Practical Free-Start Collision Attacks on 76-step SHA-1

Pierre Karpman X

[★]Inria and École polytechnique, France

[●]Nanyang Technological University, Singapore

Joint work with Thomas Peyrin and Marc Stevens

CWI, Amsterdam 2015–07–22

Title deconstruction

Practical

Free-Start

Collision

We can compute it

Not unlike a false start

As in f(x) = f(x')

Attacks

We're the baddies

on 76-step

Not quite the real thing

SHA-1

lot a cat 👶

Introduction

SHA-1 quickie

History of SHA-1 attacks

Our attack

Implementation

Results

Introduction

SHA-1 quickie

History of SHA-1 attacks

Our attack

Implementation

Results

Hash functions

Hash function

A (binary) hash function is a mapping $\mathcal{H}: \{0,1\}^* \to \{0,1\}^n$

- Many uses in crypto: hash n' sign, MAC constructions...
- It is a keyless primitive
- Sooo, what's a good hash function?

Three security notions (informal)

First preimage resistance

Given t, find m such that $\mathcal{H}(m) = t$ Best generic attack is in $\mathcal{O}(2^n)$

Second preimage resistance

Given m, find $m' \neq m$ such that $\mathcal{H}(m) = \mathcal{H}(m')$ Best generic attack is in $\mathcal{O}(2^n)$

Collision resistance

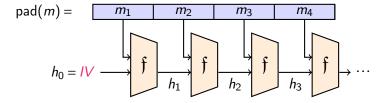
Find $m, m' \neq m$ such that $\mathcal{H}(m) = \mathcal{H}(m')$ Best generic attack is in $\mathcal{O}(2^{\frac{n}{2}})$

Merkle-Damgård construction

A domain of {0,1}* is annoying, so...

- **1** Start from a compression function $f: \{0,1\}^n \times \{0,1\}^b \to \{0,1\}^n$
- 2 Use a domain extender \approx $\Re(m_1||m_2||...||m_\ell) = \Re(\Re(...\Re(IV, m_1)...), m_\ell)$
- 3 Reduce the security of ${\mathcal H}$ to the one of ${\mathfrak f}$
 - ► $A(\mathcal{H}) \Rightarrow A(\mathfrak{f})$
 - ▶ $\neg A(f) \Rightarrow \neg A(\mathcal{H})$
 - $(A(\mathfrak{f}) \Rightarrow ???)$
 - Invalidates the security reduction, tho

MD in a picture



Additional security notions for MD

Semi-free-start collisions

The attacker may choose IV, but it must be the same for m and m'

Free-start preimages & collisions

No restrictions on IV whatsoever

Free-start preimages & collisions (variant)

Attack $\mathfrak f$ instead of $\mathcal H$

What did we do?

- ► This work: collisions on 76/80 steps of the compression function of SHA-1 (95% of SHA-1)
- And it's practical
- One inexpensive GPU is enough for fast results

Introduction

SHA-1 quickie

History of SHA-1 attacks

Our attack

Implementation

Results

The SHA-1 hash function

- Designed by the NSA in 1995 as a quick fix to SHA-0
- Part of the MD4 family
- ► Hash size is 160 bits ⇒ collision security should be 80 bits
- Message blocks are 512-bit long

SHA-1 round function

It's a 5-branch ARX Feistel

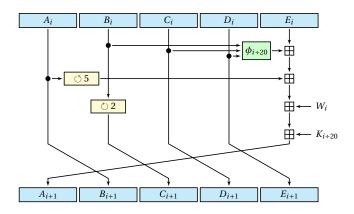
$$A_{i+1} = A_i^{\circlearrowleft 5} + \phi_{i \div 20} \big(A_{i-1}, A_{i-2}^{\circlearrowleft 2}, A_{i-3}^{\circlearrowleft 2} \big) + A_{i-4}^{\circlearrowleft 2} + W_i + K_{i \div 20}$$

with a linear message expansion:

$$W_{0...15}=M_{0...15},~W_{i\geq 16}=\left(W_{i-3}\oplus W_{i-8}\oplus W_{i-14}\oplus W_{i-16}\right)^{\circlearrowleft 1} \stackrel{\hookleftarrow \sim}{}$$
 The only difference between SHA-0 and SHA-1

80 steps in total

Round function in a picture



Introduction

SHA-1 quickie

History of SHA-1 attacks

Our attack

Implementation

Results

Wang collisions

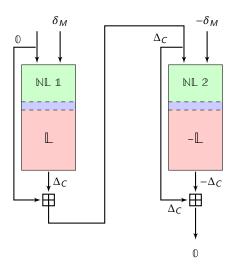
SHA-1 is not collision-resistant (Wang, Yin, Yu, 2005)

Differential collision attack

- Find a message difference that entails a good *linear* diff. path
- Construct a non-linear diff. path to bridge the IV with the linear path
- Use message modification to speed-up the attack
- Requires a pair of two-block messages

```
Attack complexity \equiv 2^{69}
Eventually improved to \equiv 2^{61} (Stevens, 2013)
```

Two-block attack in a picture



Preimage detour

SHA-1 is much more resistant to preimage attacks

- ► No attack on the full function
- ▶ Practical attacks up to $\lesssim 30$ steps ($\lesssim 37.5\%$ of SHA-1) (De Cannière & Rechberger, 2008)
- ► Theoretical attacks up to 62 steps (77.5% of SHA-1) (Espitau, Fouque, Karpman, 2015)

Our attack

Let's break stuff!



Why doing free-start again?

- Main reason is starting from a "middle" state + shift the message
- ► ⇒ Can use freedom in the message up to a later step
- ► ⇒ But no control on the IV value
- ➤ ⇒ Must ensure proper backward propagation

But then we need to...

- 1 Find a good linear part
- Construct a good shifted non-linear part
- 3 Find accelerating techniques

Let's do this for 76 steps!

- Best practical result is on 75 (we wanna beat 'em)
- ► (First step # with visible result for full SHA-1)

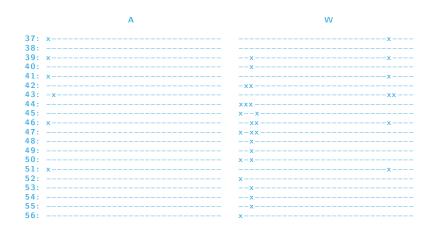
Linear part selection

Criteria:

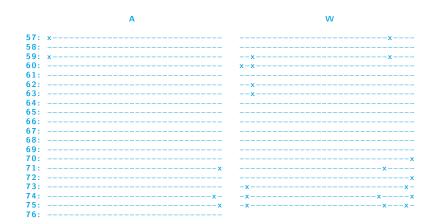
- High overall probability
- ► No (or few) differences in last five steps (= differences in IV)
- Few differences in early message words
- → Not many candidates

We picked II(55,0) (Manuel notation, 2011)

Linear path in a picture (part 1/2)



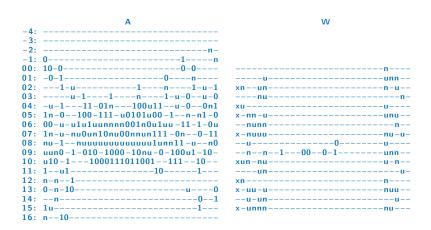
Linear path in a picture (part 2/2)



Non-linear part construction

- Start with prefix of high backward probability for the first 5 steps
- Use improved JLCA for the rest
- ► ⇒ Good overall path with "few" conditions (236 up to #36)

Non-linear path in a picture



Accelerating techniques

- Message modification: correct bad instances
- Neutral bits: generate more good instances when one's found
- We choose NBs because:
 - Easy to find
 - Easy to implement
 - Good parallelization potential (more of that later)

Neutral bits (with an offset)

- We start with an offset (remember?)
- ► ⇒ Use neutral bits with an offset too
- ► In our attack, offset = 6
 - free message words = W6...21 instead of W0...15
- ➤ ⇒ Must also consider backward propagation

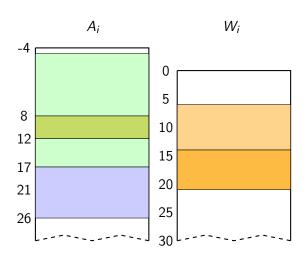
Our neutral bits

```
A18:
. . . . . . . . . . . . . . . . . xxxxx . . . . . . . . . . .
A19:
W14
W15
W16
A20:
A21:
A23 ·
A24:
W/20
A25 :
A26:
```

Let's sum up

- Initialize the state with an offset
- Initialize message words with an offset
- Use neutral bits with an offset
- → many neutral bits up to late steps (yay)
- → don't know the IV in advance (duh)
- ▶ Linear path \Rightarrow differences in the IV
- Everything done in one block
- → Attack on the compression function

Same thing in a picture



Introduction

SHA-1 quickie

History of SHA-1 attacks

Our attack

Implementation

Results

If it's practical you must run it

- Attack expected to be practical, but still expensive
- Why not using GPUs?
- One main challenge: how to deal with the branching?

Target platform

- ► Nvidia GTX-970
- ► Recent, high-end, good price/performance
- ► $13 \times 128 = 1664$ cores @ $\propto 1$ GHz
- High-level programming with CUDA
- ► Throughput for 32-bit arithmetic: all 1/cycle/core except ⊙
- ► ≈ S\$ 500

Architecture imperatives

- Execution is bundled in warps of 32 threads
- Control-flow divergence is serialized ⇒ minimize branching
- ► Hide latency by grouping threads into larger blocks
- ► But careful about register / memory usage

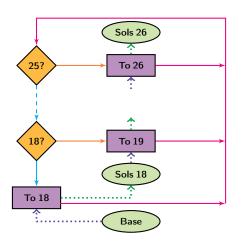
Our snippet-based approach

- 1 Store partial solutions up to some step in shared buffers
- 2 Every thread of a block loads one solution
- 3 ... tries all neutral bits for this step
- 4 ... stores successful candidates in next step buffer

Our snippet-based approach (cont.)

- ► Base solutions up to #17 generated on CPU
- ▶ Use neutral bits up to #26 on GPU
- ► Further checks up to #56 on GPU
- Final collision check on CPU

Snippets in a picture



Results

GPU results

- ► Hardware: one GTX-970 (S\$500)
- ► One partial solution up to #56 per minute on average
- ▶ \Rightarrow Expected time to find a collision \lesssim 5 days
- ► Complexity $\equiv 2^{50.25}$ SHA-1 compression function

GPU v. CPU

- ▶ On one CPU core @ 3.2 GHz, the attack takes \approx 606 days
- ► ⇒ One GPU ≡ 140 cores
- ► (To compare with = 40 (Grechnikov & Adinetz, 2011))
- For raw SHA-1 computations, ratio is 320
- \rightarrow Lose only $\times 2.3$ from the branching (not bad)

For more details...

Pierre Karpman, Thomas Peyrin, and Marc Stevens: Practical Free-Start Collision Attacks on 76-step SHA-1, CRYPTO 2015

Eprint 2015/530