 14, \sigma_3 = 10, \sigma_4 = 1$ is chosen for *Luffa*.

2.2.3 Constants

The step constants are used for several reasons as follows:

- The sub-permutation Q_j MUST be different each other.
- The step function at each step SHOULD be different each other to prevent slide attacks and to remove fixed points.
- There SHOULD not be a kind of symmetry in the input/output.

From the viewpoint of implementations, it is better to generate the constants by a simple circuit with a fixed starting variable. Therefore we designed a small constant generator with a linear update function for *Luffa*.

2.2.4 Number of Steps

In order to achieve good throughput comparable to SHA-256, the number of steps is fixed to 8.

On the other hand, we found the differential path for 8 steps with probability 2^{-224} , and proved that there is no differential path with probability more than 2^{-124} . We do not think it is necessary for the sub-permutation Q_j to achieve full security (namely, the maximum differential probability should be not more than 2^{-256} , etc.). The collision attack based on a differential path is discussed in Section 3.2.

2.2.5 Tweaks

A tweak is applied in order to differentiate each permutation. And we also expect that the tweak will break a nature of the message injection function defined over a direct product ring. Especially, it will become hard to choose same input differences for the different permutations. The operations are chosen not to increase the implementation cost in terms of both the size and the performance.

3 Security Analysis of Permutation

This section shows some known properties of the non-linear permutation Q_j .

3.1 Basic Properties

3.1.1 Sbox S

Table 1 shows the differential probabilities corresponding to input and output differences. The maximum differential probability of the Sbox S is 2^{-2} .

Table 2 shows the biases of the linear approximation defined by corresponding input and output masks. The maximum linear probability of the Sbox S is 2^{-2} .

Let x_0, x_1, x_2, x_3 and y_0, y_1, y_2, y_3 be the 4-bit input and output of the Sbox. Then the algebraic normal form of the Sbox is given by

$$\begin{aligned} y_0 &= 1 + x_2 + x_0x_1 + x_1x_3 + x_2x_3 + x_0x_1x_3, \\ y_1 &= 1 + x_0 + x_2 + x_3 + x_0x_1 + x_0x_2 + x_1x_3 + x_2x_3 + x_0x_1x_3, \\ y_2 &= 1 + x_1 + x_1x_3 + x_2x_3 + x_0x_1x_3, \\ y_3 &= x_0 + x_1 + x_2 + x_0x_1 + x_1x_2 + x_1x_3 + x_0x_1x_2. \end{aligned}$$

Table 1: The differential profile of the Sbox S

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | a | b | c | d | e | f |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 2 | 0 | 0 | 0 | 0 | 2 | 0 | 4 | 0 | 4 | 2 | 2 | 0 | 0 | 0 |
| 2 | 0 | 0 | 2 | 2 | 0 | 0 | 4 | 0 | 0 | 0 | 2 | 2 | 0 | 0 | 4 |
| 3 | 0 | 0 | 2 | 0 | 2 | 2 | 2 | 0 | 2 | 0 | 0 | 2 | 2 | 0 | 2 |
| 4 | 0 | 0 | 2 | 2 | 0 | 0 | 0 | 0 | 4 | 0 | 2 | 2 | 0 | 4 | 0 |
| 5 | 2 | 4 | 2 | 2 | 0 | 2 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 6 | 4 | 0 | 0 | 0 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 4 | 4 |
| 7 | 0 | 0 | 0 | 2 | 2 | 2 | 2 | 0 | 2 | 0 | 2 | 0 | 2 | 0 | 2 |
| 8 | 0 | 4 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 4 | 0 | 0 | 4 | 0 | 0 |
| 9 | 2 | 0 | 0 | 0 | 0 | 2 | 0 | 4 | 0 | 0 | 2 | 2 | 4 | 0 | 0 |
| a | 4 | 2 | 0 | 0 | 2 | 0 | 0 | 0 | 0 | 2 | 0 | 0 | 2 | 4 | 0 |
| b | 0 | 2 | 0 | 2 | 0 | 2 | 2 | 0 | 2 | 2 | 2 | 0 | 0 | 0 | 2 |
| c | 0 | 0 | 2 | 2 | 0 | 0 | 0 | 0 | 4 | 0 | 2 | 2 | 0 | 4 | 0 |
| d | 2 | 0 | 2 | 2 | 4 | 2 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| e | 0 | 2 | 2 | 2 | 2 | 0 | 0 | 0 | 0 | 2 | 2 | 2 | 2 | 0 | 0 |
| f | 0 | 2 | 2 | 0 | 0 | 2 | 2 | 0 | 2 | 2 | 0 | 2 | 0 | 0 | 2 |

Note that the number of monomials which appears in the polynomial representation is smaller than that of a randomly generated Sbox. Though one might claim that this Sbox is weak in terms of algebraic attacks, we have not

Table 2: The linear profile of the Sbox S

| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | a | b | c | d | e | f |
|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 4 | 4 | 4 | -4 | 0 | 0 |
| 2 | 4 | -4 | 0 | 4 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 3 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 4 | 0 | -4 | 0 | -4 | 0 | -4 |
| 4 | 2 | 2 | 0 | 2 | -4 | 0 | -2 | 0 | -2 | -2 | 0 | 2 | 0 | -4 | -2 |
| 5 | 2 | 2 | 0 | 2 | 4 | 0 | -2 | -4 | 2 | -2 | 0 | 2 | 0 | 0 | 2 |
| 6 | 2 | 2 | 0 | 2 | -4 | 0 | -2 | 0 | 2 | 2 | 0 | -2 | 0 | 4 | 2 |
| 7 | -2 | -2 | 0 | -2 | -4 | 0 | 2 | -4 | 2 | -2 | 0 | 2 | 0 | 0 | 2 |
| 8 | 0 | 0 | -4 | 0 | 0 | 4 | 0 | 2 | 2 | 2 | -2 | 2 | 2 | -2 | 2 |
| 9 | 0 | 0 | -4 | 0 | 0 | -4 | 0 | 2 | 2 | -2 | 2 | -2 | -2 | -2 | 2 |
| a | -4 | 0 | 0 | 4 | 0 | 0 | 0 | -2 | -2 | 2 | -2 | -2 | -2 | -2 | 2 |
| b | 0 | 4 | 0 | 0 | 0 | 0 | 4 | -2 | 2 | 2 | 2 | -2 | 2 | -2 | -2 |
| c | 2 | 2 | 4 | -2 | 0 | 0 | 2 | 2 | 0 | 0 | -2 | 0 | -2 | -2 | 4 |
| d | 2 | 2 | -4 | -2 | 0 | 0 | 2 | -2 | -4 | 0 | -2 | 0 | -2 | 2 | 0 |
| e | 2 | -2 | 0 | -2 | 0 | -4 | -2 | -2 | 0 | 4 | -2 | 0 | 2 | -2 | 0 |
| f | -2 | 2 | 0 | 2 | 0 | -4 | 2 | 2 | 0 | 0 | -2 | 4 | 2 | 2 | 0 |

found any practical attack on *Luffa* using this property.

3.1.2 Differential Propagation

It is easy to see that the branch number of *MixWord* is 5. We confirmed that the minimum number of active Sboxes of 4 steps is 31. Therefore the maximum differential characteristic probability (MDCP) of Q_j is upper bounded by 2^{-124} . The MDCP of 4 steps is estimated by the exhaustive 4-bit-wise truncated differential path search, but the direct path search for more than 4 steps is computationally infeasible.

We also searched for the good differential path for 8 steps to get close to the real bound. We consider only the case that the output differences of Sboxes at different positions in the same step are equal. Under this assumption, the search for a good differential path of the permutation Q_j is equivalent to search for a low weight code word of the linear code defined by the iteration of *MixWord* (See Section 2.2.2 in detail for the definition of this code). We applied Leon's probabilistic algorithm [17] to find a low weight code word and Table 3 shows the best differential path so for.

Table 3: The best known (truncated) differential path

| | | | |
|---|----------------------------------|----------------------------------|----|
| 0 | 00000000000100000000111101010100 | 00000100000101000100010100010100 | 16 |
| 1 | 00000111000010000011011000000000 | 00101101101010001110110000100010 | 22 |
| 2 | 0000000000000000000000001010000 | 00000000000100010001000100010100 | 8 |
| 3 | 00000000000000000000000000000000 | 10001000100000000010000000000000 | 4 |
| 4 | 10000000000000000000000000000010 | 00000001000000000000010001000001 | 6 |
| 5 | 10000000000000011100001000001101 | 0000000001000011100100000010011 | 16 |
| 6 | 0001111010101000000000000100000 | 00000101010000010100010001000001 | 16 |
| 7 | 00010101010000101000011110011110 | 01100000101000100000010000110110 | 24 |

The most left column in the table means the step so that r -th line (counting from 0) means the 4-bit-wise truncated differences which is input to the $r + 1$ -th step function. The most right column means the number of active Sboxes at each step. The above path has 112 active Sboxes in total so that the differential characteristic probability of the path is 2^{-224} . It should be remarked that 8 step functions cannot be considered a perfect random permutation.

Table 4: The best known differential probabilities

| Number of steps | Diff. probability |
|-----------------|-------------------|
| 4 | $\leq 2^{-62}$ |
| 5 | 2^{-100} |
| 6 | 2^{-144} |
| 7 | 2^{-176} |
| 8 | 2^{-224} |

The best known differential characteristic probabilities up to 8 steps are summarized in Table 4. Note that the best known differential paths for more than 4 steps is a part of the 8 steps differential path. The line 2-6, line 1-6, 0-6 in Table 3 correspond to the probability for 5, 6, 7 steps respectively in Table 4.

3.2 Collision Attack Based on A Differential Path

Here we discuss a differential based collision attack consisting of two message blocks. In other word, how to find a message pair $(M^{(i)}, M^{(i+1)})$ and $(M^{(i)} \oplus \Delta^{(i)}, M^{(i+1)} \oplus \Delta^{(i+1)})$ such that $MI(\text{Round}(H^{(i-1)}, M^{(i)}), M^{(i+1)}) = MI(\text{Round}(H^{(i-1)}, M^{(i)} \oplus \Delta^{(i)}), M^{(i+1)} \oplus \Delta^{(i+1)})$ and the cost are discussed.

3.2.1 General Discussion

Let p_j be the maximum differential characteristic probability of Q_j and $p = \max_j p_j$. If some input bits are chosen adequately, the differential probability with the condition tends to be higher than p . This technique is well-known as a message modification in general. In order to simplify the discussion, we assume that a bit constraint of an input improves the differential probability double, i.e., the differential characteristic probability under m bits constraints is assumed to be $2^m \cdot p$. Note that this is hard to happen in practice because some of the constraints are often conflictive with the others.

The attack with messages of two rounds requires all the output differences of Q_j are not zero. Algorithm 1 describe the procedure of the attack.

Let H_{pre} , M_{pre} and M_{online} be the numbers of messages to be used in Step 1, 2 and 3 respectively. Let H_{online} be the number of iterations of

Algorithm 1 Differential based collision attack

Step 1 Choose a good internal state $H^{(i-1)}$ (which can satisfy constraints as many as possible) by moving $M^{(i-1)}$ randomly.

Step 2 Choose a part of the message block $M^{(i)}$ to satisfy constraints.

Step 3 Move the rest of the message block $M^{(i)}$ and check if the output differences are equal. If not, go back to Step 1.

the whole procedure. The differential probability under the constraints is roughly given by $p^w \cdot H_{pre} \cdot M_{pre}$ and the total number of trials is given by $H_{online} \cdot M_{online}$. Then an inner collision will be found if the following inequality is satisfied:

$$p^w \cdot (H_{pre} \cdot M_{pre})(H_{online} \cdot M_{online}) \geq 1.$$

The calculation complexity of the attack is given by $H_{online}(H_{pre} + M_{online})$ and it should be smaller than $2^{\frac{w-1}{4}n_b}$ for the attack being faster than a birthday attack to find a collision of outputs. Therefore $H_{pre} \cdot H_{online} < H_{online}(H_{pre} + M_{online}) \leq 2^{\frac{w-1}{4}n_b}$ is the necessary condition. In addition, $M_{pre} \cdot M_{online}$ is upper-bounded by 2^{n_b} . Therefore the lower bound of the maximum differential characteristic probability of Q_j (for a successful attack) is given by

$$p \geq 2^{-\frac{n_b}{4}(1+\frac{3}{w})}.$$

For $w = 3, 4, 5$, $p \geq 2^{-128}$, 2^{-112} , $2^{-102.4}$ are the necessary conditions respectively.

Finding an inner collision should be also considered because an inner collision can be used to find a second preimage and preimage as well as to find a collision for sponge variants. The discussion for the inner collision can be done in the same manner. In this case, the collision attack is successful if the number of operations is less than $2^{n_b(w-1)/2}$. Therefore $p \geq 2^{-170.7}$, 2^{-160} , $2^{-153.6}$ are the necessary conditions for $w = 3, 4, 5$ respectively.

The currently known best differential characteristic probability of Q_j is 2^{-224} and the tweaks will make it difficult to find a tuple of good differential paths, therefore we believe that *Luffa* is secure against this attack.

3.2.2 How to Find A Collision for 5 Steps without Tweaks

As an example to find an inner collision in practice, we describe a naive attack for 5 steps using the line 1-5 in Table 3. In order to simplify the discussion, we ignore the influences of the multiplications in the message injection function and the tweaks. In addition, we consider only *Luffa*-256, i.e., $w = 3$.

Before starting the attack, the bit-wise differences should be chosen and we can easily find an iterative sequence of differences ($0x1 \rightarrow 0xa \rightarrow 0x1 \rightarrow \dots$) from Table 1. A 4-bit data $a_{j,3,l}||a_{j,2,l}||a_{j,1,l}||a_{j,0,l}$ (or $a_{j,7,l}||a_{j,6,l}||a_{j,5,l}||a_{j,4,l}$) are called a *crumble* in the following. Let us consider the crumbles corresponding to ones (we call them active crumbles) in the difference at the line 1. If all the input bits to these Sboxes are chosen to be adequate, then the input differences $0x1$ are mapped to $0xa$ with probability 1. Each input to a crumble should hold two bits of constraints in order to satisfy the above condition. In Step 1, the attacker tries to find a state $H^{(i-1)}$ whose active crumbles satisfy constraints. The number of constraints here is $22 \cdot 2 \cdot 3 = 132$. In Step 2, the attacker chooses a message $M^{(i)}$ such that the above mentioned 4-bit crumbles are all zero. The number of constraints in this step is $22 \cdot 4 = 88$. Now the attacker has $256 - 88 = 168$ bits freedom in the message $M^{(i)}$. On the other hand, there are still $(8 + 4 + 6 + 16) \cdot 3 = 102$ active Sboxes, so that 2^{204} trials are needed to find an input pair which follows the differential path. Therefore $2^{204-168} = 2^{36}$ iterations of whole steps are needed. As a result, the total complexity of the attack is estimated at $2^{36}(2^{132} + 2^{168}) \approx 2^{204}$.

We expect that more detailed analysis such as the message modification at the next step certainly reduces the calculation complexity of the attack on 5 steps without tweaks, and may allow to attack 5 steps with tweaks or 6 steps. But the rough estimate in Section 3.2.1 tells that this kind of approach never reach to more than 6 steps. A multiple differential path search also helps to improve the attack, but we think the attack on 7 steps is hard to expect.

3.3 Birthday Problem on The Unbalanced Function

The standard birthday problem assumes that the underlying set is uniformly distributed. Bellare and Kohno discussed the birthday problem for unbal-

anced distributions and proved that the collision happens more often if the underlying distribution is not uniform [4]. The “non-randomness” of the permutation Q_j may tempt to apply their result to *Luffa*. Though the distribution of the outputs of the differential of Q_j is not uniform, to find a collision is equivalent to find a preimage of zero in the differential of Q_j . Therefore we believe that the application of the result on the unbalanced birthday problem is not possible.

4 Security Analysis of Chaining

In this section, the underlying permutations are assumed to be random and the security of the chaining of *Luffa* is discussed. We found a generic attack to find an inner collision which queries to the permutation Q_j only $2^{\frac{w-1}{w+1}n_b}$ times. Therefore we cannot claim that *Luffa* has sufficient security in terms of the current stream of security proof concerning only the number of queries to a random function. On the other hand, all of the attacks we considered require not less than $2^{\frac{w-1}{2}n_b}$ calculations. We believe that none of the attacks threatens the practical security of *Luffa*.

At the beginning of this section, the basic properties of the message injection functions are given.

After that, three generic attacks to find an inner collision are presented. Finding an inner collision of a sponge function yields a second preimage and preimage as well as a collision, same applies to *Luffa*. Therefore, *Luffa* is not secure if there is an attack to find an inner collision whose calculation complexity is less than $2^{\frac{w-1}{2}n_b}$.

All attacks presented in this section concern an inner collision, but the conditions assumed for the message injection functions MI are different. The first attack does not require any property to the message injection function more than a surjectivity. The second attack is adjusted to *Luffa*. The third attack is applicable only to very simple message injection function, and does not threaten the security of *Luffa* itself. However, this attack illuminate the necessary condition for the message injection function so that we also describe the attack. The expected number of queries and the computational complexities required for the attacks are summarized in Table 5.

In the last of this section, the security of the finalization process is dis-

Table 5: The complexity of generic attacks

| MI | Num. of queries (exponent of 2^a) | Calc. complexities (exponent of 2^a) | Section |
|--------------|---|--|---------|
| Any | $\frac{w-1}{w}n_b$ | $\frac{w-1}{2}n_b$ | 4.2 |
| <i>Luffa</i> | $\frac{w-1}{w+1}n_b$ | $\geq \frac{w-1}{2}n_b$ | 4.3 |
| XORings | $\frac{n_b}{2}$ | $\frac{n_b}{2}$ | 4.4 |

cussed.

4.1 The Basic Properties of The Message Injection Functions

The message injection functions of *Luffa* are defined over a direct product ring $\text{GF}(2^8)^{32}$. They have branch numbers 4, 5, 6 for $w = 3, 4, 5$ respectively. On the other hand, they have a kind of “non-diffusing” property as follows. Let a be any 32-bit data and n be a non-negative integer less than 32. A 32-bit data concatenating $32 - n$ bits consecutive zeros and the least significant n bits of a is denoted by $0||a[n]$. For a 256-bit data $X = (x_0, \dots, x_7)$, $0||X[n]$ is defined by $(0||x_0[n], \dots, 0||x_7[n])$. Then the message injection function MI satisfies $LH_j(0||H^{(i-1)}[n]) \oplus LM_j(0||M^{(i)}[n]) = 0||X[n]$ for all j and n , where X is an n_b bit data. This property is used in the attack described in Section 4.3.

4.2 The First Naive Attack

The first attack does not require any special property to the message injection function. The attack requires about $2^{\frac{w-1}{w}n_b}$ queries to each permutation Q_j , and also requires to estimate $2^{\frac{w-1}{2}n_b}$ intermediate states. Nevertheless the number of required queries is small, the latter calculations seem dominant so that we believe that this attack does not threaten the practical security of *Luffa*.

4.2.1 Long Message Attack to Find An Inner Collision

Let us denote the transformation by the message injection function MI by

$$X_j = LH_j(H^{(i-1)}) \oplus LM_j(M^{(i)}), \quad 0 \leq j < w.$$

If LM_j are surjective, then there are subspaces V_j over $\text{GF}(2)$ of dimension $\lceil \frac{w-1}{w}n_b \rceil$ such that for any state $H^{(i-1)}$ there is a message $M^{(i)}$ such that $X_j \in V_j$ for all j .

At the pre-computation phase, the attacker queries elements of V_j to Q_j , then the queries and the corresponding answers are stored. After that, he chooses an adequate message block $M^{(1)}$ such that $LH_j(H^{(0)}) \oplus LM_j(M^{(1)}) \in V_j$ for $0 \leq j < w$, then he accesses to the storage in order to get the next state $H^{(1)}$. Iterations of this process generates amount of intermediate states without an extra query to Q_j . An inner collision will be found if the number of intermediate states becomes more than $(2^{\lceil \frac{w-1}{w}n_b \rceil})^{\frac{w}{2}} \approx 2^{\frac{w-1}{2}n_b}$ because all the outputs of the message injection function MI are included in $\prod_{j=0}^{w-1} V_j$. In terms of the number of queries, the attack requires $2^{\lceil \frac{w-1}{w}n_b \rceil}$ queries to each permutation Q_j .

4.2.2 How to Reduce The Message Length

The attack in Section 4.2.1 requires a message of about $2^{\frac{w-1}{2}n_b}$ block length. However, an attack applicable to shorter messages is more attractive in practice and the message length is upper bounded by 2^{64} bits for *Luffa-224* and *-256*, and is upper bounded by 2^{128} bits for *Luffa-384* and *-512* respectively. Here we show how to reduce the message length drastically with a negligibly small additional cost.

Assume that the attacker allows V_0 to be 1 dimension larger. Then it is possible to find two messages such that $LH_j(H^{(i-1)}) \oplus HM_j(M^{(i)}) \in V_j$ for $0 \leq j < w$. This approach allows the attacker to construct a binary tree of intermediate states instead of a long sequence. If the depth of the tree becomes larger than $\frac{w-1}{2}n_b$, i.e., the length of any chain in the tree is not less than $\frac{w-1}{2}n_b$, then it includes $2^{\frac{w-1}{2}n_b}$ states.

Note that this idea is applicable without an additional cost in the case of *Luffa-256* and *Luffa-512*, because $n_b = 256$ is not divisible by $w = 3$ nor 5.

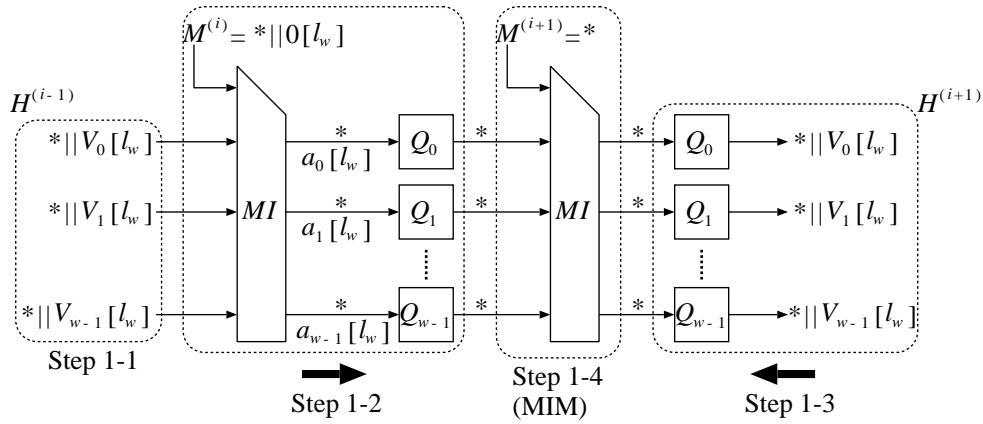


Figure 1: The first step of the meet-in-the-middle attack

4.2.3 The Efficiency of The Attack

The presented attack requires a huge memory of size $w \cdot 2^{\frac{w-1}{w} n_b}$ and accesses to the memory $w \cdot 2^{\frac{w-1}{2} n_b}$ times. If a memory access is much faster than a query to the permutation Q_j , the attack can violate the security claimed for *Luffa*. However, a memory access is very slow in general and a faster memory access such as a cache memory is very expensive. In addition, the required time to reach an objective memory address increases as the size of the memory gets larger. On the other hand, only 500 instructions are necessary to calculate the output of the permutation Q_j . We believe an access to such a huge memory is not less costly than a direct calculation of the permutation.

4.3 Meet-In-The-Middle Attack on *Luffa*

Next, a meet-in-the-middle (MIM) attack using a kind of non-diffusion property of the message injection function of *Luffa* is presented. The attack requires about $2^{\frac{w-1}{w+1} n_b}$ queries to each permutation Q_j , and also requires to apply the MIM $(2^{\frac{w-1}{w+1} n_b})^{\frac{w-1}{2}}$ times. The total calculation complexity of the second attack is considered not less than $2^{\frac{w-1}{2} n_b}$.

Algorithm 2 Meet-in-the-middle attack for *Luffa***Step 1: MIM to find a connection****1-1 (Upper side):** Fix a state $H^{(i-1)} = (*||V_0[l_w], \dots, *||V_{w-1}[l_w])$.**1-2 (Upper side):** Move $M^{(i)} = *||0[l_w]$ and query $M^{(i)}$ to each Q_j to generate inputs of *MI* at $i + 1$ -th round.**1-3 (Bottom):** For $0 \leq j < w$, move $*||V_j[l_w]$ and query to Q_j^{-1} to generate outputs of *MI* at $i + 1$ -th round.**1-4 (MIM):** Move $M^{(i+1)}$ and find an inner collision such that $Q_j(LH_j(H^{(i-1)}) \oplus LM_j(M^{(i)})) = Y_j$ for all j .**Step 2: Long message attack****2-1:** Set $i = 1$.**2-2:** Set $H^{(0)} = (V_0, \dots, V_{w-1})$.**repeat****2-3:** Apply the MIM in Step 1 at round i . Let $H^{(i+1)} = (Y_0, \dots, Y_{w-1})$ be the resultant state.**2-3:** Choose $\tilde{M}^{(i+2)} = *||0[l_w]$ such that $LH_0(H^{(i+1)}) \oplus LM_0(\tilde{M}^{(i+2)}) = V_0$. Note that this state $S^{(i+1)} = MI(H^{(i+1)}, \tilde{M}^{(i+2)})$ is a “sprig” in the chain.**2-4:** $i = i + 2$.**until** An inner collision happens in $\{S^{(i+1)}\}_i$ **4.3.1 Attack Description**

Here we use the notations defined in Section 4.1. In addition, the concatenation of any $32 - n$ bits and the least significant n bits of a is denoted by $*||a[n]$ and $*||X[n]$ is also defined in the similar manner.

Let l_w be an integer $n_b - \lceil \frac{w-1}{w+1} n_b \rceil$. Then the procedure of the attack is given in Algorithm 2 and Figure 1 shows the step 1 of the attack. It describes a long message attack, but a modification for the shorter message is possible in the similar manner to the first attack.

4.3.2 The Complexity of The Attack

Let q be the maximum number of queries to Q_j in Step 1-2 and 1-3. A inner collision is found if $q \cdot q^w \cdot 2^{n_b} \geq (2^{n_b})^{\frac{w}{2}}$, so that $q \geq 2^{\frac{w-1}{w+1} n_b}$ is necessary.

Step 2 is almost same as the long message attack presented in Section 4.2. Only the difference is to use the states generated in Step 2-3. The states are

of the form $(V_0, *||X_1[l_w], \dots, *||X_{w-1}[l_w])$, where $X_j[l_w]$ are constants, so that $(2^{\lceil \frac{w-1}{w+1} n_b \rceil})^{\frac{w-1}{2}}$ states are necessary to find an inner collision and it is still smaller than $2^{\frac{w-1}{2} n_b}$.

On the other hand, the total complexity of the attack becomes larger because the calculation complexity of Step 1 is not negligible any more. Let \mathcal{H} be the set consisting of inputs to MI in Step 1, \mathcal{X}_j be the set of the j -th output blocks. The calculation complexity to find a collision such that $LH_j(H^{(i-1)}) \oplus LM_j(M^{(i)}) = X_j$ for $0 \leq j < w$ is not less than the maximum size of the sets \mathcal{H} and \mathcal{X}_j because \mathcal{H} , \mathcal{X}_j are random sets. Therefore the complexity of Step 1 is at least $2^{\frac{w-1}{w+1} n_b}$.

The following algorithm shows a natural approach to find an inner collision in Step 1 and it requires $2^{\frac{2(w-1)}{w+1} n_b}$ calculations.

Algorithm 3 How to find a inner collision in Step 1-4

1-4-1: Choose any $H^{(i-1)} \in \mathcal{H}$ and $X_0 \in \mathcal{X}_0$.

1-4-2: Choose a message $M^{(i)} \in \mathcal{M}$ such that $X_0 = LH_0(H^{(i-1)}) \oplus LM_0(M^{(i)})$.

1-4-3: Check if $LH_j(H^{(i-1)}) \oplus LM_j(M^{(i)}) \in \mathcal{X}_j$ for $1 \leq j < w$. If yes, output $H^{(i-1)}, M^{(i)}$. Otherwise go back to Step 1-4-1.

The above discussion concludes that the total complexity of the second attack is not less than $2^{\frac{w-1}{w+1} n_b} \cdot (2^{\lceil \frac{w-1}{w+1} n_b \rceil})^{\frac{w-1}{2}} \approx 2^{\frac{w-1}{2} n_b}$. Therefore this attack does not threaten the practical security of *Luffa*.

4.4 An Application of The Multicollision Attack on A Weakened MI

In this section, an generic attack more effective than that described in Section 4.2 and 4.3 is presented. The attack is applicable only to the chaining with weak message injection function MI and it is not applicable to *Luffa*. But this attack indicates the necessity of a strong mixing function for MI . The message injection function MI of the variant is defined by

$$X_j = H_j^{(i-1)} \oplus M^{(i)}, \quad 0 \leq j < w.$$

A remarkable point of this variant is that each line is totally independent from others until the finalization is applied. The basic idea of the attack is

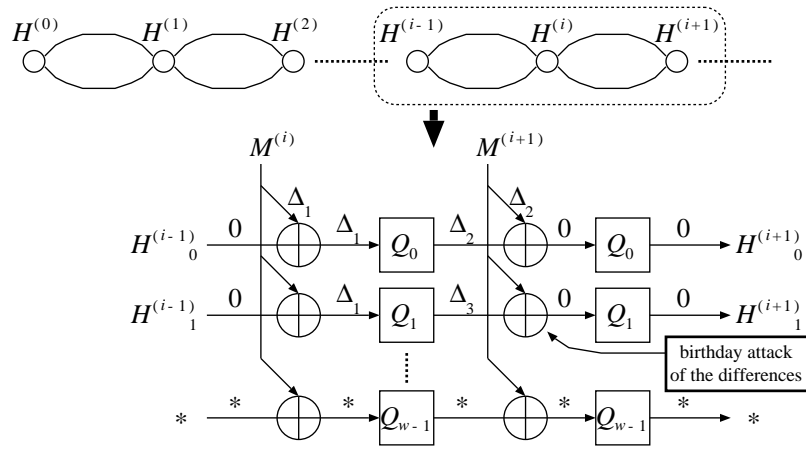


Figure 2: Partial collision chain for the weakened message injection

finding a collision for each block independently.

4.4.1 Attack Description for $w = 3$

In order to simplify the explanation, the width w is fixed to three. The applications in the case of the larger width are considered later.

The first step of the attack is finding a partial collision of first two blocks, consisting of two round inputs $M^{(i)}, M^{(i+1)}$, where the second message block $M^{(i+1)}$ is chosen to cancel the difference of $M^{(i)}$. After that a birthday attack by randomly chosen $M^{(i)}$ is applied to the second block. The calculation complexity of this step is approximately given by that of the birthday attack at the second block, so that it is about $2^{\frac{n_b}{2}}$ queries to the permutation Q_j .

The next step is constructing a partial collision chain of length $\frac{n_b}{2}$, i.e., the chain consists of n_b rounds. Then the attacker gets $2^{\frac{n_b}{2}}$ states whose first 2 blocks are all equal so that there will be a collision at the third block. This is an application of the multi-collision technique proposed by Joux [15].

4.4.2 Extension to $w > 3$

The vulnerability is caused by the fact that the intermediate values of each block can be independently calculated. This property holds even for $w > 3$,

so that recursive application of the multi-collision technique is possible. For example, the multi-collision chain consists of $(\frac{nb}{2})^2$ partial collisions allows to find a collision for $w = 4$. Both the number of queries to the permutation Q_j and the computational complexity are $n_b^w \cdot 2^{\frac{nb}{2}}$ for any w . The complexity of the attack increases if w becomes large, but very slowly.

4.5 Collision, Second Preimage, And Preimage

Bertoni *et al.* proved that the best attack on a sponge function is to find an inner collision [6]. And to find a collision of outputs, a second preimage, and a preimage all belong to the inner collision. We believe that it is also the case for *Luffa* and we have not found any serious attack to find an inner collision even though *Luffa* has no security proof so far. Therefore we think *Luffa* has the sufficient collision resistance, second preimage resistance, and preimage resistance.

4.6 Security Analysis of Finalization

4.6.1 Saturation Property of *Luffa* without A Blank Round

A saturation attack (or Square attack) was originally proposed by Daemen *et al.* as the dedicated attack on a block cipher SQUARE [9]. The basic idea of the attack is similar to truncated differential attack, but it uses another property preserved by a permutation mapping. Assume that the attacker takes all possible inputs and gets the outputs of the permutation. Then the sum of outputs becomes zero.

The message injection function MI and the consecutive permutations Q_j preserves this property. The output function OF also preserves the property. If there is no blank round and the attacker can take all values of the last message block, the sum of the outputs becomes zero. This is a kind of distinguisher and requires 2^{nb} calculations.

The first necessary condition for the attack is satisfied if the message has only a block. On the other hand, the message padding at the last block always prevents to take all values. Therefore we conclude that this property does not threaten the security of *Luffa*.

4.6.2 Slide Attack

A slide attack was originally proposed by Biryukov and Wagner [8] for block ciphers whose key scheduling is simple. Gorski *et al.* pointed out a slide attack is applicable under the keyed hashing context [13] and showed that it can break some sponge-like hash functions. Their attack is applicable if the round function is very thin and the finalization is not well considered. In the case of *Luffa*, the round function is much stronger than that of the broken sponge variants. In addition, the message padding and the fixed message inputs to a blank round are defined to avoid the sliding property. Therefore we believe that a kind of slide attacks does not threaten the security of *Luffa*.

5 Implementation Aspects

5.1 Performance Figures for Software Implementations

Here, we show the software performances of the *Luffa* hash family. The hash family was implemented on several processors; the Intel[®] Core[™]2 [14], the Atmel[®] AVR[®] [1], Renesas[®] H8[®] [19] and ARM[®] ARM9 [3] processors. We have evaluated the speed and memory usage of the hash family on the processors.

As for speed performance, we have measured two types of figures; execution time to hash a one-block message and execution time to hash a very long message. For the former, “one-block” means that the padded message consists only of a 256-bit block. That is, the length of the original message is no more than 255 bits. For the latter, the length of the message is set to be so large that a contribution of the finalization function to a throughput speed is negligible. More precisely, the figure is the execution time of an invocation of the round function divided by 32.

There are small differences between the throughputs for a short message and a very long message. However, it should be pointed out that the throughput becomes worst if the message length is exactly 256 bits. And it is about three times slower than when the message is very long, because *Luffa* calls its round function 3 times in this case. This feature will have an influence only for some applications such as an HMAC for very short messages.

Table 6: Execution time and memory requirements on AVR[®] ATmega8515

| Bit length of hash value | Execution time | | Memory requirements | |
|-----------------------------|------------------------------------|---------------------------------|---|----------------|
| | One-block msg. (cycles/message) | Very long msg. (cycles/byte) | Code size + constant data (bytes) | RAM (bytes) |
| 224 | 23,402 | 693.8 | 640 + 120 | 132 |
| 256 | 23,458 | 695.7 | 640 + 120 | 132 |
| 384 | 64,123 | 1,007.3 | 722 + 160 | 164 |
| 512 | 87,217 | 1,366.8 | 786 + 200 | 196 |

5.1.1 8-bit Processors

Luffa has been implemented for the Atmel[®] AVR[®] processor and the Renesas[®] H8[®] processor in assembly languages.

Atmel[®] AVR[®] Processor Our target processor model is ATmega8515. The processor has 8192 bytes of flash memory and 512 bytes of SRAM. We used Atmel[®]'s AVR Studio[®] as a development environment and measured execution time on the AVR Simulator of ATmega8515 bundled with the AVR Studio[®].

The execution time and memory requirements of an assembly code are shown in Table 6. In the table, the second column lists the execution time to hash a one-block message. The third column lists the execution time to hash a very long message. The fourth and fifth columns are for the sizes of the implementation. There is no setup of the algorithm in *Luffa*, hence the setup time is zero.

Renesas[®] H8[®] Processor Our target processor model is H8[®] 38024F (H8/300L core). The processor has 32 kbytes of flash memory and 1 kbytes of SRAM. We used Renesas[®]'s High-performance Embedded Workshop (HEW) as a development environment and measured execution time on the Simulator of H8/300L bundled with HEW.

The execution time and memory requirements of an assembly code are shown in Table 7. In the table, the second column lists the execution time to hash a one-block message. The third column lists the execution time to

Table 7: Execution time and memory requirements on Renesas[®]H8/300L

| Bit length of hash value | Execution time | | Memory requirements | |
|-----------------------------|------------------------------------|---------------------------------|---|----------------|
| | One-block msg. (cycles/message) | Very long msg. (cycles/byte) | Code size + constant data (bytes) | RAM (bytes) |
| 224 | 50,320 | 1,559.9 | 780 + 120 | 142 |
| 256 | 50,320 | 1,559.9 | 780 + 120 | 142 |
| 384 | 138,508 | 2,095.0 | 932 + 160 | 176 |
| 512 | 183,596 | 2,789.3 | 1,066 + 200 | 208 |

hash a very long message. The fourth and fifth columns are for the sizes of the implementation. There is no setup of the algorithm in *Luffa*, hence the setup time is zero.

5.1.2 32-bit Processors

Here, we will show some performance figures of *Luffa* for 32-bit processors. The hash family has been implemented for the Intel[®] Core[™]2 Duo processor in C and assembly languages and for the ARM[®] ARM9 processor in C language.

Intel[®] Core[™]2 Duo Processor Our target processor model is the Intel[®] Core[™]2 Duo E6600 2.4GHz processor in 32-bit mode, which will be used by NIST. We have measured speed performances of three different types of codes, a C code obeying the ANSI C grammar, a C code using Visual C++[®] SSE intrinsics and an assembly code. Both of the C codes obey the NIST API.

Table 8: 32-bit platforms used for measurement

| Programming language | Processor | Memory | OS | Compiler or assembler |
|----------------------|--|---------|---|-------------------------------------|
| C | Core [™] 2 Duo E6600 (2.4GHz) | 2GBytes | Windows Vista [®] 32-bit Edition | Visual Studio [®] 2005 C++ |
| Assembly | Core [™] 2 Duo E6600 (2.4GHz) | 2GBytes | Ubuntu [®] Linux [®] 8.04 32-bit distribution | gnu as |

Table 9: Throughput speed and execution time to hash a very long message on Core™2 Duo in 32-bit mode

| Bit length of hash value | ANSI C | | C using SSE intrinsics | | Assembly | |
|-----------------------------|-----------------------------|-----------------|-----------------------------|-----------------|-----------------------------|-----------------|
| | Exec. time (cycles/byte) | Speed (Mbps) | Exec. time (cycles/byte) | Speed (Mbps) | Exec. time (cycles/byte) | Speed (Mbps) |
| 224 | 33.9 | 565 | 19.1 | 1,001 | 13.9 | 1,381 |
| 256 | 33.4 | 574 | 19.2 | 995 | 13.9 | 1,381 |
| 384 | 45.2 | 424 | 22.2 | 864 | 15.7 | 1,218 |
| 512 | 59.7 | 321 | 35.1 | 546 | 25.5 | 752 |

Table 10: Execution time to hash a one-block message on Core™2 Duo in 32-bit mode

| Bit length of hash value | ANSI C | C using SSE intrinsics | Assembly |
|-----------------------------|------------------------------------|------------------------------------|------------------------------------|
| | Execution time (cycles/message) | Execution time (cycles/message) | Execution time (cycles/message) |
| 224 | 1,248 | 905 | 445 |
| 256 | 1,303 | 915 | 445 |
| 384 | 3,129 | 1,699 | 1,009 |
| 512 | 4,153 | 2,591 | 1,633 |

In accordance with the test environment of NIST, we have used the Microsoft®'s Visual Studio® 2005 C++ compiler and Windows Vista® Ultimate 32-bit operating system for the measurement of the C codes. The assembly code was measured on a 32-bit Linux® distribution. These environments are shown on Table 8.

We measured two types of figures, the execution time to hash a very long message and the execution time to hash a one-block message. The figures are shown at Table 9 and 10. In Table 9, a throughput speed is also listed. Note that the results of the C codes include overheads coming from the NIST API, but that of the assembly code does not. Also note that there is no setup of the algorithm in *Luffa*, hence the setup time is zero.

Since the fast implementation of *Luffa* does not use look-up tables, only small size of memory is required to implement *Luffa*. Therefore, a relation of time-memory trade-off in 32-bit implementation is relatively weak.

Table 11: Execution time on ARM[®] ARM926EJ-S[™] in C language

| Bit length of hash value | Execution time | |
|-----------------------------|------------------------------------|---------------------------------|
| | One-block msg. (cycles/message) | Very long msg. (cycles/byte) |
| 224 | 4,418 | 100.1 |
| 256 | 4,476 | 100.1 |
| 384 | 10,524 | 137.1 |
| 512 | 13,677 | 176.6 |

ARM[®] ARM9 Processor Our target processor model is ARM[®] ARM-926EJ-S[™]. We used ARM[®] RealView[®] Development Suite as a development environment and measured execution time on the Simulator of ARM926EJ-S[™]. The model ARM926EJ-S[™] has an instruction cache and a data cache. The size of each cache can be from 4 kbytes to 128 kbytes. In our measurement, the simulator is set to have 4 kbytes for each.

The execution time of an ANSI C code is shown in Table 11. This code is slightly different from the ANSI C code used on the Intel[®]Core[™]2 Duo processor. In the table, the second column lists the execution time to hash a one-block message. The third column lists the execution time to hash a very long message. Note that the result of the C code includes overheads coming from the NIST API. Also note that there is no setup of the algorithm in *Luffa*, hence the setup time is zero.

5.1.3 64-bit Processor

Here, we show some performance figures of *Luffa* for a 64-bit processor. The figures are for the Intel[®] Core[™]2 Duo processor, which will be used by NIST.

Intel[®] Core[™]2 Duo Processor Our target processor model is the Intel[®] Core[™]2 Duo E6600 2.4GHz Processor in 64-bit mode. We have measured speed performances of three different types of codes, an C code obeying the ANSI C grammar, a C code using Visual C++[®] SSE intrinsics and an assembly code. Both of the C codes obey the NIST API.

Table 12: 64-bit platforms used for measurement

| Programming language | Processor | Memory | OS | Compiler or assembler |
|----------------------|---|---------|---|-------------------------------------|
| C | Core TM 2 Duo E6600 (2.4GHz) | 2GBytes | Windows Vista [®] 64-bit Edition | Visual Studio [®] 2005 C++ |
| Assembly | Core TM 2 Duo E6600 (2.4GHz) | 2GBytes | Ubuntu [®] Linux [®] 8.04 64-bit distribution | gnu as |

Table 13: Throughput speed and execution time to hash a very long message on CoreTM2 Duo in 64-bit mode

| Bit length of hash value | ANSI C | | C using SSE intrinsics | | Assembly | |
|--------------------------|--------------------------|--------------|--------------------------|--------------|--------------------------|--------------|
| | Exec. time (cycles/byte) | Speed (Mbps) | Exec. time (cycles/byte) | Speed (Mbps) | Exec. time (cycles/byte) | Speed (Mbps) |
| 224 | 32.0 | 598 | 16.5 | 1,158 | 13.4 | 1,428 |
| 256 | 32.0 | 598 | 16.7 | 1,146 | 13.4 | 1,428 |
| 384 | 39.0 | 491 | 19.1 | 1,004 | 15.2 | 1,263 |
| 512 | 50.3 | 381 | 30.8 | 623 | 23.2 | 827 |

Table 14: Execution time to hash a one-block message on CoreTM2 Duo in 64-bit mode

| Bit length of hash value | ANSI C | C using SSE intrinsics | Assembly |
|--------------------------|---------------------------------|---------------------------------|---------------------------------|
| | Execution time (cycles/message) | Execution time (cycles/message) | Execution time (cycles/message) |
| 224 | 1,203 | 762 | 430 |
| 256 | 1,211 | 762 | 430 |
| 384 | 2,705 | 1,451 | 973 |
| 512 | 3,480 | 2,213 | 1,485 |

In accordance with the test environment of NIST, we have used the Microsoft[®]'s Visual Studio[®] 2005 C++ compiler and Windows Vista[®] Ultimate 64-bit operating system for the measurement of the C codes. The assembly code was measured on a 64-bit Linux[®] distribution. These environments are listed on Table 12.

We measured two types of figures, the execution time to hash a very long message and the execution time to hash a one-block message. The figures are shown at Table 13 and 14. In Table 13, a throughput speed is also listed.

Note that the results of the C codes include overheads coming from the NIST API, but that of the assembly code does not. Also note that there is no setup of the algorithm in *Luffa*, hence the setup time is zero.

Since the fast implementation of *Luffa* does not use look-up tables, only small size of memory is required to implement *Luffa*. Therefore, a relation of time-memory trade-off in 64-bit implementation is relatively weak.

5.2 Performance Figures for Hardware Implementations

The hardware performance evaluation of the *Luffa* hash family was done by synthesizing the proposed designs using 0.13 μm CMOS standard cell library. The code was first written in GEZEL [12] and tested for the functionality using the test vectors provided by the software implementations. The GEZEL code was then compiled to VHDL and synthesized using the Synopsys[®]Design Vision[™] tool [20]. The synthesis results are given in Table 15. Similar to the software implementations, the notation of “one-block message” is used when the padded message consists only of a single 256-bit block. The last column shows the throughput for hashing a very long message.

Our goal was to show that the family of cryptographic hash functions *Luffa* can be implemented efficiently in hardware. We targeted both, the compact and the high-throughput implementations. As can be seen from Table 15, the most compact implementation is obtained for *Luffa*-224/256 algorithm and consumes approximately 10 kGE. Note that our only goal for the compact implementation was to have a small die size, regardless of the circuit speed. Hence, we fixed the frequency to 100 MHz and synthesized our design. On the other hand, the high-throughput designs are synthesized regardless of the gate count and show that the *Luffa* hash algorithm can achieve the throughput of more than 12.5 Gbps.

The compact implementation was made using only one non-linear permutation block. Inside the permutation we used a single Sbox and a single MixWord block. This approach resulted in a large number of cycles (891), while on the other hand it efficiently reduced the final gate count. We used three 256-bit registers to maintain the internal state.

Table 15: Hardware evaluation of the *Luffa* hash family.

| Design | Area [GE] | Frequency [MHz] | # of cycles per round | Throughput [Mbps] | |
|------------------------------------|--------------|--------------------|--------------------------|----------------------|----------------------|
| | | | | One-block message | Very long message |
| <i>Luffa</i> -224/256 [†] | 10, 157 | 100 | 891 | 28.7 | 28.7 |
| <i>Luffa</i> -224/256 [‡] | 26, 849 | 444 | 9 | 12, 642.0 | 12, 642.0 |
| <i>Luffa</i> -384 [‡] | 34, 985 | 444 | 9 | 12, 642.0 | 12, 642.0 |
| <i>Luffa</i> -512 [‡] | 44, 163 | 444 | 9 | 12, 642.0 | 12, 642.0 |

[†] Compact implementation.

[‡] High-throughput implementations.

For the high-throughput implementations, the goal was to minimize the critical path. We used w permutation blocks in parallel and each of them contained 64 Sboxes and 4 MixWord blocks. The straightforward implementation resulted in the critical path of 3.05 ns and the cycle count of 8. The critical path was placed from the input of the message injection function to the output of the permutation block. As the message injection function is performed only once at the beginning of every round, we moved the state registers at the input of the permutation blocks. This resulted in the faster design, shortening the critical path to only 2.25 ns . One more clock cycle had to be spent in order to perform the complete round, but in total the final throughput got increased for about 20 %.

The throughput for “one-block” message as well as for the very long message was calculated according to the following equation:

$$\text{Throughput} = \frac{\text{Frequency}}{\# \text{ of Cycles}} \times 256 .$$

However, it should be pointed out that, for some special cases, the throughput drops down for about three times. For example, if the message length before padding is exactly 256 bits, then after the padding is done, the message will be composed of two 256-bit blocks. Now, the round function has to be executed once per each block and one more blank round needs to be performed at the end. This special case will have an influence only for some applications such as an HMAC for “one-block” messages.

Regarding the hardware implementation, one can further explore the different levels of parallelisms and make a plenty of trade-offs by trading the speed for the circuit size and vice versa. Especially challenging part remains the compact implementation of the hash functions in general and hence, we expect more research effort in that direction.

5.3 Security against Side Channel Attacks

A side channel attack observes physical information leakage in addition to the input and the output of the cipher in order to recover the secret data. There is no secret for a naive hash function so that it is not necessary to consider the threat of side channel attacks. However, some applications such as the HMAC use a secret information. In such applications, side channel attacks also should be considered. Hereinafter, we discuss the abstract property of the sub-permutations of *Luffa* against side channel attacks both in software and hardware implementations.

Tsunoo *et al.* proposed a cache timing attack for a software implementation of DES [21]. The attack observes the delay of the operation caused by the difference between cache hits and cache misses. This attack can be widely applicable to ciphers in which Sboxes are implemented by reference tables and amount of papers have been published to improve the attack and the countermeasure. Besides, a bit slice permutation does not refer the cache and implements the Sboxes by a set of logical instructions. This feature of a bit slice permutation avoids any kind of attacks based on cache timing.

A differential power analysis (DPA) observes the power consumption of certain part of operations depending on the secret information [16]. It is hard to perfectly avoid the DPA, instead amount of techniques to increase the cost to apply the attack have been proposed. The most likely approach is obscuring the power consumption of each operation. We believe that the cost to obscure the operations of *Luffa* is not costly in comparison to MD-like hash functions.

Trademarks

- ARM[®] and RealView[®] are registered trademarks and ARM926EJ-S[™] is a trademark of ARM Limited in the United States and/or other countries.
- Atmel[®], AVR[®], and AVR Studio[®] are registered trademarks of Atmel Corporation in the United States and/or other countries.
- Intel[®] is a registered trademark and Core[™] is a trademark of Intel Corporation in the U.S. and other countries.
- Linux[®] is a registered trademark of Linus Torvalds in the U.S. and other countries.
- Microsoft[®], Windows Vista[®], and Visual Studio[®] are registered trademarks of Microsoft Corporation in the United States and/or other countries.
- Renesas[®] and H8[®] are registered trademarks of Renesas Technology Corporation in the United States and/or other countries.
- Synopsys[®] is a registered trademark and Design Vision[™] is a trademark of Synopsys, Inc. in the United States and/or other countries.
- Ubuntu[®] is a registered trademark of Canonical Ltd. in the United States and/or other countries.

References

- [1] Atmel Corporation, “8-bit AVR Instruction Set,” available at http://www.atmel.com/dyn/resources/prod_documents/doc0856.pdf.
- [2] R. Anderson, E. Biham and L. Knudsen, “Serpent: A Proposal for the Advanced Encryption Standard,” available at <http://www.cl.cam.ac.uk/~rja14/serpent.html>.
- [3] Reference Manuals of the ARM[®] architectures and processors are available at <http://infocenter.arm.com/help/index.jsp>.

-
- [4] M. Bellare and T. Kohno, “Hash function balance and its impact on birthday attacks,” *Advances in Cryptology - Eurocrypt’04*, Lecture Notes in Computer Science, Vol. 3027, Springer-Verlag, pp. 401–418, 2004.
- [5] G. Bertoni, J. Daemen, M. Peeters and G. Van Assche, “Sponge Functions,” Ecrypt Hash Workshop 2007.
- [6] G. Bertoni, J. Daemen, M. Peeters and G. Van Assche, “On the Indifferentiability of the Sponge Construction,” *Advances in Cryptology, Eurocrypt’08*, Lecture Notes in Computer Science, Vol. 4965, Springer-Verlag, pp. 181–197, 2008.
- [7] E. Biham, R. Anderson and L. Knudsen, “Serpent: A New Block Cipher Proposal,” *Fast Software Encryption, FSE’97*, Lecture Notes in Computer Science, Vol. 1372, Springer-Verlag, pp. 222–238, 1998.
- [8] A. Biryukov and D. Wagner, “Slide Attacks,” *Fast Software Encryption, FSE’99*, Lecture Notes in Computer Science, Vol. 1636, Springer-Verlag, pp. 245–259, 1999.
- [9] J. Daemen, L. Knudsen, V. Rijmen, “The Block Cipher Square,” *Fast Software Encryption, FSE’97*, Lecture Notes in Computer Science, Vol. 1267, Springer-Verlag, pp. 149–165, 1997.
- [10] J. Daemen, M. Peeters, G. Van Assche and V. Rijmen, “Nessie Proposal: NOEKEON,” available at <http://gro.noekeon.org/>.
- [11] National Institute of Standards and Technology, “Secure Hash Standard,” FIPS 180-2.
- [12] GEZEL, http://rijndael.ece.vt.edu/gezel2/index.php/Main_Page.
- [13] M. Gorski, S. Lucks and T. Peyrin, “Slide Attacks on Hash Functions,” Cryptology ePrint Archive 2008/263, 2008.
- [14] Intel Corporation, “Intel[®] 64 and IA-32 Architectures Software Developer’s Manual,” available at <http://www.intel.com/products/processor/manuals/index.htm>.

- [15] A. Joux, “Multicollisions in iterated hash functions. Application to cascaded constructions,” *Advances in Cryptology, CRYPTO’04*, Lecture Notes in Computer Science, Vol. 3152, Springer-Verlag, pp. 306–316, 2004.
- [16] P. Kocher, J. Jaffe, and B. Jun, “Introduction to differential power analysis and related attacks,” 1998. Available at <http://www.cryptography.com/dpa/technical/index.html>.
- [17] J. S. Leon, “A Probabilistic Algorithm for Computing Minimum Weights of Large Error-Correcting Codes,” *IEEE Trans. on Information Theory*, Vol. 34, No. 5, pp. 1354–1359, 1988.
- [18] D. A. Osvik, “Speeding up Serpent,” *The 3rd AES Conference, Proceedings*, pp. 317–329, 2000.
- [19] Renesas Technology Corporation, “H8/300L Series Software Manual,” available at http://documentation.renesas.com/eng/products/mpumcu/rej09b0214_h83001.pdf.
- [20] Synopsis, <http://www.synopsys.com/>.
- [21] Y. Tsunoo, T. Saito, T. Suzaki, M. Shigeri, H. Miyauchi, “Cryptanalysis of DES Implemented on Computers with Cache,” *Cryptographic Hardware and Embedded Systems, CHES’03*, Lecture Notes in Computer Science, Vol. 2779, Springer-Verlag, pp. 62–76, 2003.
- [22] C. De Cannière, H. Sato, and D. Watanabe, “Hash Function *Luffa*, Specification,” NIST SHA3 competition, 2008.