FSBday:

Implementing Wagner's generalized birthday attack against the SHA-3* candidate FSB

Daniel J. Bernstein[1], Tanja Lange[2], Ruben Niederhagen[3], Christiane Peters[2], and Peter Schwabe[2] **

[1] Department of Computer Science
University of Illinois at Chicago, Chicago, IL 60607–7045, USA
djb@cr.yp.to

[2] Department of Mathematics and Computer Science
Technische Universiteit Eindhoven, P.O. Box 513, 5600 MB Eindhoven, Netherlands
tanja@hyperelliptic.org, c.p.peters@tue.nl, peter@cryptojedi.org

[3] Lehrstuhl für Betriebssysteme, RWTH Aachen University
Kopernikusstr. 16, 52056 Aachen, Germany
ruben@polycephaly.org

Abstract. The hash function FSB is one of the candidates submitted to NIST's competition to find the new standard hash function, SHA-3. The compression function of FSB is based on error correcting codes. In this paper we show how to use Wagner's generalized birthday attack to find collisions in FSB's compression function. In particular, we present details on our implementation attacking FSB_{48}, a toy version of FSB which was proposed by the FSB submitters as a training case for FSB. Our attack does not make use of any properties of the particular linear code used within FSB. FSB_{48} was chosen as a target where generalized birthday attacks would be one of the strongest attacks and which could be attacked in practice.

We show how to adapt this attack so that it runs on our computer cluster of only 10 PCs which provides far less memory than the usual implementation of generalized birthday attacks would require. This situation is very interesting for estimating the security of systems against distributed attacks using contributed off-the-shelf PCs.

For the SHA-3 competition this result is meaningful in that it allows to assess the security of FSB against the strongest non-structural attack; it does not provide any insight in the security of this particular choice of linear code.

Keywords: SHA-3, Birthday, FSB – Wagner, not much Memory

---

* SHA-2 will soon retire, see [9]

** This work was supported by the National Science Foundation under grant ITR-0716498, by the European Commission under Contract ICT-2007-216499 CACE, and by the European Commission under Contract ICT-2007-216646 ECRYPT II. Permanent ID of this document: ded1984108ff55330edb8631e7bc410c. Date: 2009.06.17.

1 Introduction

NIST's SHA-3 competition is attracting a lot of attention. Currently there are 51 submissions in the first round. This paper describes an attack against the SHA-3 candidate FSB [1]. FSB's compression function is based on error-correcting codes. It has been shown before for other code-based hash functions that Wagner's generalized birthday attack [11] is an efficient tool for finding collisions in the compression function. Unless structural attacks using properties of the specific linear code are possible, generalized birthday attacks are the most powerful attacks known against code-based hash functions.

In this paper we describe an implementation of the generalized birthday attack against a reduced-size version FSB_{48} which was suggested as a training case by the designers of FSB. The attack has not finished at the time of writing this document but we give performance figures for running the code for this attack. Our results allow to estimate how expensive a similar attack would be for full-size FSB.

We will show that a straightforward implementation of the attack would need more than 20 TB of storage. However, we are running the attack on the Coding and Cryptography Computer Cluster (CCCC) at Technische Universiteit Eindhoven which has a total hard-disk space of only 7 TB. We detail how we deal with this restricted background storage, by applying ideas described by Bernstein in [5] and introducing a compression technique of partial results.

We also explain the algorithmic measures we took to make the attack run as fast as possible, carefully balancing our code to use available RAM, network throughput, hard-disk throughput and computing power.

We are to the best of our knowledge the first to describe a full implementation of Wagner's generalized birthday attack. We plan to put all code described in this paper into the public domain to maximize reusability of our results.

Organization of the paper. In Section 2 we give a short introduction to Wagner's generalized birthday attack and Bernstein's adaptation of this attack to storage-restricted environments. Section 3 describes the FSB hash function to the extent necessary to understand our attack methodology. In Section 4 we describe our attack strategy which has to match the restricted hard-disk space of our computer cluster. Section 5 details the measures we applied to make the attack run as efficiently as possible dealing with the bottlenecks mentioned above. We evaluate the overall cost of our attack in Section 6, and give cost estimates for a similar attack against full-size FSB in Section 7.

Naming conventions. Throughout the paper we will denote list j on level i as $L_{i,j}$. For both levels and lists we start counting at zero.

Logarithms denoted as lg are logarithms to the base 2.

Additions of list elements or constants used in the algorithm are additions modulo 2.

In units such as GB, TB, PB and EB we will always assume base 1024 instead of 1000. In particular we give 700 GB as the size of a hard disk advertised as 750 GB.

2 Wagner's Generalized Birthday Attack

The generalized birthday problem, given 2^{i-1} lists containing B-bit strings, is to find 2^{i-1} elements — exactly one in each list — whose xor equals 0.

The special case $i = 2$ is the classic birthday problem: given two lists containing B-bit strings, find two elements — exactly one in each list — whose xor equals 0. In other words, find an element of the first list that equals an element of the second list.

This section describes a solution to the generalized birthday problem due to Wagner [11]. Wagner also considered generalizations to operations other than xor, and to the case of k lists when k is not a power of 2.

2.1 The tree algorithm

Wagner's algorithm builds a binary tree as described in this subsection starting from the input lists $L_{0,0}, L_{0,1}, \ldots, L_{0,2^{i-1}-1}$. See Figure 1 in Section 4. The speed and success probability of the algorithm are analyzed under the assumption that each list contains $2^{B/i}$ elements chosen uniformly at random.

On level 0 take the first two lists $L_{0,0}$ and $L_{0,1}$ and compare their list elements on their least significant B/i bits. Given that each list contains about $2^{B/i}$ elements we can expect $2^{B/i}$ pairs of elements which are equal on those least significant B/i bits. We take the xor of both elements on all their B bits and put the xor into a new list $L_{1,0}$. Similarly compare the other lists — always two at a time — and look for elements matching on their least significant B/i bits which are xored and put into new lists. This process of *merging* yields 2^{i-2} lists containing each about $2^{B/i}$ elements which are zero on their least significant B/i bits. This completes level 0.

On level 1 take the first two lists $L_{1,0}$ and $L_{1,1}$ which are the results of merging the lists $L_{0,0}$ and $L_{0,1}$ as well as $L_{0,2}$ and $L_{0,3}$ from level 0. Compare the elements of $L_{1,0}$ and $L_{1,1}$ on their least significant $2B/i$ bits. As a result of the xor'ing in the previous level, the last B/i bits are already known to be 0, so it suffices to compare the next B/i bits. Since each list on level 1 contains about $2^{B/i}$ elements we again can expect about $2^{B/i}$ elements matching on B/i bits. We build the xor of each pair of matching elements and put it into a new list $L_{2,0}$. Similarly compare the remaining lists on level 1.

Continue in the same way until level $i-2$. On each level j we consider the elements on their least significant $(j+1)B/i$ bits of which jB/i bits are known to be zero as a result of the previous merge. On level $i-2$ we get two lists containing about $2^{B/i}$ elements. The least significant $(i-2)B/i$ bits of each element in both lists are zero. Comparing the elements of both lists on their $2B/i$ remaining bits gives 1 expected match, i.e., one xor equal to zero. Since each element is the xor of elements from the previous steps this final xor is the xor of 2^{i-1} elements from the original lists and thus a solution to the generalized birthday problem.

2.2 Wagner in memory-restricted environments

A 2007 paper [5] by Bernstein includes two techniques to mount Wagner's attack on computers which do not have enough memory to hold all list entries. Various special cases of the same techniques also appear in a 2005 paper [3] by Augot, Finiasz, and Sendrier and in a 2009 paper [8] by Minder and Sinclair.

Clamping through precomputation. Suppose that there is space for lists of size only 2^b with $b < B/i$. Bernstein suggests to generate $2^{b \cdot (B-ib)}$ entries and only consider those of which the least significant $B - ib$ bits are zero.

This idea can be generalized as follows: The least significant $B - ib$ bits can have an arbitrary value, this *clamping value* does not even have to be the same on all lists as long as the *sum* of all clamping values is zero. This will be important if an attack does not produce a collision. We can then simply restart the attack with different clamping values.

Clamping through precomputation may be limited by the maximal number of entries we can generate per list. Furthermore, halving the available storage space increases the precomputation time by a factor of 2^i.

Note that clamping some bits through precomputation might be a good idea even if enough memory is available as we can reduce the amount of data in later steps and thus make those steps more efficient.

After the precomputation step we apply Wagner's tree algorithm to lists containing bit strings of length B' where B' equals B minus the number of clamped bits. For performance evaluation we will only consider lists on level 0 *after* clamping through precomputation and then use B instead of B' for the number of bits in these entries.

Repeating the attack. Another way to mount Wagner's attack in memory-restricted environments is to carry out the whole computation with smaller lists leaving some bits at the end "uncontrolled". We can then deal with the lower success probability by repeatedly running the attack with different clamping values.

In the context of clamping through precomputation we can simply vary the clamping values used during precomputation. If for some reason we cannot clamp any bits through precomputation we can apply the same idea of changing clamping values in an arbitrary merge step of the tree algorithm. Note that any solution to the generalized birthday problem can be found by some choice of clamping values.

Expected number of runs. Wagner's algorithm, without clamping through precomputation, produces an expected number of exactly one collision. However this does not mean that running the algorithm necessarily produces a collision.

In general, the expected number of runs of Wagner's attack is a function of the number of remaining bits in the entries of the two input lists of the last merge step and the number of elements in these lists.

Assume that b bits are clamped on each level and that lists have length 2^b. Then the probability to have at least one collision after running the attack once is

$$P_{\text{success}} = 1 - \left(\frac{2^{B-(i-2)b} - 1}{2^{B-(i-2)b}} \right)^{2^{2b}},$$

and the expected number of runs $E(R)$ is

$$E(R) = \frac{1}{P_{\text{success}}}. \tag{2.1}$$

For larger values of $B - ib$ the expected number is about 2^{B-ib}. It is common to model the total time for one run as being linear in the total list size; i.e., $(i-1)2^{i-1}2^{B-ib}2^b$. Here $i-1$ is the number of levels, 2^{i-1} is the number of lists, 2^{B-ib} is approximately the number of runs, and 2^b is the number of entries per list.

Using Pollard iteration. If because of memory restrictions the number of uncontrolled bits is high, it may be more efficient to use a variant of Wagner's attack that uses Pollard iteration [7, Chapter 3, exercises 6 and 7].

Assume that $L_0 = L_1$, $L_2 = L_3$, etc., and that combinations $x_0 + x_1$ with $x_0 = x_1$ are excluded. The output of the generalized birthday attack will then be a collision between two distinct elements of $L_0 + L_2 + \cdots$.

We can instead start with only 2^{i-2} lists $L_0, L_2, \ldots$ and apply the usual Wagner tree algorithm, with a nonzero clamping constant to enforce the condition that $x_0 \neq x_1$. The number of clamped bits before the last merge step is now $(i-3)b$. The last merge step produces 2^{2b} possible values, the smallest of which has an expected number of $2b$ leading zeros, leaving $B - (i-1)b$ uncontrolled.

Think of this computation as a function mapping clamping constants to the final $B - (i-1)b$ uncontrolled bits and apply Pollard iteration to find a collision between the output of two such computations; combination then yields a collision of 2^{i-1} vectors.

As Pollard iteration has square-root running time, the expected number of runs for this variant is $2^{B/2-(i-1)b/2}$, each taking time 2^b, so the exptected running time is

$$t = 2^{B/2-(i-1)b/2+b} \tag{2.2}$$

The Pollard variant of the attack becomes more efficient than plain Wagner with repeated runs if $B > (i+2)b$.

3 The FSB Hash Function

In this section we briefly describe the construction of the FSB hash function. Since we are going to attack the function we omit details which are necessary for implementing the function but do not influence the attack. The second part of this section gives a rough description of how to apply Wagner's generalized birthday attack to find collisions of the compression function of FSB.

3.1 Details of the FSB hash function

The Fast Syndrome Based hash function (FSB) was introduced by Augot, Finiasz and Sendrier in 2003. See [2], [3], and [1]. The security of FSB's compression function relies on the difficulty of the "Syndrome Decoding Problem" from coding theory.

The FSB hash function processes a message in three steps: First the message is converted by a so-called domain extender into suitable inputs for the compression function which digests the inputs in the second step. In the third and final step the Whirlpool hash function designed by Barreto and Rijmen [4] is applied to the output of the compression function in order to produce the desired length of output.

Our goal in this paper is to investigate the security of the compression function. We do not describe the domain extender, the conversion of the message to inputs for the compression function, or the last step involving Whirlpool.

The compression function. The main parameters of the compression function are called n, r and w. We consider n strings of length r which are chosen uniformly at random and can be written as an $r \times n$ binary matrix H. Note that the matrix H can be seen as the parity check matrix of a binary linear code. The FSB proposal [1] actually specifies a particular structure of H for efficiency; we do not consider attacks exploiting this structure.

An n-bit string of weight w is called *regular* if there is exactly a single 1 in each interval $[(i-1)\frac{n}{w}, i\frac{n}{w} - 1]_{1 \le i \le w}$. We will refer to such an interval as a *block*. The input to the compression function is a regular n-bit string of weight w.

The compression function works as follows. The matrix H is split into w blocks of n/w columns. Each non-zero entry of the input bit string indicates exactly one column in each block. The output of the compression function is an r-bit string which is produced by computing the xor of all the w columns of the matrix H indicated by the input string.

Preimages and collisions. A preimage of an output of length r of one round of the compression function is a regular n-bit string of weight w. A collision occurs if there are $2w$ columns of H — exactly two in each block — which add up to zero.

Finding preimages or collisions means solving two problems coming from coding theory: finding a preimage means solving the Regular Syndrome Decoding problem and finding collisions means solving the so-called 2-regular Null-Syndrome Decoding problem. Both problems were defined and proved to be NP-complete in [3].

Parameters. We follow the notation in [1] and write $\text{FSB}_{\texttt{length}}$ for the version of FSB which produces a hash value of length `length`. Note that the output of the compression function has r bits where r is considerably larger than `length`.

NIST demands hash lengths of 160, 224, 256, 384, and 512 bits, respectively. Therefore the SHA-3 proposal contains five versions of FSB: FSB_{160}, FSB_{224}, FSB_{256}, FSB_{384}, and FSB_{512}. The proposal also contains FSB_{48}, which is a reduced-size version of FSB and the main attack target in this paper. The binary matrix H for FSB_{48} has dimension $192 \times 3 \cdot 2^{17}$; i.e., r equals 192 and n is $3 \cdot 2^{17}$. In each round a message chunk is converted into a regular $3 \cdot 2^{17}$-bit string of Hamming weight $w = 24$. The matrix H contains 24 blocks of length 2^{14}. Each 1 in the regular bit string indicates exactly one column in a block of the matrix H. The output of the compression function is the xor of those 24 columns.

3.2 Attacking the compression function of FSB_{48}

Coron and Joux pointed out in [6] that Wagner's generalized birthday attack can be used to find preimages and collisions in the compression function of FSB. The following paragraphs present a slightly streamlined version of the attack of [6] in the case of FSB_{48}.

Determining the number of lists for a Wagner attack on FSB_{48}. A collision for FSB_{48} is given by 48 columns of the matrix H which add up to zero; the collision has exactly two columns per block. Each block contains 2^{14} columns and each column is a 192-bit string.

We choose 16 lists to solve this particular 48-sum problem. Each list entry will be the xor of three columns coming from one and a half blocks. This ensures that we do not have any overlaps, i.e., more than two columns coming from one matrix block in the end. We assume that taking sums of the columns of H does not bias the distribution of 192-bit strings. Applying Wagner's attack in a straightforward way means that we need to have at least $2^{\lceil 192/5 \rceil}$ entries per list. By clamping away 39 bits in each step we expect to get at least one collision after one run of the tree algorithm.

Building lists. We build 16 lists containing 192-bit strings each being the xor of three distinct columns of the matrix H. We select each triple of three columns from one and a half blocks of H in the following way:

List $L_{0,0}$ contains the sums of columns i_0, j_0, k_0, where columns i_0 and j_0 come from the first block of 2^{14} columns, and column k_0 is picked from the following block with the restriction that it is taken from the first half of it. Since we cannot have overlapping elements we get 2^{27} sums of columns i_0 and j_0 coming from the first block. These two columns are then added to all possible columns k_0 coming from the first 2^{13} elements of the second block of the matrix H. In total we get roughly 2^{40} elements for $L_{0,0}$.

We note that by splitting every second block in half we neglect several solutions of the 48-xor problem. For example, a solution involving two columns from the first half of the second block cannot be found by this algorithm. We justify our choice by noting that fewer lists would nevertheless require more storage and a longer precomputation phase to build the lists.

The second list $L_{0,1}$ contains sums of columns i_1, j_1, k_1, where column i_1 is picked from the second half of the second block of H and j_1 and k_1 come the third block of 2^{14} columns. This again yields about 2^{40} elements.

Similarly, we construct the lists $L_{0,2}$, $L_{0,3}$,..., $L_{0,15}$.

For each list we generate more than twice the amount needed for a straightforward attack as explained above. In order to reduce the amount of data for the following steps we note that about $2^{40}/4$ elements are likely to be zero on their least significant two bits. Clamping those two bits away should thus yield a list of 2^{38} bit strings. Note that since we know the least significant two bits of the list elements we can ignore them and regard the list elements as 190-bit strings. Now we expect that a straightforward application of Wagner's attack to 16 lists with about $2^{190/5}$ elements yields a collision after completing the tree algorithm.

Note on complexity in the FSB proposal. The SHA-3 proposal estimates the complexity of Wagner's attack as described above as $2^{r/i}r$ where 2^{i-1} is the number of lists used in the algorithm. This does not take memory into account, and in general is an underestimate of the work required by Wagner's algorithm; i.e., attacks of this type against FSB are more difficult than claimed by the FSB designers.

Note on information-set decoding. The FSB designers say in [1] that Wagner's attack is the fastest known attack for finding preimages, and for finding collisions for small FSB parameters, but that another attack — information-set decoding — is better than Wagner's attack for finding collisions for large FSB parameters.

In general information-set decoding can be used to find an n-bit string of Hamming weight 48. Information-set decoding will not take into account that we look for a regular n-bit string; it has to be run repeatedly until its output happens to be regular. Thus, the running times given in [1] provide certainly lower bounds for information-set decoding, but in practice they are not likely to hold.

4 Attack Strategy

This section we will discuss the necessary measures we took to mount the attack on our cluster. We will start with an evaluation of available and required storage.

4.1 Storage requirements

How large is a list entry? The obvious way of representing a list entry is as a 192-bit string, the xor of three columns of the matrix. Bits we already know to be zero of course do not have to be stored so in each level of the tree the number of bits per entry decreases by the number of bits clamped in the previous level. Ultimately we are not interested in the *value* of the entry — we know already that in a successful attack it will be all-zero at the end — but in the column positions in the matrix that lead to this all-zero value. So we have to store the positions alongside the value; unlike storage requirements for *values* the number of bytes for *positions* increases with increasing levels, and becomes dominant for higher levels.

We have considered two strategies to reduce the number of bits per entry: *compressed positions* and *dynamic recomputation*. The compressed-positions approach does not store full positions but e.g. positions modulo 256. After the attack has successfully finished the full position information can be computed by checking which of the possible positions lead to the appropriate intermediate results on each level.

Dynamic recomputation reduces the storage requirements by not storing the entry value at all but recomputing it every time it is needed from the positions.

For our implementation we decided to use this second approach: As discussed in Section 3 we have 2^{40} possibilities to choose columns to produce entries of a list, so we can encode the positions in 40 bits (5 bytes).

In each level the size of a single entry doubles (because the number of positions doubles), the expected number of entries per list remains the same but the number of lists halves, so the total amount of data is the same on each level when using dynamic recomputation.

What list size can we handle? As described in Section 3.2 we can start with 16 lists of size 2^{38}, each containing bit strings of length $r' = 190$. However, storing 16 lists with 2^{38} entries, each entry encoded in 5 bytes requires 20480 GB of storage space. The Coding and Cryptography Computer Cluster at Eindhoven University of Technology only has a total hard disk space of 7 TB, so we have to adapt our attack strategy to this limitation.

On the first level we have 16 lists and as we need at least 5 bytes per list entry we can handle at most $7 \cdot 2^{40}/2^4/5 = 1.36 \times 2^{36}$ entries per list. Some of the disk space is used for operating system and so a straight-forward implementation would use lists of size 2^{36}.

We can generate at most 2^{40} entries per list so following [5] we could clamp 4 bits during list generation, giving us 2^{36} values for each of the 16 lists. These values have a length of 188 bits represented through 5 bytes holding the positions from the matrix. Clamping 36 bits in each of the 3 steps leaves two lists of length 2^{36} with 80 non-zero bits. According to (2.1) we thus expect to run the attack 256.5 times until we find a collision.

We could compensate for this relatively low success probability by repeatedly running the attack with different clamping constants but that would make the attack very inefficient. Instead we will now describe how we can handle lists of size 2^{37}.

4.2 The strategy

The idea is to first use 8 nodes to compute list $L_{3,0}$ (see Figure 1), the result list of the left half tree. During this computation we can use lists with 2^{37} entries which need a total storage of 5120 GB. This means that during list generation we only clamp 3 bits and then during each merge clamp 37 bits yielding a list $L_{3,0}$ containing entries with 78 non-zero bits.

Each entry in the resulting list $L_{3,0}$ contains 24 positions from the matrix and thus needs 40 bytes. We can however compress this data by now storing the

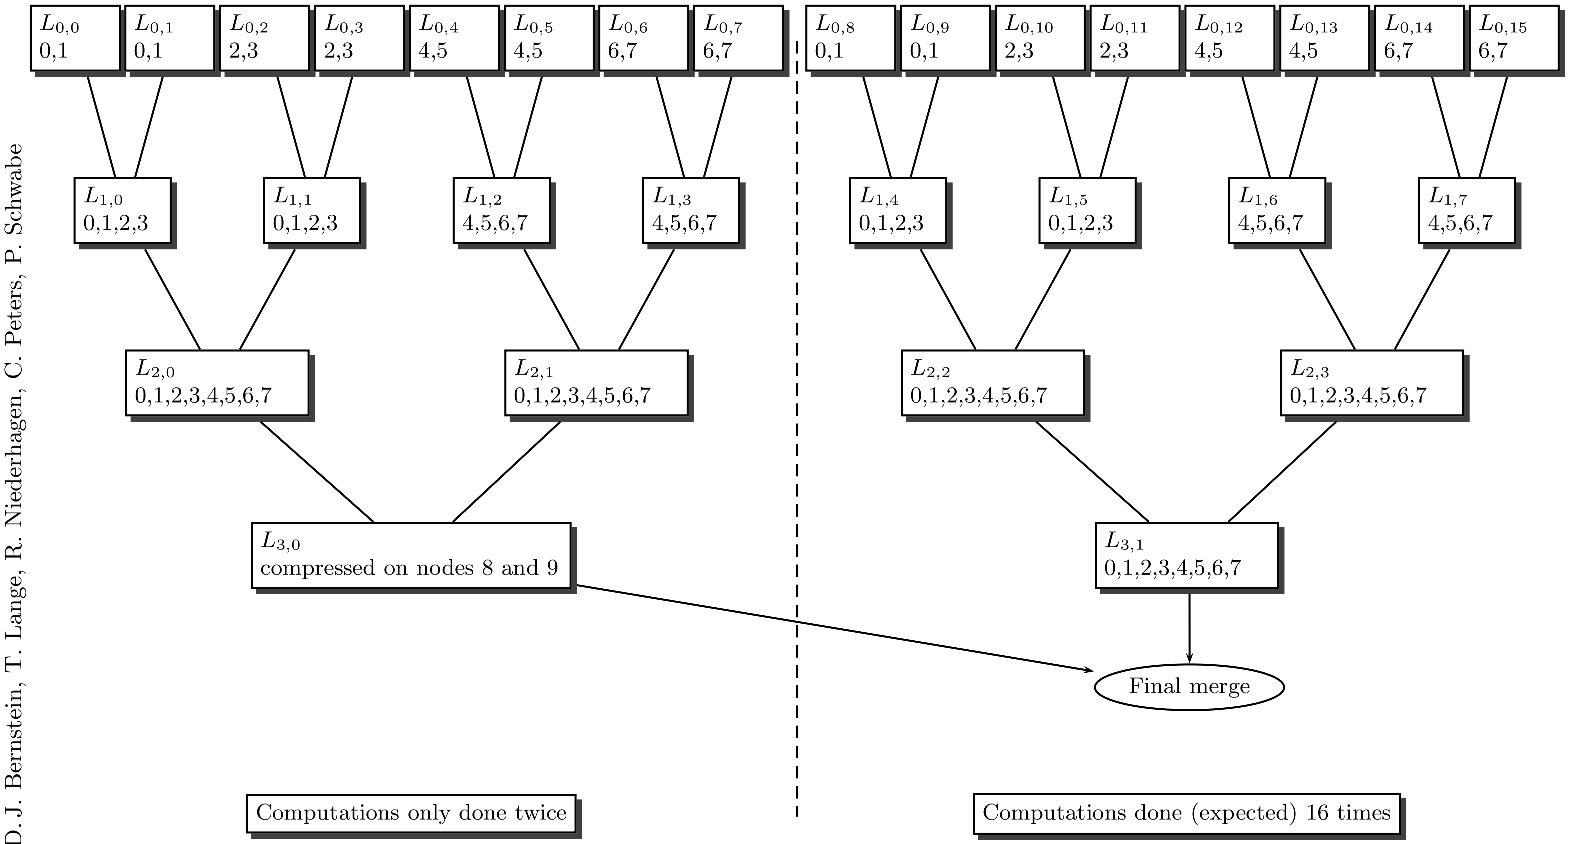

Fig. 1. Structure of the attack: in each box the upper line denotes the list, the lower line gives the nodes holding fractions of this list

value instead of the positions. The remaining number of non-zero bits per entry is 78 (10 bytes), so we can compress list $L_{3,0}$ by a factor of four and send the resulting 1280 GB of compressed data to the two remaining nodes to sort and store.

Then we proceed with the right half-tree until we have list $L_{3,1}$ and for the final merge compare with the values stored on the two remaining nodes. If we find a collision we have to do the left half-tree computation again to reconstruct the positions from the value. Observe that during this final merge we now only have 4 uncontrolled bits (two lists of length 2^{37} containing entries with 78 non-zero bits), so this approach decreases the expected number of runs to 16.5.

If we do not find a collision in the final merge we could start again with computation of the left half-tree with different clamping constants. A more efficient approach is to keep the data of the left half-tree and just change clamping constants in the right half-tree in a way that they add up to zero until we find a collision with the values from the left half-tree. So we need to compute $L_{3,0}$ just once at the beginning and again after finding a collision. We expect to have 16.5 computations of $L_{3,1}$, each using different clamping constants, and expect to have 16.5 final merge steps.

5 Implementing the Attack

The computation platform for this particular implementation of Wagner's generalized birthday attack on FSB is a ten-node cluster of conventional desktop PCs. Each node has an Intel Core 2 Quad Q6600 CPU with a clock rate of 2.40GHz and direct fully cached access to 8 GB of RAM. About 700 GB mass storage are provided by a Western Digital SATA hard disk with 20 GB reserved for system and user data. The nodes are connected via switched Gigabit Ethernet using Marvell PCI-E adapter cards.

We chose MPI as communication model for the implementation. This choice has several virtues:

- MPI provides an easy interface to start the application on all nodes and to initialize the communication paths.
- MPI offers synchronous message-based communication primitives.
- MPI is a broadly accepted standard for HPC applications and is provided on a multitude of different platforms.

We decided to use MPICH2 which is an implementation of the MPI 2.0 standard from the University of Chicago. MPICH2 provides an Ethernet-based back end for the communication with remote nodes and a fast shared-memory-based back end for local data exchange.

We implemented two micro-benchmarks to measure hard disk and network throughput. The results of these benchmarks are shown in Figure 2. Note that we measure hard disk throughput directly on the device, circumventing the filesystem, to reach peak performance of the hard disk. We measured both sequential and randomized access to the disk.

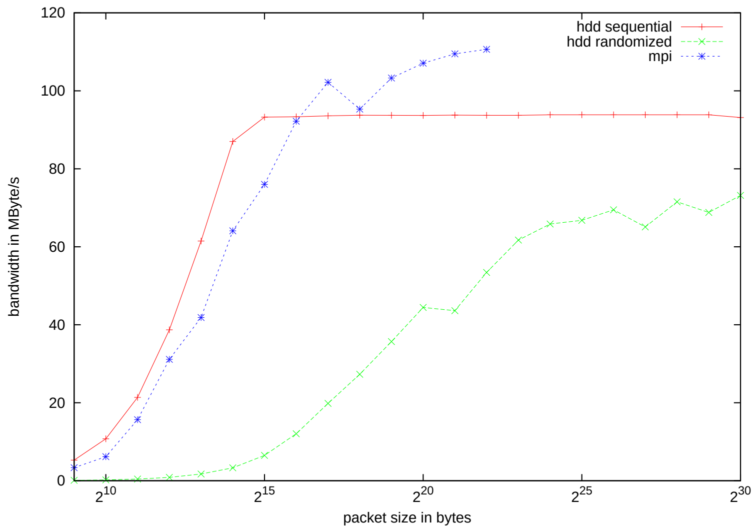

Fig. 2. Micro-benchmarks measuring hard disk and network throughput.

The rest of this section explains how we parallelized and streamlined Wagner's attack to make the best of the available hardware.

5.1 Parallelization

Most of the time in the attack is spent computing the list entries for one half tree. This is done by initially generating the list entries with the current clamping constant on level 0. Then each list is sorted and afterwards merged with its neighboring list giving the entries for the next level. The sorting and merging is repeated until level 3 produces one final list.

This algorithm is parallelized by distributing fractions of lists over the nodes in a way that each node can perform sort and merge locally on two lists. On each level of the computation, each node contains fractions of two lists; on level 0 each node contains two half-lists, on level 1 each node contains two quarter-lists, etc. The lists on level j are split between the nodes according to $(j + 1)$ bits of the value. For example, on level 0, node 0 contains all entries of lists 0 and 1 ending with a zero bit (in the bits not controlled by initial clamping), and node 1 contains all entries of lists 0 and 1 ending with a one bit.

Therefore from the view of one node, on each level the fractions of both lists are loaded from hard disk, the entries are sorted and the two lists are merged. The newly generated list is split into its fractions and these fractions are sent over the network to their associated nodes. There the data is received and stored onto the hard drive. All this is implemented in a way that allows a continuous data stream to be generated and processed (see Figure 3).

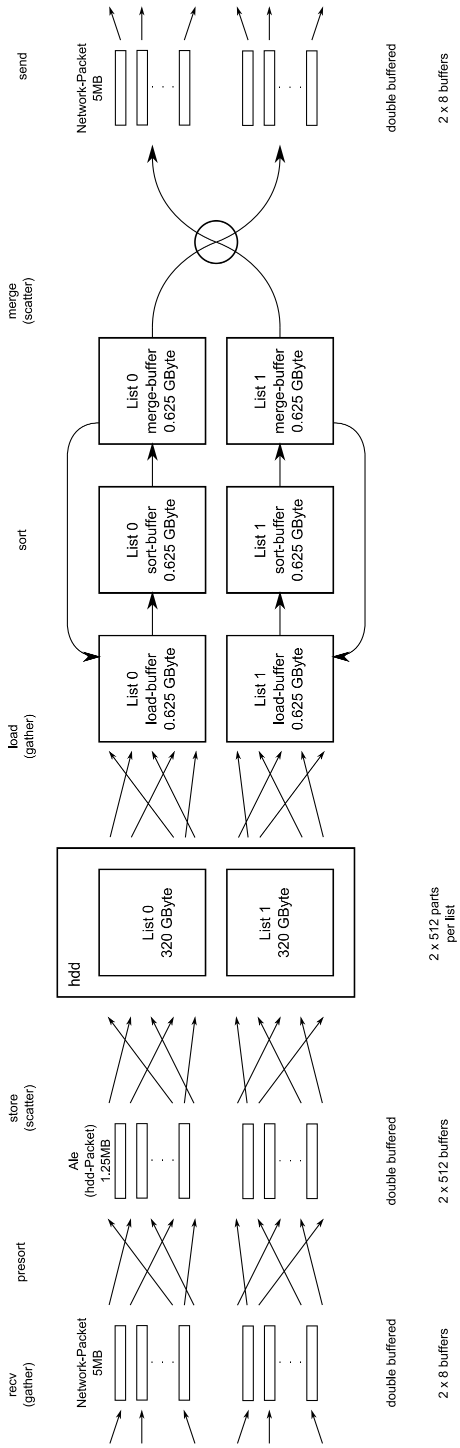

Fig. 3. Data flow during the computation within one half-tree

To be able to perform the sort in memory, incoming data is presorted into one of 512 parts looking at the nine least significant bits of the current sort range. This leads to an expected part size of 640 MB (0.625 GB) which can be loaded into main memory at once to be sorted further. The benefit of presorting the entries before storing them is:

1. We can sort a whole fraction, that exceeds the size of the memory, by sorting its presorted parts independently.
2. Two adjacent parts of the two lists on one node (with the same presort-bits) can be merged directly after they are sorted.

The merge is implemented straightforwardly. If blocks of entries in both lists share the same value then all possible combinations are generated: specifically, if a b-bit string appears in the compared positions in c_1 entries in the first list and c_2 entries in the second list then all $c_1 c_2$ xors appear in the output list.

5.2 Efficient implementation

Cluster computation imposes three main bottlenecks:

- the computational power and memory latency of the CPUs for computation-intensive applications
- limitations of network throughput and latency for communication-intensive applications
- hard disk throughput and latency for data-intensive applications

Wagner imposes hard load on all of these components: a large amount of data needs to be sorted, merged and distributed over the nodes occupying as much storage as possible. Therefore demand for optimization is primarily determined by the slowest component in terms of data throughput; latency generally can be hidden by pipelining and data prefetch.

Our benchmarks show that, for sufficiently large packets, the performance of the system is mainly bottlenecked by hard-disk throughput. Since the throughput of MPI over Gigabit Ethernet is higher than the hard-disk throughput for packet sizes larger than 2^{16} bytes and since the same amount of data has to be sent that needs to be stored, no performance penalty is expected by the network for this size of packets.

Therefore our first implementation goal was to design an interface to the hard disk that allows for maximum hard-disk throughput. The second goal was to optimize the implementation of sort and merge algorithms up to a level where the hard disks are kept busy at peak throughput.

Since we do not need any caching-, journaling- or even filing-capabilities of conventional filesystems, we implemented a primitive filesystem, which we call *AleSystem*, which provides fast and direct access to the hard disk. Each cluster node has one large unformatted data partition `sda1`, which is directly opened by the AleSystem using native Linux file I/O. Caching is deactivated by using the open flag O_DIRECT: data is not read for a long time and does not benefit from caching. Data is stored in portions of *ales*, all administrative information is stored directly in RAM. On sequential access, the throughput of the AleSystem reaches about 90 MB/s which is roughly the maximum that the hard disk allows for.

Since our cluster nodes are driven by quad-core CPUs, the speed of the computation is primarily based on multi-threaded parallelization. On the one side, the receive-/presort-/store-, on the other side, the load-/sort-/merge-/send-tasks are pipelined. At the current state of the implementation, we have several threads for sending/receiving data and for running the AleSystem. The core of the implementation is given by five threads which process the main computation. There are two threads which have the task to presort incoming data (one thread for each list). Furthermore, sorting is parallelized with two threads (one thread for each list) and for the merge task we have one more thread.

Given this task distribution, the size of necessary buffers can be defined. The micro-benchmarks show that bigger buffers generally lead to higher throughput. However, the sum of all buffer sizes is limited by the size of the available RAM. For the list parts we need 6 buffers, each 640MB, adding up to 3.75 GB. We need two times 2×8 network buffers for double-buffered send and receive, which results in 32 network buffers. To presort the entries double-buffered into 512 parts of two lists, we need 2048 ales. The size of network packets as well as ales must be a multiple of 5 bytes because the size of each entry is a multiple of 5 bytes. Therefore we chose a size of 5 MB for the network packets summing up to 160 MB and a size of 1.25 MB for the ales giving a memory demand of 2.5 GB. Over all, our implementation requires about 6.5 GB of RAM leaving enough space for the system and additional data as stack and the administrative data for the AleSystem.

Using our rough splitting of tasks to threads, we reach an average CPU usage of about 60% up to 80% peak. At the current optimization state, our average hard disk throughput is about 40 MB/s. The hard disk micro-benchmark (see figure 2) shows, that an average throughput between 45 MB/s and 50 MB/s should be feasible for packet sizes of 1.25 MB. Since sorting is the most complex task, we will further parallelize sorting to be able to use 100% of the CPU if the hard disk allows for higher data transfer. We expect that further parallelization of the sort task will increase CPU data throughput on sort up to about 50 MB/s. That should suffice for maximum hard disk throughput.

6 Results

Based on benchmarks of the current state of the implementation we will give estimates of how long the whole attack against FSB_{48} will take on the Coding and Cryptography Computer Cluster.

We have not yet implemented compression of list $L_{3,0}$ and the final merging step. The rest of the implementation described in Sections 4 and 5 is complete. In under 33 hours the implementation successfully generated one half tree and stored list $L_{3,1}$ on disk. This total time was divided as follows: List generation took 2.54 hours, the first sort-and-merge step took 9.57 hours, the second sort-and-merge step took 10.04 hours and the third sort-and-merge step took 10.77 hours.

As explained in Section 4 we expect 18.5 half-tree computations, accounting for a total time of 610.5 hours. At the end of computation of the left half-tree we have to compress the lists and send them to nodes 8 and 9 – this does not require any more computation, network traffic or disk access than sending, presorting and storing the data in the right half-tree, so we do not expect it to take any longer.

During the last merge step we have fewer computations to perform because we can skip dynamic recomputation of values in the left half-tree. As before, network traffic and hard disk access is no larger than for the other merging steps, so we expect this part to take no longer than 10 hours. Even if we take a pessimistic estimate of 12 hours, performing this last step 16.5 times will take just 198 hours, yielding an expected total time of the attack of 808.5 hours, i.e., 34 days.

6.1 Time-storage tradeoffs

As described in Section 4, the main restriction on the attack strategy was the total amount of background storage.

If we had 12.8 TB of storage at hand we could handle lists of size 2^{38}, again using the compression technique for the left half-tree. As described in Section 4 this would give us exactly one expected collision in the last merge step and thus reduce the number of required half-tree computations from 18.5 to 3.58, the expected number of last merge steps from 16 to 1.58 (see (2.1)). With a total storage of 20.48 TB we could omit half-tree compression and thus further reduce the expected number of half-tree computations to 2.58.

Increasing the size of the background storage even further would eventually allow to store list entry values alongside the positions and thus eliminate the need for dynamic recomputation. However, the performance of the attack is bottlenecked by hard-disk throughtput rather than CPU time so we don't expect any improvement through this measure.

On clusters with even less background storage the computation time will (asymptotically) increase by a factor of 16 with each halving of the storage size. For example a cluster with 3.2 TB of storage can only handle lists of size 2^{36}. The attack would then require 258.5 half-tree computations and 256.5 last merge steps.

Of course the time required for one half-tree computation depends on the amount of data. As long as the performance is bottlenecked mainly by hard-disk (or network) throughput the running time is linearly dependent on the amount of data, i.e. a Wagner computation involving 2 half-tree computations with lists of size 2^{38} is about 4.5 times as fast as a Wagner computation involving 18 half-tree computations with lists of size 2^{37}.

7 Scalability Analysis

The attack described in this paper including the variants discussed in Section 6 are much more expensive in terms of time and especially memory than a brute-force attack against the 48-bit hash function FSB_{48}.

This section gives estimates of the power of Wagner's attack against the larger versions of FSB, demonstrating that the FSB design overestimated the power of the attack. Table 1 gives the parameters of all FSB hash functions.

	n	w	r
FSB_{48}	3×2^{17}	24	192
FSB_{160}	5×2^{18}	80	640
FSB_{224}	7×2^{18}	112	896
FSB_{256}	2^{21}	128	1024
FSB_{384}	23×2^{16}	184	1472
FSB_{512}	31×2^{16}	248	1987

Table 1. Parameters of all FSB hash functions

A straightforward Wagner attack against FSB_{160} uses 16 lists of size 2^{127} containing elements with 632 bits. The entries of these lists are generated as xors of 10 columns from 5 blocks, yielding 2^{135} possiblities to generate the entries. Precomputation includes clamping of 8 bits. Each entry then requires 135 bits of storage so each list occupies more than 2^{131} bytes. For comparison, the largest currently available storage systems offer a few petabytes (2^{50} bytes) of storage.

To limit the amount of memory we can instead generate, e.g., 32 lists of size 2^{60}, where each list entry is the xor of 5 columns from 2.5 blocks, with 7 bits clamped during precomputation. Each list entry then requires 67 bits of storage.

Clamping 60 bits in each step leaves 273 bits uncontrolled so the Pollard variant of Wagner's algorithm (see Section 2.2) becomes more efficient than the plain attack. This attack generates 16 lists of size 2^{60}, containing entries which are the xor of 5 columns from 5 distinct blocks each. This gives us the possibility to clamp 10 bits through precomputation, leaving $B = 630$ bits for each entry on level 0.

The number of bytes sorted by this attack is approximately 2^{220} (see (2.2)). This is substantially faster than a brute-force collision attack on the compression function, but is clearly much slower than a brute-force collision attack on the hash function, and even slower than a brute-force *preimage* attack on the hash function.

Similar statements hold for the other full-size versions of FSB. Table 2 gives rough estimates for the time complexity of Wagner's attack without storage restriction and with storage restricted to a few hundred exabytes (2^{60} entries per list). These estimates only consider the number and size of lists being a power of 2 and the number of bits clamped in each level being the same. The estimates ignore the time complexity of precomputation.

Although fine-tuning the attacks might give small speedups compared to the estimates, it is clear that the compression function of FSB is oversized, assuming that Wagner's algorithm in a somewhat memory-restricted environment is the most efficient attack strategy.

	Number of lists	Size of lists	Bits per entry	Total storage	Total bytes sorted
FSB_{160}	16	2^{127}	632	$17 \cdot 2^{111}$	2^{127}
	16 (Pollard)	2^{60}	630	$9 \cdot 2^{64}$	2^{220}
FSB_{224}	16	2^{177}	884	$24 \cdot 2^{181}$	2^{177}
	16 (Pollard)	2^{60}	858	$13 \cdot 2^{64}$	2^{339}
FSB_{256}	16	2^{202}	1010	$27 \cdot 2^{206}$	2^{202}
	16 (Pollard)	2^{60}	972	$14 \cdot 2^{64}$	2^{382}
	32 (Pollard)	2^{56}	1024	$18 \cdot 2^{60}$	2^{400}
FSB_{384}	16	2^{291}	1453	$39 \cdot 2^{295}$	2^{291}
	32 (Pollard)	2^{60}	1467	$9 \cdot 2^{65}$	$2^{613.5}$
FSB_{512}	16	2^{393}	1962	$53 \cdot 2^{397}$	2^{393}
	32 (Pollard)	2^{60}	1956	$12 \cdot 2^{65}$	2^{858}

Table 2. Estimates for the cost of generalized birthday attacks against the compression function of FSB. Storage is measured in bytes. Time is modeled by the total number of bytes sorted.

References

1. Daniel Augot, Matthieu Finiasz, Philippe Gaborit, Stéphane Manuel, and Nicolas Sendrier. SHA-3 Proposal: FSB, 2009. `http://www-rocq.inria.fr/secret/CBCrypto/index.php?pg=fsb`.
2. Daniel Augot, Matthieu Finiasz, and Nicolas Sendrier. A fast provably secure cryptographic hash function, 2003. `http://eprint.iacr.org/2003/230`.
3. Daniel Augot, Matthieu Finiasz, and Nicolas Sendrier. A family of fast syndrome based cryptographic hash functions. In Ed Dawson and Serge Vaudenay, editors, *Mycrypt*, volume 3715 of *LNCS*, pages 64–83. Springer, 2005.
4. Paulo S. L. M. Barreto and Vincent Rijmen. The WHIRLPOOL Hashing Function. `http://www.larc.usp.br/~pbarreto/WhirlpoolPage.html`.
5. Daniel J. Bernstein. Better price-performance ratios for generalized birthday attacks. In *Workshop Record of SHARCS'07: Special-purpose Hardware for Attacking Cryptographic Systems*, 2007. `http://cr.yp.to/papers.html#genbday`.
6. Jean-Sébastien Coron and Antoine Joux. Cryptanalysis of a provably secure cryptographic hash function, 2004. `http://eprint.iacr.org/2004/013`.

7. Donald E. Knuth. *The Art of Computer Programming. Vol. 2, Seminumerical Algorithms*. Addison-Wesley Publishing Co., Reading, Mass., third edition, 1997. Addison-Wesley Series in Computer Science and Information Processing.
8. Lorenz Minder and Alistair Sinclair. The extended k-tree algorithm. In Claire Mathieu, editor, *SODA*, pages 586–595. SIAM, 2009.
9. Michael Naehrig, Christiane Peters, and Peter Schwabe. SHA-2 will soon retire. To appear. `http://cryptojedi.org/users/peter/index.shtml#retire`.
10. David Wagner. A generalized birthday problem (extended abstract). In Moti Yung, editor, *Advances in cryptology–CRYPTO 2002: 22nd annual international cryptology*, volume 2442 of *Lecture Notes in Computer Science*, pages 288–304. Springer, Berlin, 2002. See also newer version [11]. `http://www.cs.berkeley.edu/~daw/papers/genbday.html`.
11. David Wagner. A generalized birthday problem (extended abstract) (long version), 2002. See also older version [10]. `http://www.cs.berkeley.edu/~daw/papers/genbday.html`.