


Practical Preimages for Maraca

Sebastiaan Indestege
`sebastiaan.indestege@esat.kuleuven.be`

COSIC, ESAT/SCD, K.U. Leuven, Belgium

Dagstuhl
12 January 2009

Cryptographic Hash Functions




$$h : \{0, 1\}^* \mapsto \{0, 1\}^w$$

Desired properties

- Collision resistance, (Second) preimage resistance, ...
- Efficiently computable, *i.e.*, fast!

Cryptographic Hash Functions

Preimage Resistance




- Given y , “hard” to find x s.t. $h(x) = y$.
- Exhaustive search $\mathcal{O}(2^w)$



Maraca

- SHA-3 submission (*not selected for round 1*)
- Designer: Robert J. Jenkins Jr.
- Previous best attack: Canteaut and Naya-Plasencia, collision attack (2^{237})


Description of Maraca



- All sizes: **1024 bits**
- **Acc_i**: XOR of (up to) 3 (rotated) W_i 's from the past 46
- **Finalisation**: pad, 46 rounds **zero** W_i , 28 **perm**'s, truncate

Description of Maraca

The Maraca permutation



Maraca's S-box

- 8×8 bit S-box
- Three output bits are **linear** functions of input bits
- And not very non-linear otherwise, either...

Attacking Maraca

Maraca's S-box

- 8×8 bit S-box
- Three output bits are **linear** functions of input bits
- And not very non-linear otherwise, either...

Idea

- Impose **affine conditions** on the S-box inputs, to turn it into an **affine function**
 - 7 conditions: trivial, always works
 - Maraca S-box: can be done with **only 3 conditions**

Attacking Maraca

Exploiting the Maraca S-box

First attempt

- 1024 degrees of freedom per round
- 3 permutations \times 128 S-boxes \times 3 conditions = 1152
- **not enough** degrees of freedom...

Attacking Maraca

Exploiting the Maraca S-box

Second attempt

- Try to make the conditions **dependent**
(*then some conditions come for free*)

Attacking Maraca

Exploiting the Maraca S-box

Second attempt

- Try to make the conditions **dependent**
(*then some conditions come for free*)
- Treat the S-boxes differently for each permutation layer
 - ① 3 conditions per S-box (even+odd)
 - ② 4 conditions per S-box (even+odd)
 - ③ 1 (even) resp. 0 **conditions** (odd) per S-box

Attacking Maraca

Exploiting the Maraca S-box

Second attempt

- Try to make the conditions **dependent**
(then some conditions come for free)
- Treat the S-boxes differently for each permutation layer
 - ① 3 conditions per S-box (even+odd)
 - ② 4 conditions per S-box (even+odd)
 - ③ 1 (even) resp. **0 conditions** (odd) per S-box
- Total: 960 conditions per round
- Net result: **64 degrees of freedom per round**

Attacking Maraca

A preimage attack

Attack strategy (simplified)

- ① Strip off last 30 **perm**'s
- ② Restrict message to an **affine space**,
for which Maraca becomes an **affine function**
- ③ With enough rounds, we should get a **basis** for the output
- ④ Find **preimages** by XOR'ing precomputed messages

Attacking Maraca

A preimage attack

Practical issues

- 1 How to solve this linear system over $\text{GF}(2)$?

Attacking Maraca

A preimage attack

Practical issues

- 1 How to solve this linear system over $\text{GF}(2)$?
 - Size is more than $(750\,000)^2$?
 - That's 65.5 gigabytes just to store the matrix. . .
 - Gauss elimination: cubic time complexity, so $\approx 2^{58.5}$ operations. . .

Attacking Maraca

A preimage attack

Practical issues

- ① How to solve this linear system over $\text{GF}(2)$?
 - Size is more than $(750\,000)^2$?
 - That's 65.5 gigabytes just to store the matrix. . .
 - Gauss elimination: cubic time complexity, so $\approx 2^{58.5}$ operations. . .
- ② And then **contradictions** appear every once in a while. . .

Attacking Maraca

A preimage attack

Solutions

- ① Too large system of equations?
 - The linear system is quite structured
 - Simultaneously **build** and **solve** it!
 - And distribute it on a cluster of computers

Attacking Maraca

A preimage attack

Solutions

- ① Too large system of equations?
 - The linear system is quite structured
 - Simultaneously **build** and **solve** it!
 - And distribute it on a cluster of computers
- ② Contradictions?
 - Resort to a trivial linearisation of an S-box when contradictions occur
 - Costs some degrees of freedom, but only happens sporadically

Practical Results

Precomputation

- Used 32 AMD Opteron CPU's in a cluster with an Infiniband network
- Time: \pm 8 hours on 32 nodes (*i.e.* 10.7 CPU-days)
- Memory: \pm 20 GB of RAM (*distributed*)
- Result: a 94 MB data file

Practical Results

Precomputation

- Used 32 AMD Opteron CPU's in a cluster with an Infiniband network
- Time: ± 8 hours on 32 nodes (*i.e. 10.7 CPU-days*)
- Memory: ± 20 GB of RAM (*distributed*)
- Result: a 94 MB data file

Online phase

- Invert truncation arbitrarily (so we get multi-preimages)
- Compute a number of rounds backwards (166 perm's)
- XOR the correct messages from the data file
- Done! (takes just seconds on a PC, and always succeeds)

Conclusion

- Practical preimage attack on Maraca
- Fully implemented and verified

