Cryptanalysis of LASH

Ron Steinfeld[1], Scott Contini[1], Krystian Matusiewicz[1], Josef Pieprzyk[1], Jian Guo[2], San Ling[2], and Huaxiong Wang[1,2]

[1] Advanced Computing – Algorithms and Cryptography,
Department of Computing, Macquarie University
{rons,scontini,kmatus,josef,hwang}@ics.mq.edu.au
[2] Division of Mathematical Sciences,
School of Physical & Mathematical Sciences
Nanyang Technological University.
{guojian,lingsan,hxwang}@ntu.edu.sg

Abstract. We show that the LASH-x hash function is vulnerable to attacks that trade time for memory, including collision attacks as fast as $2^{(4x/11)}$ and preimage attacks as fast as $2^{(4x/7)}$. Moreover, we briefly mention heuristic lattice based collision attacks that use small memory but require very long messages that are expected to find collisions much faster than $2^{x/2}$. All of these attacks exploit the designers' choice of an all zero IV.

We then consider whether LASH can be patched simply by changing the IV. In this case, we show that LASH is vulnerable to a $2^{(7x/8)}$ preimage attack. We also show that LASH is trivially not a PRF when any subset of input bytes is used as a secret key. None of our attacks depend upon the particular contents of the LASH matrix – we only assume that the distribution of elements is more or less uniform.

Additionally, we show a generalized birthday attack on the final compression of LASH which requires $O\left(x2^{\frac{x}{2(1+\frac{107}{105})}}\right) \approx O(x2^{x/4})$ time and memory. Our method extends the Wagner algorithm to truncated sums, as is done in the final transform in LASH.

1 Introduction

The LASH hash function [2] is based upon the provable design of Goldreich, Goldwasser, and Halevi (GGH) [6], but changed in an attempt to make it closer to practical. The changes are:

1. Different parameters for the $m \times n$ matrix and the size of its elements to make it more efficient in both software and hardware.
2. The addition of a final transform [7] and a Miyaguchi-Preneel structure [9] in attempt to make it resistant to faster than generic attacks.

The LASH authors note that if one simply takes GGH and embeds it in a Merkle-Damgård structure using parameters that they want to use, then there are faster than generic attacks. More precisely, if the hash output is x bits, then they roughly describe attacks which are of order $2^{x/4}$ if n is larger than approximately m^2, or $2^{(7/24)x}$ otherwise[3]. These attacks require an amount of memory of the same order as the computation time. The authors hope that adding the second changes above prevent faster than generic attacks. The resulting proposals are called LASH-x, for LASH with an x bit output.

Although related to GGH, LASH is *not* a provable design: no security proof has been given for it. Both the changes of parameters from GGH and the addition of the Miyaguchi-Preneel structure and final transform prevent the GGH security proof from being applied.

Our Results. In this paper, we show:

- LASH-x is vulnerable to collision attacks which trade time for memory (Sect. 4). This breaks the LASH-x hash function in as little as $2^{(4/11)x}$ work (i.e. nearly a cube root attack). Using similar techniques, we can find preimages in $2^{(4/7)x}$ operations. These attacks exploit LASH's all zero IV, and thus can be avoided by a simple tweak to the algorithm.
- Even if the IV is changed, the function is still vulnerable to a short message (1 block) preimage attack that runs in time/memory $O(2^{(7/8)x})$ – faster than exhaustive search (Sect. 5). Our attack works for *any* IV.
- LASH is not a PRF (Sect. 3.1) when keyed through any subset of the input bytes. Although the LASH authors, like other designers of heuristic hash functions, only claimed security goals of collision resistance and preimage resistance, such functions are typically used for many other purposes [5] such as HMAC [1] which requires the PRF property.
- LASH's final compression (including final transform) can be attacked in $O\left(x2^{\frac{x}{2(1+\frac{107}{105})}}\right) \approx O(x2^{x/4})$ time and memory. To do this, we adapt Wagner's generalized birthday attack [12] to the case of truncated sums (Sect. 6). As far as we are aware, this is the fastest known attack on the final LASH compression.

We also explored collision attacks for very long messages using lattice reduction techniques; experiments for LASH-160 suggest that such attacks can find collisions for LASH-160 in significantly less than 2^{80} time

[3] The authors actually describe the attacks in terms of m and n. We choose to use x which is more descriptive.

and with very little memory. Due to lack of space, we could not include these results here – refer to Sect. 6.2 for a brief summary.

Before we begin, we would like to make a remark concerning the use of large memory. Traditionally in cryptanalysis, memory requirements have been mostly ignored in judging the effectiveness of an attack. However, recently some researchers have come to question whether this is fair [3, 4, 13]. To address this issue in the context of our results, we point out that the design of LASH is motivated by the assumption that GGH is insufficient due to attacks that use large memory and run faster than generic attacks [2]. We are simply showing that LASH is also vulnerable to such attacks so the authors did not achieve what motivated them to change GGH. We also remark that a somewhat more efficient cost-based analysis is possible [11], but page limits prevent us from providing the analysis here.

After doing this work, we have learnt that a collision attack on the LASH compression function was sketched at the Second NIST Hash Workshop [8]. The attack applies to a certain class of circulant matrices. However, after discussions with the authors [10], we determined that the four concrete proposals of x equal to 160, 256, 384, and 512 are not in this class (although certain other values of x are). Furthermore, the attack is on the compression function only, and does not seem to extend to the full hash function.

We remark that our zero IV attacks apply also to the FFT hash function[4] [8] if it were to be used in Merkle-Damgård mode with a zero IV, giving collisions/preimages with complexity $O(2^{x/3})/O(2^{x/2})$, even if the internal state is longer than output length x. However, our preimage attack for non zero IV would not apply due to the prime modulus.

2 Description of LASH

2.1 Notation

Let us define $\mathsf{rep}(\cdot) : \mathbb{Z}_{256} \rightarrow \mathbb{Z}_{256}^8$ as a function that takes a byte and returns a sequence of elements $0, 1 \in \mathbb{Z}_{256}$ corresponding to its binary representation in the order of most significant bit first. For example, $\mathsf{rep}(128) = (1, 0, 0, 0, 0, 0, 0, 0)$. We can generalize this notion to sequences of bytes. The function $\mathsf{Rep}(\cdot) : \mathbb{Z}_{256}^m \rightarrow \mathbb{Z}_{256}^{8 \cdot m}$ is defined as $\mathsf{Rep}(s) = \mathsf{rep}(s_1)||\ldots||\mathsf{rep}(s_m)$, e.g. $\mathsf{Rep}((192, 128)) = (1, 1, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0,$

[4] The proposal only specified the compression function and not the full hash function.

$0, 0, 0)$. Moreover, for two sequences of bytes we define $\oplus$ as the usual bitwise XOR of the two bitstrings.

We index elements of vectors and matrices starting from zero.

2.2 The LASH-x Hash Function

The LASH-x hash function maps an input of length less than 2^{2x} bits to an output of x bits. Four concrete proposals were suggested in [2]: $x = 160$, 256, 384, and 512.

The hash is computed by iterating a compression function that maps blocks of $n = 4x$ bits to $m = x/4$ bytes ($2x$ bits). The measure of n in bits and m in bytes is due to the original paper. Always $m = n/16$. Below we describe the compression function, and then the full hash function.

Compression Function of LASH-x. The compression function is of the form $f : \mathbb{Z}_{256}^{2m} \to \mathbb{Z}_{256}^{m}$. It is defined as

$$f(r, s) = (r \oplus s) + H \cdot [\mathsf{Rep}(r)||\mathsf{Rep}(s)]^T \ , \tag{1}$$

where $r = (r_0, \ldots, r_{m-1})$ and $s = (s_0, \ldots, s_{m-1})$ are column vectors[5] belonging to $\mathbb{Z}_{256}^{m}$. The vector r is called the *chaining variable*. The matrix H is a circulant matrix of dimensions $m \times (16m)$ defined as $H_{j,k} = a_{(j-k) \bmod 16m}$, where $a_i = y_i \bmod 2^8$, and y_i is defined as $y_0 = 54321$, $y_{i+1} = y_i^2 + 2 \pmod{2^{31} - 1}$ for $i > 0$. Our attacks do not use any properties of this sequence. In some cases, our analysis will split the matrix H into a left half H_L and a right half H_R (each of size $m \times 8m$), where H_L is multiplied by the bits of $\mathsf{Rep}(r)$ and H_R by the bits of $\mathsf{Rep}(s)$.

A visual diagram of the LASH-160 compression function is given in Figure 1, where t is $f(r, s)$.

The Full Function. Given a message of l bits, padding is first applied by appending a single '1'-bit followed by enough zeros to make the length a multiple of $8m = 2x$. The padded message consists of $\kappa = \lceil (l+1)/8m \rceil$ blocks of m bytes. Then, an extra block b of m bytes is appended that contains the encoded bit-length of the original message, $b_i = \lfloor l/2^{8i} \rfloor \pmod{256}$, $i = 0, \ldots, m-1$.

Next, the blocks $s^{(0)}, s^{(1)}, \ldots, s^{(\kappa)}$ of the padded message are fed to the compression function in an iterative manner as follows: $r^{(0)} := (0, \ldots, 0)$ and then $r^{(j+1)} := f(r^{(j)}, s^{(j)})$ for $j = 0, \ldots, \kappa$. The $r^{(0)}$ is called the IV.

[5] In this paper, we sometimes abuse notation when there is no confusion in the text: r and s can be both sequences of bytes as well as column vectors.

$$\begin{array}{c}40\\ \text{bytes}\end{array}\left\{\left(\begin{bmatrix} \\ r \\ \\ \end{bmatrix} \oplus \begin{bmatrix} \\ s \\ \\ \end{bmatrix}\right) + \overbrace{\begin{bmatrix} & & \\ \cdots & H & \cdots \\ & & \end{bmatrix}}^{640 \text{ columns}} \cdot \begin{bmatrix} \vdots \\ \mathsf{Rep}(r) \\ \vdots \\ - \\ \vdots \\ \mathsf{Rep}(s) \\ \vdots \end{bmatrix} = \begin{bmatrix} \\ t \\ \\ \end{bmatrix}\right.$$

Fig. 1. Visualizing the LASH-160 compression function.

Finally, the last chaining value $r^{(\kappa+1)}$ is sent through a *final transform* which takes only the 4 most significant bits of each byte to form the final hash value h. Precisely, the ith byte of h is $h_i = 16\lfloor r_{2i}/16\rfloor + \lfloor r_{2i+1}/16\rfloor$ $(0 \leq i < m/2)$.

3 Initial Observations

3.1 LASH is Not a PRF

In some applications (e.g. HMAC) it is required that the compression function (parameterized by its IV) should be a PRF. Below we show that LASH does not satisfy this property by exhibiting a differential that holds with probability 1, independent of the IV.

Assume that r is the secret parameter fixed beforehand and unknown to us. We are presented with a function $g(\cdot)$ which may be $f(r, \cdot)$ or a random function and by querying it we have to decide which one we have.

First we write $H = [H_L||H_R]$ and so (1) can be rewritten as

$$f(r, s) = (r \oplus s) + H_L \cdot \mathsf{Rep}(r)^T + H_R \cdot \mathsf{Rep}(s)^T \ \ .$$

Setting $s = 0$, we get $f(r, 0) = r + H_L \cdot \mathsf{Rep}(r)^T$. Now, for $s' = (128, 0, \ldots, 0)$ we have $\mathsf{Rep}(s') = 100\ldots0$ and so

$$f(r, s') = (r_0 \oplus 128, r_1, \ldots, r_{m-1}) + H_L \cdot \mathsf{Rep}(r)^T + H_R[\cdot, 0] \ \ ,$$

where $H_R[\cdot, 0]$ denotes the first column of the matrix H_R. One can readily compute the difference between $f(r, s')$ and $f(r, 0)$:

$$f(r, s') - f(r, 0) = H_R[\cdot, 0] + (128, 0, \ldots, 0)^T.$$

Regardless of the value of the secret parameter r, the output difference is a fixed vector equal to $H_R[\cdot, 0] + (128, 0, \ldots, 0)^T$. Thus, using only two queries we can distinguish with probability $1-2^{-8m}$ the LASH compression function with secret IV from a randomly chosen function.

The same principle can be used to distinguish LASH even if most of the bytes of s are secret as well. In fact, it is enough for us to control only one byte of the input to be able to use this method and distinguish with probability $1-2^{-8}$.

3.2 Absorbing the Feed-Forward Mode

According to [2], the feed-forward operation is motivated by Miyaguchi-Preneel hashing mode and is introduced to thwart some possible attacks on the plain matrix-multiplication construction. In this section we show two conditions under which the feed-forward operation can be described in terms of matrix operations and consequently absorbed into the LASH matrix multiplication step to get a simplified description of the compression function. The first condition requires one of the compression function inputs to be *known*, and the second requires a special subset of input messages.

First Condition: Partially Known Input. Suppose the r portion of the (r, s) input pair to the compression function is known and we wish to express the output $g(s) \stackrel{\text{def}}{=} f(r, s)$ in terms of the unknown input s. We observe that each $(8i + j)$th bit of the feedforward term $r \oplus s$ (for $i = 0, \ldots, m-1$ and $j = 0, \ldots, 7$) can be written as

$$\mathsf{Rep}(r \oplus s)_{8i+j} = \mathsf{Rep}(r)_{8i+j} + (-1)^{\mathsf{Rep}(r)_{8i+j}} \cdot \mathsf{Rep}(s)_{8i+j}.$$

Hence the value of the ith byte of $r \oplus s$ is given by

$$\sum_{j=0}^{7} \left(\mathsf{Rep}(r)_{8i+j} + (-1)^{\mathsf{Rep}(r)_{8i+j}} \cdot \mathsf{Rep}(s)_{8i+j} \right) \cdot 2^{7-j} =$$

$$\left(\sum_{j=0}^{7} \mathsf{Rep}(r)_{8i+j} \cdot 2^{7-j} \right) + \left(\sum_{j=0}^{7} (-1)^{\mathsf{Rep}(r)_{8i+j}} \cdot \mathsf{Rep}(s)_{8i+j} \cdot 2^{7-j} \right).$$

The first integer in parentheses after the equal sign is just the ith byte of r, whereas the second integer in parentheses is linear in the bits of s with known coefficients, and can be absorbed by appropriate additions to elements of the matrix H_R (defined in Section 2.2). Hence we have an 'affine' representation for $g(s)$:

$$g(s) = (D' + H_R) \cdot \mathsf{Rep}(s)^T + \underbrace{r + H_L \cdot \mathsf{Rep}(r)^T}_{m \times 1 \text{ vector}}, \tag{2}$$

where

$$D' = \begin{bmatrix} J_0 & 0_8 & \dots & 0_8 & 0_8 \\ 0_8 & J_1 & \dots & 0_8 & 0_8 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0_8 & 0_8 & \dots & J_{m-2} & 0_8 \\ 0_8 & 0_8 & \dots & 0_8 & J_{m-1} \end{bmatrix} .$$

Here, for $i = 0, \dots, m-1$, we define the 1×8 vectors $0_8 = [0, 0, 0, 0, 0, 0, 0, 0]$ and

$$J_i = [2^7 \cdot (-1)^{\mathsf{Rep}(r)_{8i}}, 2^6 \cdot (-1)^{\mathsf{Rep}(r)_{8i+1}}, \dots, 2^1 \cdot (-1)^{\mathsf{Rep}(r)_{8i+6}}, 2^0 \cdot (-1)^{\mathsf{Rep}(r)_{8i+7}}] .$$

Second Condition: Special Input Subset. Observe that when bytes of one of the input sequences (say, r) are restricted to values $\{0, 128\}$ only (i.e. only MS bit in each byte can be set), the XOR operation behaves like the byte-wise addition modulo 256. In other words, if $r^* = 128 \cdot r'$ where $r' \in \{0, 1\}^m$ then

$$\begin{aligned} f(r^*, s) &= r^* + s + H \cdot [\mathsf{Rep}(r^*)||\mathsf{Rep}(s)]^T \\ &= (D + H) \cdot [\mathsf{Rep}(r^*)||\mathsf{Rep}(s)]^T . \end{aligned} \tag{3}$$

The matrix D recreates values of r^* and s from their representations, similarly to matrix D' above. Then the whole compression function has the linear representation $f(r', s) = H' \cdot [r'||\mathsf{Rep}(s)]^T$ for a matrix H' which is matrix $D+H$ after removing $7m$ columns corresponding to the 7 LS bits of r for each byte. The resulting restricted function compresses $m + 8m$ bits to $8m$ bits using only matrix multiplication without any feed-forward mode.

4 Attacks Exploiting Zero IV

Collision Attack. In the original LASH paper, the authors describe a "hybrid attack" against LASH without the appended message length and final transform. Their idea is to do a Pollard or parallel collision search in such a way that each iteration forces some output bits to a fixed value (such as zero). Thus, the number of possible outputs is reduced from the standard attack. If the total number of possible outputs is S, then a collision is expected after about $\sqrt{S}$ iterations. Using a combination of table lookup and linear algebra, they are able to achieve $S = 2^{(14m/3)}$ in their paper. Thus, the attack is not effective since a collision is expected in about $2^{(7m/3)} = 2^{(7x/12)}$ iterations, which is more than the $2^{x/2}$ iterations one gets from the standard birthday attack on the full LASH function (with the final output transform).

Here, exploiting the zero IV, we describe a similar but simpler attack on the full function which uses table lookup only. Our messages will consist of a number of all-zero blocks followed by one "random" block. Regardless of the number of zero blocks at the beginning, the output of the compression function immediately prior to the length block being processed is determined entirely by the one "random" block. Thus, we will be using table lookup to determine a message length that results in a hash output value which has several bits in certain locations set to some predetermined value(s).

Refer to the visual diagram of the LASH-160 compression function in Fig. 1. Consider the case of the last compression, where the value of r is the output from the previous iteration and the value of s is the message length being fed in. The resulting hash value will consist of the most-significant half-bytes of the bytes of t. Our goal is to quickly determine a value of s so that the most significant half-bytes from the bottom part of t are all approximately zero.

Our messages will be long but not extremely long. Let α be the maximum number of bytes necessary to represent (in binary) any s that we will use. So the bottom $40 - \alpha$ bytes of s are all 0 bytes, and the bottom $320 - 8\alpha$ bits of $\mathsf{Rep}(s)$ are all 0 bits. As before, we divide the matrix H into two halves, H_L and H_R. Without specifying the entire s, we can compute the bottom $40-\alpha$ bytes of $(r \oplus s) + H_L \cdot \mathsf{Rep}(r)$. Thus, if we precomputed all possibilities for $H_R \cdot \mathsf{Rep}(s)$, then we can use table lookup to determine a value of s that hopefully causes h (to be chosen later) most-significant half-bytes from the bottom part of t to be 0. See the diagram in Fig. 2. The only restriction in doing this is $\alpha + h \leq 40$.

We additionally require dealing with the padding byte. To do so, we restrict our messages to lengths congruent to 312 mod 320. Then our "random" block can have anything for the first 39 bytes followed by `0x80` for the 40th byte which is the padding. We then ensure that only those lengths occur in our table lookup by only precomputing $H_R \cdot \mathsf{Rep}(s)$ for values of s of the form $320i + 312$. Thus, we have $\alpha = \lceil \frac{\log 320 + c}{8} \rceil$ assuming we take all values of i less than 2^c. We will aim for $h = c/4$, i.e. setting the bottom $c/4$ half-bytes of t equal to zero. The condition $\alpha + h \leq 40$ is then satisfied as long as $c \leq 104$, which will not be a problem.

Complexity. Pseudocode for the precomputation and table lookup can be found in Table 1 of [11]. With probability $1 - \frac{1}{e} \approx 0.632$, we expect to find a match in our table lookup. Assume that is the case.

$$\left(\overbrace{\begin{bmatrix}.\\.\\.\\.\\.\end{bmatrix}}^{r} \oplus \overbrace{\begin{bmatrix}\ell\\0\\0\\0\\0\end{bmatrix}}^{s}\right) + \overbrace{\begin{bmatrix} & | & \\ & | & \\ H_L & | & H_R \\ & | & \\ & | & \end{bmatrix}}^{H} \cdot \overbrace{\begin{bmatrix}.\\.\\.\\-\\ \ell\\0\\0\end{bmatrix}}^{\mathsf{Rep}(r||s)} = \overbrace{\begin{bmatrix}.\\.\\.\\0|.\\0|.\end{bmatrix}}^{t}$$

Fig. 2. Visualizing the final block of the attack on the LASH-160 compression function. Diagram is not to scale. Table lookup is done to determine the values at the positions marked with ℓ. Places marked with 0 are set to be zero by the attacker (in the t vector, this is accomplished with the table lookup). Places marked with '.' are outside of the attacker's control.

Due to rounding error, each of the bottom $c/4$ most significant half-bytes of t will either be 0 or -1 (`0xf` in hexadecimal). Thus there are $2^{c/4}$ possibilities for the bottom $c/4$ half-bytes, and the remaining $m - c/4 = x/4 - c/4$ half-bytes ($x - c$ bits) can be anything. So the size of the output space is $S = 2^{x-c+c/4} = 2^{x-3c/4}$. We expect a collision after we have about $2^{x/2-3c/8}$ outputs of this form. Note that with a Pollard or parallel collision search, we will not have outputs of this form a fraction of about $1/e$ of the time. This only means that we have to apply our iteration a fraction of $1/(1 - \frac{1}{e}) \approx 1.582$ times longer, which has negligible impact on the effectiveness of the attack. Therefore, we ignore such constants. Balancing the Pollard search time with the precomputation time, we get an optimal value with $c = (4/11)x$, i.e. a running time of order $2^{(4/11)x}$ LASH-x operations. The lengths of our colliding messages will be order $\leq 2^{c+\log 2x}$ bits. For instance, this attack breaks LASH-160 in as few as 2^{58} operations using 2^{58} storage. The lengths of our colliding messages will be order $\leq 2^{c+\log 2x}$ bits.

Experimental Results. We used this method to find collisions in LASH-160 truncated to the first 12 bytes of the hash: see [11]. The result took one week of cpu time on a 2.4GHz PC with $c = 28$.

Preimage Attack. The same lookup technique can be used for preimage attacks. One simply chooses random inputs and hashes them such that the looked up length sets some of the output hash bits to the target. This involves 2^c precomputation, 2^c storage, and $2^{x-3c/4}$ expected computation time, which balances time/memory to $2^{(4x/7)}$ using the optimal parameter setting $c = 4x/7$.

5 Short Message Preimage Attack on LASH with Arbitrary IV

The attacks in the previous section crucially exploit a particular parameter choice made by the LASH designers, namely the use of an all zero Initial Value (IV) in the Merkle-Damgård construction. Hence, it is tempting to try to 'repair' the LASH design by using a non-zero (or even random) value for the IV. In this section, we show that for any choice of IV, LASH-x is vulnerable to a preimage attack faster than the desired security level of $O(2^x)$. Our preimage attack takes time/memory $O(2^{(7x/8)})$, and produces preimages of short length ($2x$ bits). We give a general description of the attack below with parameter choices for LASH-160 in parentheses. Figure 3, which illustrates the attack for LASH-160, will be particularly useful in understanding our algorithm. For ease of description, we ignore the padding bit, but the reader should be able to see that this can be dealt with as well.

The Attack. Let $f : \mathbb{Z}_{256}^{2m} \to \mathbb{Z}_{256}^{m}$ denote the internal LASH compression function and $f_{out} : \mathbb{Z}_{256}^{2m} \to \mathbb{Z}_{16}^{m}$ denote the final compression function, i.e. the composition of f with the final transform applied to the output of f. Given a target value t_{out} whose LASH preimage is desired, the inversion algorithm finds a single block message $s_{in} \in \mathbb{Z}_{256}^{m}$ hashing to t_{out}, i.e. satisfying

$$f(r_{in}, s_{in}) = r_{out} \qquad \text{(first compression)}$$

and

$$f_{out}(r_{out}, s_{out}) = t_{out} \qquad \text{(final compression)}\ ,$$

where s_{out} is the $8m$-bit (320-bit for LASH-160) binary representation of the length block, and $r_{in} = IV$ is an arbitrary known value. The inversion algorithm proceeds as follows:

Step 1: *Find 2^m (2^{40} for LASH-160) inverses of the final compression.* Using the precomputation-based preimage attack on the final compression function f_{out} described in the previous section (with straightforward modifications to produce the preimage using bits of r_{out} rather than s_{out} and precomputation parameter $c_{out} = (20/7)m$ ($c_{out} \approx 114$ for LASH-160)), compute a list L of 2^m preimage values of r_{out} satisfying $f_{out}(r_{out}, s_{out}) = t_{out}$.

Step 2: *Search for a preimage s_{in} that maps to t_{out}.* Let $c = 3.5m$ ($c = 140$ for LASH-160) be a parameter (later we show that choosing $c = 3.5m$ is optimal). Split the $8m$-bit input s_{in} to be determined into two disjoint parts $s_{in}(1)$ (of length $6m - c$ bit) and $s_{in}(2)$ (of length $2m + c$ bit), i.e. $s_{in} = s_{in}(1)||s_{in}(2)$ (For LASH-160, we have $s_{in}(1)$ of length 100 bits and $s_{in}(2)$ of length 220 bits). We loop through all possibilities for the list L and the set of inputs $s_{in}(1)$ (a total of $2^m \cdot 2^{6m-c} = 2^{7m-c}$ possibilities, or 2^{140} possibilities for LASH-160). For each such possibility, we run the *internal compression function 'hybrid' partial inversion algorithm* described below to compute a matching 'partial preimage' value for $s_{in}(2)$, where by 'partial preimage' we mean that the compression function output $f(r_{in}, s_{in})$ matches target r_{out} on a fixed set of $m+c = 4.5m$ bits (180 bits for LASH-160). We leave the remaining $3.5m$ bits (140 bits for LASH-160) to chance. Thus, for each such computed partial preimage $s_{in} = s_{in}(1)||s_{in}(2)$ and corresponding r_{out} value, we check whether s_{in} is a *full* preimage, i.e. whether $f(r_{in}, s_{in}) = r_{out}$ holds, and if so, output desired preimage s_{in}.

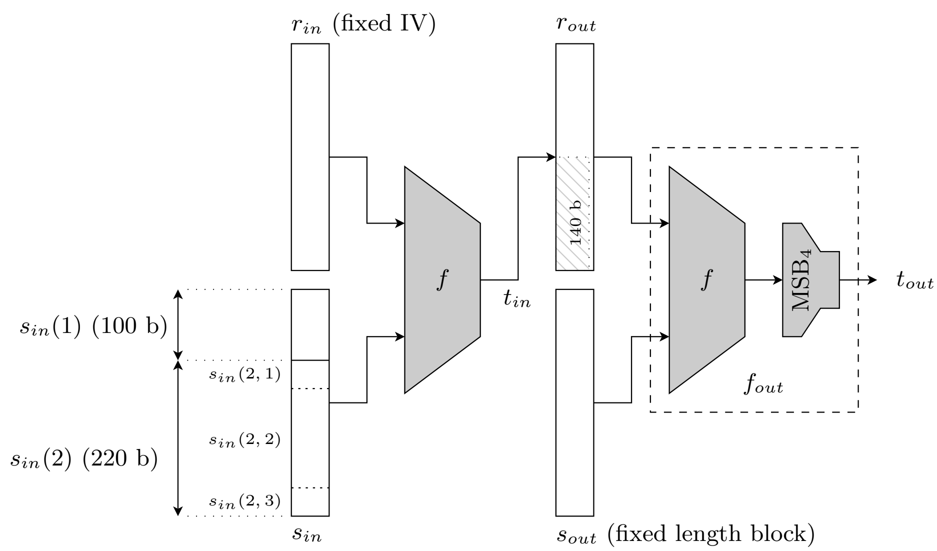

Fig. 3. Illustration of the preimage attack applied to LASH-160.

Internal Compression Function 'Hybrid' Partial Inversion Algorithm. For integer parameter c, the internal compression function 'hybrid' partial inversion algorithm is given a $8m$-bit target value t_{in} ($= r_{out}$), an $8m$-bit input r_{in}, and the $(6m - c)$-bit value $s_{in}(1)$, and computes a $(2m + c)$-bit value for $s_{in}(2)$ such that $f(r_{in}, s_{in})$ matches t_{in} on the top

$c/7$ bytes as well as on the LS bit of all remaining bytes. Hence, it matches on a total of $m + c$ bits. (For LASH-160, both t_{in} and r_{in} are 320 bits, $s_{in}(1)$ is 100 bits, $s_{in}(2)$ is 220 bits, and $f(r_{in}, s_{in})$ matches t_{in} on all of the bytes in the top half of t_{in} as well as all least significant bits). Some preliminaries are necessary before we explain the algorithm.

From Section 3.2 we know that for known r_{in}, the Miyaguchi-Preneel feedforward term $(r_{in} \oplus s_{in})$ can be absorbed into the matrix by appropriate modifications to the matrix and target vector, i.e. the inversion equation

$$(r_{in} \oplus s_{in}) + H \cdot [\mathsf{Rep}(r_{in})||\mathsf{Rep}(s_{in})]^T = t_{in} \bmod 256, \tag{4}$$

where H is the LASH matrix, can be reduced to an equivalent linear equation

$$H' \cdot [\mathsf{Rep}(s_{in})]^T = t'_{in} \bmod 256, \tag{5}$$

for appropriate matrix H' and vector t'_{in} easily computed from the known H, t_{in}, and r_{in}.

We require some notation and a one-time precomputation. We divide $s_{in}(2)$ into 3 parts $s(2,1)$ (length m bits = 40 bits for LASH-160), $s(2,2)$ (length c bits = 140 bits for LASH-160) and $s(2,3)$ (length m bits = 40 bits for LASH-160). For $i = 1, 2, 3$ let $H'(2,i)$ denote the submatrix of matrix H' from (5) consisting of the columns indexed by the bits in $s(2,i)$ (e.g. $H'(2,1)$ consists of the m columns of H' indexed by the m bits of $s(2,1)$). Similarly, let $H'(1)$ denote the submatrix of H' consisting of the columns of H' indexed by the m bits of $s_{in}(1)$.

The one-time precomputation pairs up values of $s(2,2)$ with $s(2,3)$ such that, after multiplying by the corresponding columns of H, the result has 0's in all m least significant bits. To do this, for each of 2^c possible values of $s(2,2)$, we find by linear algebra over $GF(2)$, a matching value for $s(2,3)$ such that

$$[H'(2,2)\ H'(2,3)] \cdot [\mathsf{Rep}(s(2,2))||\mathsf{Rep}(s(2,3))]^T = [0^m]^T \bmod 2. \tag{6}$$

The entry $s(2,2)||s(2,3)$ is stored in a hash table, indexed by the string of c bits obtained by concatenating 7 MS bits of each of the top $c/7$ bytes of vector y.

We are now ready to describe the internal compression function 'hybrid' partial inversion algorithm, based on [2] (a combination of table lookup and linear algebra), to find $s_{in}(2)$ such that the left and right hand sides of (5) match on the desired $m + c$ bits (180 bits for LASH-160). The gist of the algorithm is to use linear algebra to match the least significant bits and table lookup to match other bits. It works as follows:

– Solving Linear Equations: Compute $s(2,1)$ such that

$$H'(2,1)\cdot[\mathsf{Rep}(s(2,1))]^T = t'_{in} - H'(1)\cdot[\mathsf{Rep}(s_{in}(1))]^T \bmod 2. \quad (7)$$

Note that adding (6) and (7) implies that $H' \cdot [\mathsf{Rep}(s_{in}(1))||\mathsf{Rep}(s_{in}(2))]^T = t'_{in} \bmod 2$ with $s_{in}(2) = s(2,1)||s(2,2)||s(2,3)$ for any entry $s(2,2)||s(2,3)$ from the hash table. In other words, all least significant bits will be matched regardless of which entry is taken from the hash table.

– Lookup Hash Table: Find the $s(2,2)||s(2,3)$ entry indexed by the c-bit string obtained by concatenating the 7 MS bits of each of the top $c/7$ bytes of the vector $t'_{in} - H'(2,1)\cdot[\mathsf{Rep}(s(2,1))]^T - H'(1)\cdot[\mathsf{Rep}(s_{in}(1))]^T \bmod 256$. This implies that vector $H' \cdot [\mathsf{Rep}(s_{in}(1))||\mathsf{Rep}(s_{in}(2))]^T$ matches t'_{in} on all top $c/7$ bytes, as well as on the LS bits of all bytes, as required.

Correctness of Attack. For each of 2^m target values r_{out} from list L, and each of the $2^{2.5m}$ possible values for $s_{in}(1)$, the partial preimage inversion algorithm returns $s_{in}(2)$ such that $f(r_{in}, s_{in})$ matches r_{out} on a fixed set of $m+c$ bits. Modelling the remaining bits of $f(r_{in}, s_{in})$ as uniformly random and independent of r_{out}, we conclude that $f(r_{in}, s_{in})$ matches r_{out} on all $8m$ bits with probability $1/2^{8m-(m+c)} = 1/2^{7m-c} = 1/2^{3.5m}$ (using $c = 3.5m$) for each of the $2^{2.5m} \times 2^m = 2^{3.5m}$ runs of the partial inversion algorithm. Assuming that each of these runs is independent, the expected number of runs which produce a full preimage is $2^{3.5m} \times 1/2^{3.5m} = 1$, and hence we expect the algorithm to succeed and return a full preimage.

Complexity. The cost of the attack is dominated by the second step, where we balance the precomputation time/memory $O(2^c)$ of the hybrid partial preimage inversion algorithm with the expected number 2^{7m-c} of runs to get a full preimage. This leads (with the optimum parameter choice $c = 3.5m$) to time/memory cost $O(2^{3.5m}) = O(2^{(7x/8)})$, assuming each table lookup takes constant time. To see that second step dominates the cost, we recall that the first step with precomputation parameter c_{out} uses a precomputation taking time/memory $O(2^{c_{out}})$, and produces a preimage after an expected $O(2^{4m-3c_{out}/4})$ time using $c_{out} + (4m - 3c_{out}/4) = 4m + c_{out}/4$ bits of r_{out}. Hence, repeating this attack 2^m times using m additional bits of r_{out} to produce 2^m distinct preimages is expected to take $O(\max(2^{c_{out}}, 2^{5m-3c_{out}/4}))$ time/memory using $5m + c_{out}/4$ bits of r_{out}. The optimal choice for c_{out} is $c_{out} = (20/7)m \approx 2.89m$, and with this choice the first step takes $O(2^{(20/7)m}) = o(2^{3.5m})$ time/memory and uses $(40/7)m < 8m$ bits of r_{out} (the remaining bits of r_{out} are set to zero).

6 Attacks on the Final Compression Function

This section presents collision attacks on the final compression function f_{out} (including the output transform). For a given $r \in \mathbb{Z}_{256}^m$, the attacks produce $s, s' \in \mathbb{Z}_{256}^m$ with $s \neq s'$ such that $f_{out}(r, s) = f_{out}(r, s')$. To motivate these attacks, we note that they can be converted into a 'very long message' collision attack on the full LASH function, similar to the attack in Sect. 4. The two colliding messages will have the same final non-zero message block, and all preceding message blocks will be zero. To generate such a message pair, the attacker chooses a random $(8m-8)$-bit final message block (common to both messages), pads with a 0x80 byte, and applies the internal compression function f (with zero chaining value) to get a value $r \in \mathbb{Z}_{256}^m$. Then using the collision attack on f_{out} the attacker finds two distinct length fields $s, s' \in \mathbb{Z}_{256}^m$ such that $f_{out}(r, s) = f_{out}(r, s')$. Moreover, s, s' must be congruent to $8m-8 \pmod{8m}$ due to the padding scheme. For LASH-160, we can force s, s' to be congruent to $8m - 8 \pmod{64}$ by choosing the six LS bits of the length, so this leaves a $1/5^2$ chance that both inputs will be valid.

The lengths s, s' produced by the attacks in this section are very long (longer than $2^{x/2}$). However, we hope the ideas here can be used for future improved attacks.

6.1 Generalized Birthday Attack on the Final Compression

The authors of [2] describe an application of Wagner's generalized birthday attack [12] to compute a collision for the internal compression function f using $O(2^{2x/3})$ time and memory. Although this 'cubic root' complexity is lower than the generic 'square-root' complexity of the birthday attack on the full compression function, it is still higher than the $O(2^{x/2})$ birthday attack complexity on the *full* function, due to the final transformation outputting only half the bytes. Here we describe a variant of Wagner's attack for finding a collision in the *final* compression including the final transform (so the output bit length is x bits). The asymptotic complexity of our attack is $O\left(x2^{\frac{x}{2(1+\frac{107}{105})}}\right)$ time and memory – slightly better than a 'fourth-root' attack. For simplicity, we can call the running time $O(x2^{x/4})$.

The basic idea of our attack is to use the linear representation of f_{out} from Sect. 3.2 and apply a variant of Wagner's attack [12], modified to carefully deal with additive carries in the final transform. As in Wagner's original attack, we build a binary tree of lists with 8 leaves. At the ith level of the tree, we merge pairs of lists by looking for pairs of entries (one from each list) such that their sums have $7-i$ zero MS bits in selected output bytes, for $i = 0, 1, 2$. This ensures that the list at the root level has 4 zero MS bits on the selected bytes (these 4 MS bits are the output bits), accounting for the effect of carries during the merging process. More precise details are given below.

The attack. The attack uses inputs r, s for which the internal compression function f has a linear representation absorbing the Miyaguchi-Preneel feedforward (see Section 3.2). For such inputs, which may be of length up to $9m$ bit (recall: $m = x/4$), the *final* compression function $f' : \mathbb{Z}_{256}^{9m} \to \mathbb{Z}_{16}^{m}$ has the form

$$f'(r) = MS_4(H' \cdot [\mathsf{Rep}(r)]^T), \tag{8}$$

where $MS_4 : \mathbb{Z}_{256}^{m} \to \mathbb{Z}_{16}^{m}$ keeps only the 4 MS bits of each byte of its input, concatenating the resulting 4 bit strings (note that we use r here to represent the whole input of the linearised compression function f' defined in Section 3.2). Let $Rep(r) = (r[0], r[1], \ldots, r[9m-1]) \in \mathbb{Z}_{256}^{9m}$ with $r[i] \in \{0, 1\}$ for $i = 0, \ldots, 9m-1$. Let $\ell \approx \lfloor \frac{4m}{2(1+107/105)} \rfloor$ (notice that $8\ell < 9m$). We refer to each component $r[i]$ of r as an *input bit*. We choose a subset of 8ℓ input bits from r and partition the subset into 8 substrings $r^i \in \mathbb{Z}_{256}^{\ell}$ ($i = 1, \ldots, 8$) each containing ℓ input bits, i.e. $r = (r^1, r^2, \ldots, r^8)$. The linearity of (8) gives

$$f'(r) = MS_4(H'_1 \cdot [r^1]^T + \cdots + H'_8 \cdot [r^8]^T),$$

where, for $i = 1, \ldots, 8$, H'_i denotes the $m \times \ell$ submatrix of H' consisting of the ℓ columns indexed $(i-1)\cdot\ell, (i-1)\cdot\ell+1, \ldots, i\cdot\ell-1$ in H'. Following Wagner [12], we build 8 lists $L_1, \ldots, L_8$, where the ith list L_i contains all 2^ℓ possible candidates for the pair (r^i, y^i), where $y^i \stackrel{\text{def}}{=} H'_i \cdot [r^i]^T$ (note that y^i can be easily computed when needed from r^i and need not be stored). We then use a binary tree algorithm described below to gradually merge these 8 lists into a single list L^3 containing 2^ℓ entries of the form $(r, y = H' \cdot [r]^T)$, where the 4 MS bits in each of the first α bytes of y are zero, for some α, to be defined below. Finally, we search the list L^3 for a pair of entries which match on the values of the 4 MS bits of the last $m-\alpha$ bytes of the y portion of the entries, giving a collision for f' with the output being α zero half-bytes followed by $m-\alpha$ random half-bytes.

The list merging algorithm operates as follows. The algorithm is given the 8 lists $L_1, \ldots, L_8$. Consider a binary tree with $c = 8$ leaf nodes at level 0. For $i = 1, \ldots, 8$, we label the ith leaf node with the list L_i. Then, for each jth internal node n_j^i of the tree at level $i \in \{1,2,3\}$, we construct a list L_j^i labelling node n_j^i, which is obtained by merging the lists L_A^{i-1}, L_B^{i-1} at level $i-1$ associated with the two parent nodes of n_j^i. The list L_j^i is constructed so that for $i \in \{1,2,3\}$, the entries (r', y') of all lists at level i have the following properties:

- $(r', y') = (r'_A || r'_B, y'_A + y'_B)$, where (r'_A, y'_A) is an entry from the left parent list L_A^{i-1} and (r'_B, y'_B) is an entry from the right parent list L_B^{i-1}.
- If $i \geq 1$, the $\lceil \ell/7 \rceil$ bytes of y' at positions $0, \ldots, \lceil \ell/7 \rceil - 1$ each have their $(7-i)$ MS bits all equal to zero.
- If $i \geq 2$, the $\lceil \ell/6 \rceil$ bytes of y' at positions $\lceil \ell/7 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1$ each have their $(7-i)$ MS bits all equal to zero.
- If $i = 3$, the $\lceil \ell/5 \rceil$ bytes of y' at positions $\lceil \ell/7 \rceil + \lceil \ell/6 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil + \lceil \ell/5 \rceil - 1$ each have their $(7-i) = 4$ MS bits all equal to zero.

The above properties guarantee that all entries in the single list at level 3 are of the form $(r, y = H' \cdot [\mathsf{Rep}(r)]^T)$, where the first $\alpha = \lceil \ell/7 \rceil + \lceil \ell/6 \rceil + \lceil \ell/5 \rceil$ bytes of y all have $7 - 3 = 4$ MS bits equal to zero, as required.

To satisfy the above properties, we use a hash table lookup procedure, which aims, when merging two lists at level i, to fix the $7-i$ MS bits of some of the sum bytes to zero. This procedure runs as follows, given two lists L_A^{i-1}, L_B^{i-1} from level $i-1$ to be merged into a single list L^i at level i:

- Store the first component r'_A of all entries (r'_A, y'_A) of L_A^{i-1} in a hash table T_A, indexed by the hash of:
  - If $i = 1$, the 7 MS bits of bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ of y'_A, i.e. string $(MS_7(y'_A[0]), \ldots, MS_7(y'_A[\lceil \ell/7 \rceil - 1]))$.
  - If $i = 2$, the 6 MS bits of bytes $\lceil \ell/7 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1$ of y'_A, i.e. string $(MS_6(y'_A[\lceil \ell/7 \rceil]), \ldots, MS_6(y'_A[\lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1]))$.
  - If $i = 3$, the 5 MS bits of bytes $\lceil \ell/7 \rceil + \lceil \ell/6 \rceil, \ldots, \alpha - 1$ of y'_A, i.e. string $(MS_5(y'_A[\lceil \ell/7 \rceil + \lceil \ell/6 \rceil]), \ldots, MS_5(y'_A[\alpha - 1]))$.
- For each entry (r'_B, y'_B) of L_B^{i-1}, look in hash table T_A for matching entry (r'_A, y'_A) of L_A^{i-1} such that:
  - If $i = 1$, the 7 MS bits of corresponding bytes in positions $0, \ldots, \lceil \ell/7 \rceil - 1$ add up to zero modulo $2^7 = 128$, i.e. $MS_7(y'_A[j]) \equiv -MS_7(y'_B[j]) \bmod 2^7$ for $j = 0, \ldots, \lceil \ell/7 \rceil - 1$.
  - If $i = 2$, the 6 MS bits of corresponding bytes in positions $\lceil \ell/7 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1$ add up to zero modulo $2^6 = 64$, i.e.

$MS_6(y'_A[j]) \equiv -MS_6(y'_B[j]) \bmod 2^6$ for $j = \lceil \ell/7 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1$.

- If $i = 3$, the 5 MS bits of corresponding bytes in positions $\lceil \ell/7 \rceil + \lceil \ell/6 \rceil, \ldots, \alpha - 1$ add up to zero modulo $2^5 = 32$, i.e. $MS_5(y'_A[j]) \equiv -MS_5(y'_B[j]) \bmod 2^5$ for $j = \lceil \ell/7 \rceil + \lceil \ell/6 \rceil, \ldots, \alpha - 1$.

– For each pair of matching entries $(r'_A, y'_A) \in L_A^{i-1}$ and $(r'_B, y'_B) \in L_B^{i-1}$, add the entry $(r'_A \| r'_B, y'_A + y'_B)$ to list L^i.

Correctness. The correctness of the merging algorithm follows from the following simple fact:

Fact If $x, y \in \mathbb{Z}_{256}$, and the k MS bits of x and y (each regarded as the binary representation of an integer in $\{0, \ldots, 2^k - 1\}$) add up to zero modulo 2^k, then the $(k-1)$ MS bits of the byte $x + y$ (in $\mathbb{Z}_{256}$) are zero.

Thus, if $i = 1$, the merging lookup procedure ensures, by the Fact above, that the $7 - 1 = 6$ MS bits of bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ of $y'_A + y'_B$ are zero, whereas for $i \geq 2$, we have as an induction hypothesis that the $7 - (i - 1)$ MS bits of bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ of both y'_A and y'_B are zero, so again by the Fact above, we conclude that the $7 - i$ MS bits of bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ of $y'_A + y'_B$ are zero, which proves inductively the desired property for bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ for all $i \geq 1$. A similar argument proves the desired property for all bytes in positions $0, \ldots, \alpha - 1$. Consequently, at the end of the merging process at level $i = 3$, we have that all entries (r, y) of list L^3 have the $7 - 3 = 4$ MS bits of bytes $0, \ldots, \alpha - 1$ being zero, as required.

Asymptotic Complexity. The lists at level $i = 0$ have $|L^0| = 2^\ell$ entries. To estimate the expected size $|L^1|$ of the lists at level $i = 1$, we model the entries (r^0, y^0) of level 0 lists as having uniformly random and independent y^0 components. Hence for any pair of entries $(r_A^0, y_A^0) \in L_A^0$ and $(r_B^0, y_B^0) \in L_B^0$ from lists L_A^0 L_B^0 to be merged, the probability that the 7 MS bits of bytes $0, \ldots, \lceil \ell/7 \rceil - 1$ of y_A^0 and y_B^0 are negatives of each other modulo 2^7 is $\frac{1}{2^{\lceil \ell/7 \rceil \times 7}}$. Thus, the total expected number of matching pairs (and hence entries in the merged list L^1) is

$$|L^1| = \frac{|L_A^0| \times |L_B^0|}{2^{\lceil \ell/7 \rceil \times 7}} = \frac{2^{2\ell}}{2^{\lceil \ell/7 \rceil \times 7}} = 2^{\ell + O(1)}.$$

Similarly, for level $i = 2$, we model bytes $\lceil \ell/7 \rceil, \ldots, \lceil \ell/7 \rceil + \lceil \ell/6 \rceil - 1$ as uniformly random and independent bytes, and with the expected sizes $|L^1| = 2^{\ell + O(1)}$ of the lists from level 1, we estimate the expected size $|L^2|$

of the level 2 lists as:

$$|L^2| = \frac{|L_A^1| \times |L_B^1|}{2^{\lceil \ell/6 \rceil \times 6}} = 2^{\ell + O(1)},$$

and a similar argument gives also $|L^3| = 2^{\ell+O(1)}$ for the expected size of the final list. The entries (r, y) of L^3 have zeros in the 4 MS bits of bytes $0, \ldots, \alpha - 1$, and random values in the remaining $m - \alpha$ bytes. The final stage of the attack searches $|L^3|$ for two entries with identical values for the 4 MS bits of each of these remaining $m - \alpha$ bytes. Modelling those bytes as uniformly random and independent we have by a birthday paradox argument that a collision will be found with high constant probability as long as the condition $|L^3| \geq \sqrt{2^{4(m-\alpha)}}$ holds. Using $|L^3| = 2^{\ell+O(1)}$ and recalling that $\alpha = \lceil \ell/7 \rceil + \lceil \ell/6 \rceil + \lceil \ell/5 \rceil = (1/7 + 1/6 + 1/5)\ell + O(1) = \frac{107}{210}\ell + O(1)$, we obtain the attack success requirement

$$\ell \geq \frac{4m}{2(1 + \frac{107}{105})} + O(1) \approx \frac{x}{4} + O(1).$$

Hence, asymptotically, using $\ell \approx \lfloor \frac{x}{2(1+107/105)} \rfloor$, the asymptotic memory complexity of our attack is $O(x 2^{\frac{x}{2(1+\frac{107}{105})}}) \approx O(x 2^{x/4})$ bit, and the total running time is also $O(x 2^{\frac{x}{2(1+\frac{107}{105})}}) \approx O(x 2^{x/4})$ bit operations. So asymptotically, we have a 'fourth-root' collision finding attack on the final compression function.

Concrete Example. For LASH-160, we expect a complexity in the order of 2^{40}. In practice, the $O(1)$ terms increase this a little. Table 1 summarises the requirements at each level of the merging tree for the attack with $\ell = 42$ (note that at level 2 we keep only 2^{41} of the 2^{42} number of expected list entries to reduce memory storage relative to the algorithm described above). It is not difficult to see that the merging tree algorithm can be implemented such that at most 4 lists are kept in memory at any one time. Hence, we may approximate the total attack memory requirement by 4 times the size of the largest list constructed in the attack, i.e. $2^{48.4}$ bytes of memory. The total attack time complexity is approximated by $\sum_{i=0}^{3} |L^i| \approx 2^{43.3}$ evaluations of the linearised LASH compression function f', plus $\sum_{i=0}^{3} 2^{3-i}|L^i| \approx 2^{46}$ hash table lookups. The resulting attack success probability (of finding a collision on the 72 random output bits among the 2^{37} entries of list L^3) is estimated to be about $1 - e^{-0.5 \cdot 2^{37}(2^{37}-1)/2^{160-88}} \approx 0.86$. The total number of input bits used to form the collision is $8\ell = 336$ bit, which is less than the number $9m = 360$ bit available with the linear representation for the LASH compression function.

Table 1. Concrete Parameters of an attack on final compression function of LASH-160. For each level i, $|L^i|$ denotes the expected number of entries in the lists at level i, 'Forced Bytes' is the number of bytes whose $7-i$ MS bits are forced to zero by the hash table lookup process at this level, 'Zero bits' is four times the total number of output bytes whose 4 MS bits are guaranteed to be zero in list entries at this level, 'Mem/Item' is the memory requirement (in bit) per list item at this level, 'log(Mem)/List' is the base 2 logarithm of the total memory requirement (in bytes) for each list at this level (assuming that our hash table address space is twice the expected number of list items).

Level (i)	$log(\|L^i\|)$	**Forced Bytes**	**Zero bits**	**Mem/Item**, bit	log(`Mem`)**/List**, Byte
0	42	6	0	42	45.4
1	42	7	24	84	46.4
2	41	9	52	168	46.4
3	37		88	336	43.4

6.2 Heuristic Lattice-Based Attacks on the Final Compression

We investigated the performance of two heuristic lattice-based methods for finding collisions in truncated versions of the final compression function of LASH. The first reduces finding collisions to a lattice Shortest Vector Problem (SVP). The second uses the SVP as a preprocessing stage and applies a cycling attack with a lattice Closest Vector Problem (CVP) solved at each iteration. Due to lack of space, a detailed description of these attacks and the experimental results obtained can be found elsewhere [11]. The lattice-based attacks have the advantage of requiring very little memory. Preliminary experimental results for LASH-160 [11] give a time complexity estimate below 2^{68} for the CVP-based attack, significantly lower than the desired 2^{80}. Using our SVP-based attack, we found a collision in the final LASH-160 compression function truncated to 120 bits, with time complexity below 2^{36} (much less than the expected 2^{60}).

7 Conclusions

The LASH-x hash function was constructed by taking the GGH provable design [6] and duct taping on components that were intended for heuristic hashing. It is thus a combination of several techniques which are individually sound when applied to ideal components. Our work illustrates that this design strategy does not necessarily yield a secure result when applied to concrete components. In summary, we showed that LASH-x does not meet the designers' security goals nor does it meet other security goals that are typically assumed in heuristic hashing [5].

Acknowledgements: This research was supported in part by Australian Research Council (under grants DP0663452, DP0558773 and DP0665035) and by Singapore Ministry of Education (under Tier 2 grant T206B2204). Ron Steinfeld's research was supported by the Macquarie University Research Fellowship.

References

1. M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In *Advances in Cryptology – CRYPTO '96*, volume 1109 of *LNCS*, pages 1–15. Springer, 1996.
2. K. Bentahar, D. Page, M.-J. O. Saarinen, J. H. Silverman, and N. Smart. LASH. Second Cryptographic Hash Workshop, August, 24–25 2006.
3. D. J. Bernstein. Circuits for integer factorization: A proposal. Web page, http://cr.yp.to/papers/nfscircuit.pdf.
4. D. J. Bernstein. What output size resists collisions in a xor of independent expansions? ECRYPT Hash Workshop, May 2007.
5. S. Contini, R. Steinfeld, J. Pieprzyk, and K. Matusiewicz. A critical look at cryptographic hash function literature. ECRYPT Hash Workshop, May 2007.
6. O. Goldreich, S. Goldwasser, and S. Halevi. Collision-free hashing from lattice problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 3(042), 1996.
7. S. Lucks. Failure-friendly design principle for hash functions. In *Advances in Cryptology – ASIACRYPT '05*, volume 3788 of *LNCS*, pages 474–494. Springer, 2005.
8. V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. Provably Secure FFT Hashing (+ comments on "probably secure" hash functions). Second Cryptographic Hash Workshop, August, 24–25 2006.
9. A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.
10. C. Peikert. Private Communication, August 2007.
11. R. Steinfeld, S. Contini, K. Matusiewicz, J. Pieprzyk, J. Guo, S. Ling, and H. Wang. Cryptanalysis of LASH. Cryptology ePrint Archive, Report 2007/430, 2007. http://eprint.iacr.org/.
12. D. Wagner. A generalized birthday problem. In *Advances in Cryptology – CRYPTO '02*, volume 2442 of *LNCS*, pages 288–303. Springer, 2002.
13. M. J. Wiener. The full cost of cryptanalytic attacks. *J. Cryptol.*, 17(2):105–124, 2004.