

Generic collision attacks on hash-functions

Chris Mitchell

Royal Holloway, University of London

1

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

2

Hash-functions

- Cryptographic hash-functions are a key element of all widely used signature schemes.
- They are also used as a component of many asymmetric encryption schemes.
- Such functions map an arbitrary length bit string to a fixed length bit string (e.g. of length 128, 160, 224 or 256 bits).
- If the output length is n bits, then we refer to an n -bit hash-function.

3

Examples

- Examples of widely used hash-functions include:
 - RIPEMD-128 and RIPEMD-160 (standardised in ISO/IEC 10118-3);
 - MD5 (now discredited);
 - SHA-1 (a NIST and ISO/IEC standard, but now looking a little shaky);
 - SHA-256 (NIST and ISO/IEC standardised).

4

Necessary properties

- A cryptographic hash-function h is normally expected to have the following properties:
 - *pre-image resistance (the 'one way' property)*: given an arbitrary n -bit block y it is computationally infeasible to find an input x such that $h(x) = y$;
 - *second pre-image resistance*: given an arbitrary bit string x , it is computationally infeasible to find a distinct bit-string $x' \neq x$ such that $h(x') = h(x)$.
- In many cases (e.g. when h is used as part of a signature scheme) the following property is also required:
 - *collision resistance*: it is computationally infeasible to find two distinct bit-strings x and x' ($x \neq x'$) such that $h(x) = h(x')$.

5

A trivial collision attack

- The Birthday Paradox tells us that, if we choose around \sqrt{m} elements from a set of size m (with replacement), then there is a good chance of choosing one element twice.
- Hence if we compute $h(x)$ for around $\sqrt{(2^n)} = 2^{n/2}$ messages, then there is a good chance of finding a collision.
- Hence, given $2^{n/2}$ computations and $2^{n/2}$ storage, we can find a collision for any n -bit hash-function.

6

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

7

Background

- A memoryless technique for finding hash-function collisions was proposed by Quisquater and Delescaille (1989) – we refer to this as the QD attack.
- This was refined by van Oorschot and Wiener (1994) to produce a technique which is both parallelisable and produces collisions for ‘meaningful’ messages [see also section 9.7.1 of the HAC].
- This requires $O(2^{n/2})$ computational steps and trivial memory for an n -bit hash-function.


8

Functions and functional graphs

- Suppose $f: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}$.
- Then f defines a directed graph (the *functional graph*) with:
 - m vertices (the values $1, 2, \dots, m$); and
 - m directed edges, where $a \rightarrow b$ is an edge if and only if $f(a) = b$.
- This graph is made up of a number of connected sub-graphs (components), where each component will contain a single directed cycle and some directed trees attached to (pointing to) the cycle.

9

Example



The above example of a functional graph for $m=13$ is taken from the HAC.

10

Random functions

- Suppose $f: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}$ is randomly chosen from the class of all such functions. [We call f a random function].
- The expected length of a cycle in the functional graph is $\sqrt{(\pi m/8)}$, and the expected maximum cycle length is around $1.7\sqrt{m}$.
- The expected length of a 'tail' (i.e. a path from a point to a cycle) is also $\sqrt{(\pi m/8)}$, and the expected maximum tail length is around $0.8\sqrt{m}$.
- The largest connected component will contain $O(2m/3)$ vertices, i.e. the majority of the vertices will be in one large connected component.

11

The QD attack I

- The QD attack involves observing that, for an n -bit hash-function h , we can consider the action of h on the set of all n -bit blocks:
- We can model this restricted version of h as a random function,

$$h: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}, \text{ where } m = 2^n.$$
- Hence, if we choose a random block x , then the sequence:

$$x, h(x), h^2(x), h^3(x), \dots$$
 will define a path in the functional graph of (the restricted version of) h .

12

QD attack II

- This path will eventually reach a cycle, and will then repeatedly go round the cycle.
- We just need to find a way to notice when the path round the cycle has been completed.
- This can be done with *distinguished points*.
- That is, call an n -bit block x a *distinguished block* if the leftmost $n/2-s$ bits of x are all zeros (for some smallish s , e.g. $s=8$).

(Actually, any fixed pattern of $n/2-s$ bits will do equally well – however, looking for zeros is usually easiest).

13

QD attack III

- The attack proceeds as follows:

choose a random block x ;

for $i = 1, 2, \dots$

compute $h^i(x)$;

if $h^i(x)$ is a distinguished block **then**

if ($h^i(x) = h^j(x)$ for any recorded *distinguished pair*
 $(j, h^j(x)), j < i$) **then stop**;

record the *distinguished pair* $(i, h^i(x))$

14

QD attack IV

- The algorithm will give a pair (i, j) such that $h^i(x) = h^j(x)$ and $1 \leq j < i$.
- This means that, either:
 - $x = h^{i-j}(x)$ [not a collision but an interesting fact in itself, since we have found a pre-image for x], or
 - for some $t < j$: $h^{i-t}(x) = h^{j-t}(x)$ & $h^{i-t-1}(x) \neq h^{j-t-1}(x)$, i.e. we have a collision.
- If the first event occurs, start again. The first event will occur if and only if x is on a cycle (probability very small).

15

QD attack complexity

- The storage requirement is clearly trivial, i.e. $O(2^s)$.
- Each cycle will probably contain at least one distinguished block (choose s large enough to ensure this, bearing in mind the average cycle length).
- Since the expected path length and expected cycle size are both $\sqrt{(\pi m/8)}$, where $m=2^n$, the number of computational steps is clearly $O(2^{n/2})$.

16

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

17

Iterated random functions

- We now need to consider some further properties of random functions.
- In particular, we are interested in the properties of the image set of functions $f: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}$.

18

Image sets I

Theorem Suppose f_1, f_2, \dots are independent random functions, where

$$f_i: \{1, 2, \dots, m\} \rightarrow \{1, 2, \dots, m\}.$$

Then the expected size of the image set of the compound function $f_t \bullet f_{t-1} \bullet \dots \bullet f_1$ (i.e. the function made up of f_1 , followed by f_2, \dots , followed by f_t) is bounded above by $2m/t$ (for large image sets).

Corollary If one iteratively applies 2^w random functions to a set, then the size of the image set (for a large image set) is around $1/2^{w-1}$ times the size of the domain.

19

Image sets II

The expected sizes of the image sets for small t (and large m) are as follows:

t (number of functions)	size of image set
1	$0.63m$
2	$0.47m$
3	$0.37m$
4	$0.31m$
5	$0.27m$

20

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

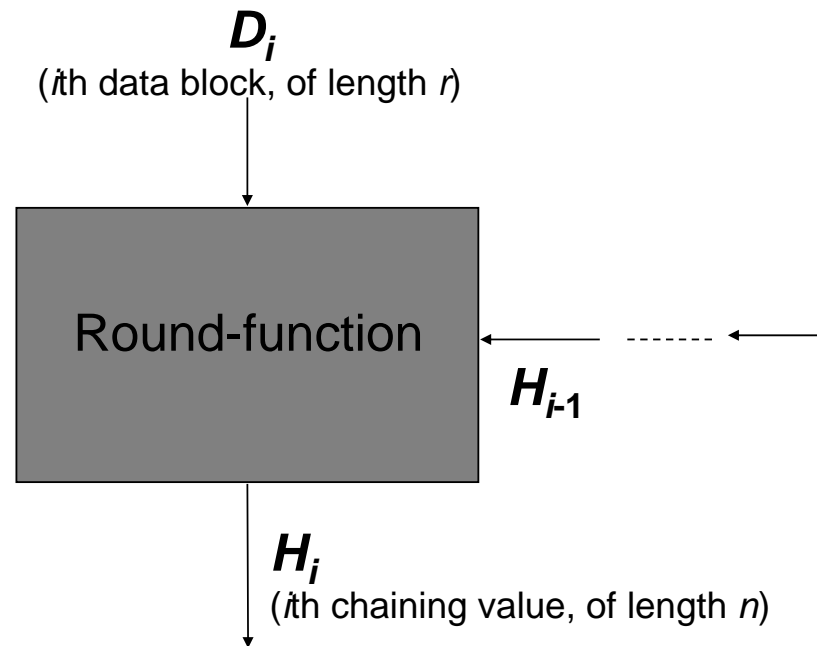
21

The Merkle meta-method

- All the widely used hash-functions are based on Merkle's meta-method.
- That is, they involve:
 1. padding the data
 2. splitting the padded data into blocks of a fixed length (say r);
 3. setting an initial chaining value to a fixed value;
 4. iteratively applying a round function (or compression function), which takes an n -bit chaining value and an r -bit data block as input, and outputs an n -bit chaining value.
- The hash-code is the final chaining value (of length n bits)

22

Merkle meta-method



23

Example parameters

- MD5, SHA-1 and RIPEMD-160 have $r=512$ and $n=160$.
- SHA-256 has $r=512$ and $n=256$.

24

The length block

- All the well-known hash-functions have a similar padding scheme.
- This involves adding up to $r-1$ zeros to the end of the message to make sure the length is a multiple of r .
- Then an additional r -bit block is added to the end, which encodes the original message length.

25

Effect of the length block

- When computing $h^i(x)$ (for $i=1,2,\dots$) in the QD attack, the input to h has a fixed length (n bits).
- Hence, in each case, the padded message being hashed will contain two r -bit blocks (assuming $n \leq r$), the first of which will contain the n -bit input followed by $r-n$ zeros, and the second of which will be fixed (the length block).
- Since the initial chaining value is fixed, the first iteration of the round function can be treated as a random function from a set of size $m=2^n$ onto itself.
- Since the r -bit length block is fixed, the second iteration of the round function can also be thought of as a random function.
- That is, the values of $h^i(x)$ for $i=1,2,\dots$, will be drawn from a set of size 0.47×2^r (and not 0.63×2^r).

26

QD attack effects I

- This will mean that the expected size of the cycle will be smaller than the previous estimates would suggest, which means that the attack is actually more effective than would otherwise be the case.
- Indeed, suppose the attacker actually computed $g^i(x)$ instead of $h^i(x)$, where:

$$g(y) = h(y || 0^{r-n} || b_1 || b_2 || \dots || b_v)$$
for some v , where 0^{r-n} denotes $r-n$ zeros, b_i is a distinct r -bit block for every i , and $||$ denotes concatenation.
- Thus computing g involves computing the round function a total of $v+2$ times, which can be modelled as the iterative performance of $v+2$ independent random functions from a set of size $m=2^n$ onto itself.

27

QD attack effects II

- That is, from Theorem A, the output of $g^i(x)$ will be drawn from a set of size at most $2^{n+1}/(v+2)$.
- Suppose $v = 2^u - 2$.
- In this case (by heuristic arguments) the expected size of the cycle which the computation of $g^i(x)$ will proceed round will be $O(2^{(n+1-u)/2})$.
- That is, the complexity of the attack has now been reduced by a factor of $2^{(u-1)/2}$, i.e. a collision can be found using $O(2^{(n+1-u)/2})$ hash computations.

28

QD attack effects III

- However, we have cheated!
- The reduction in the number of hash computations is at the cost of hashing messages whose padded versions will contain 2^u r -bit blocks, so the overall attack complexity will actually increase as u increases.
- This means it is necessary to be careful about assessing attack complexity solely in terms of hash function computations.

29

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

30

Applications of iterated hash-functions

- Multiple iterations of hash-functions have been proposed for use in a variety of applications.
- For example:
 - Lamport proposed the use of hash-chains for user authentication;
 - hash-functions are used for random bit generation, by iteratively hashing a random 'seed';
 - hash-functions are often used as a simple and quick method of detecting accidental errors in a data string D , by storing $h(D)$ with D (for a hash-function h).

31

Multiple iterations

- If an m -bit hash-function is iterated a large number of times, then we know that the size of the domain will be $O(\sqrt{m})$.
- This means that there is a potential danger of loss of up to half of the entropy.
- The effect will, to a large extent, depend on the hash-function length, and the way that the iterations are computed.

32

Example scenario I

- Suppose a series of 64-bit values are required (e.g. as pseudo-random values).
- Suppose a 128-bit hash-function is iteratively applied to a 128-bit seed.
- Could take the first 64 bits of the output, before rehashing.
- However, the full 128 bits must be rehashed, and not just 64.
- If not, then after $O(2^{32})$ hash applications, the sequence of bits output will cycle with period $O(2^{32})$.
- That is only ≈ 32 bits of entropy will remain.

33

Example scenario II

- However, if the full 128 bits are hashed every time, then, after 2^{32} iterations, around 96 bits of entropy will remain.
- This is likely to be satisfactory for most purposes.
- Of course, after 2^{64} iterations, 'only' ≈ 64 bits of entropy will remain.

34

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

35

Theory

- Some additional work is required to properly compute the expected sizes of the cycles and tails for the functional graph of compound random functions.
- It would also help to know the expected size of the largest component of the functional graph for a 'compound' random function.
- This should then yield theoretical results on the performance of variants of the QD attack.

36

Practical work I

- It would be of considerable interest to conduct some practical experiments to test whether the theoretical results on multiple iterations of hash functions actually hold for practical schemes.
- ‘Toy’ experiments could be carried out with ‘shortened’ versions of well-known functions.
- For example, one can define SHA-1_k to be the k -bit hash-function derived by taking the first k bits of a SHA-1 hash.
- It would be interesting to see how easy it is to find collisions for SHA-1_k for modest values of k (e.g. $k=80, 90, 100$).

37

Practical work II

- These experiments could be conducted for:
 - the QD attack;
 - the van Oorschot-Wiener attack.

38

Agenda

1. Hash-functions and collision attacks
2. Memoryless strategy for finding collisions
3. Properties of random functions
4. Impact on memoryless strategy
5. Other observations
6. Future work
7. Acknowledgements

39

Acknowledgements

- This is joint work with Sean Murphy.
- Thanks also to Laurence O'Toole and Peter Wild for useful discussions.

40