Hash Functions



Federico Pintore¹

¹Mathematical Institute, University of Oxford

Outline





Outline



Pash Functions: Additional Applications

Bash Functions: Constructions

 Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π.

- Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π.
- Instead of using cryptosystems that have no proofs at all, an alternative approach is to "idealise" the cryptographic hash functions!

- Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π.
- Instead of using cryptosystems that have no proofs at all, an alternative approach is to "idealise" the cryptographic hash functions!
- Let us consider a hash function that is *truly random*.

- Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π.
- Instead of using cryptosystems that have no proofs at all, an alternative approach is to "idealise" the cryptographic hash functions!
- Let us consider a hash function that is *truly random*.
- Additionally, assume that this random function is public, and it can be evaluated only querying it as an oracle (or a black box)!

- Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π.
- Instead of using cryptosystems that have no proofs at all, an alternative approach is to "idealise" the cryptographic hash functions!
- Let us consider a hash function that is *truly random*.
- Additionally, assume that this random function is public, and it can be evaluated only querying it as an oracle (or a black box)!
- If you don't idealise hash functions in your proofs, then your cryptosystem is said to be secure in the *standard model*, otherwise, it is *only* secure in the random oracle model (ROM).

- Let Π be a cryptosystem using a hash function H. Sometimes it is NOT enough for H to be collision resistant/preimage resistant to be able to write a security proof of Π .
- Instead of using cryptosystems that have no proofs at all, an alternative approach is to "idealise" the cryptographic hash functions!
- Let us consider a hash function that is *truly random*.
- Additionally, assume that this random function is public, and it can be evaluated only querying it as an oracle (or a black box)!
- If you don't idealise hash functions in your proofs, then your cryptosystem is said to be secure in the *standard model*, otherwise, it is *only* secure in the random oracle model (ROM).
- In the real world, each ideal hash function is instantiated by an *appropriate* hash function.

What do we mean by appropriate hash functions?

- What do we mean by appropriate hash functions?
- No clear definition!

- What do we mean by appropriate hash functions?
- No clear definition!
- Concrete hash functions are deterministic and fixed, they *cannot* behave like random functions!

- What do we mean by appropriate hash functions?
- No clear definition!
- Concrete hash functions are deterministic and fixed, they *cannot* behave like random functions!
- What does a proof in the random oracle buy us?

- What do we mean by appropriate hash functions?
- No clear definition!
- Concrete hash functions are deterministic and fixed, they *cannot* behave like random functions!
- What does a proof in the random oracle buy us?
- Perhaps, the scheme doesn't have "inherent design flaws"! (the only possible attacks are those due to weaknesses in the used hash functions)

- What do we mean by appropriate hash functions?
- No clear definition!
- Concrete hash functions are deterministic and fixed, they *cannot* behave like random functions!
- What does a proof in the random oracle buy us?
- Perhaps, the scheme doesn't have "inherent design flaws"! (the only possible attacks are those due to weaknesses in the used hash functions)
- Why is it widely used?

- What do we mean by appropriate hash functions?
- No clear definition!
- Concrete hash functions are deterministic and fixed, they *cannot* behave like random functions!
- What does a proof in the random oracle buy us?
- Perhaps, the scheme doesn't have "inherent design flaws"! (the only possible attacks are those due to weaknesses in the used hash functions)
- Why is it widely used?
- So far, there have been no successful *real-world* attacks on *real-world* schemes that are proven secure in the ROM. Additionally, schemes that are proven secure in the ROM are usually efficient.

- In ROM **security definitions**, the probability is taken over the random choice of *H*, whereas in the real world, you instantiate *H* by a deterministic function.
- In ROM **security proofs**, the adversary needs to query *H*, instead of execute *H* itself.

- In ROM **security definitions**, the probability is taken over the random choice of *H*, whereas in the real world, you instantiate *H* by a deterministic function.
- In ROM **security proofs**, the adversary needs to query *H*, instead of execute *H* itself.
- If *x* has not been queried yet to *H*, then the value *H*(*x*) is still considered uniform.

- In ROM **security definitions**, the probability is taken over the random choice of *H*, whereas in the real world, you instantiate *H* by a deterministic function.
- In ROM **security proofs**, the adversary needs to query *H*, instead of execute *H* itself.
- If *x* has not been queried yet to *H*, then the value *H*(*x*) is still considered uniform.
- **Extractability:** When A queries x to H in a proof by reduction, the challenger learns x.

- In ROM **security definitions**, the probability is taken over the random choice of *H*, whereas in the real world, you instantiate *H* by a deterministic function.
- In ROM **security proofs**, the adversary needs to query *H*, instead of execute *H* itself.
- If *x* has not been queried yet to *H*, then the value *H*(*x*) is still considered uniform.
- **Extractability:** When A queries x to H in a proof by reduction, the challenger learns x.
- **Programmability:** In a proof by reduction, the challenger sets the (uniformly distributed) values of *H*(*x_i*) to answer the adversary's queries!

Outline



Pash Functions: Additional Applications

Bash Functions: Constructions

7 of 34

- **Fingerprinting**: The digest *H*(*x*) of a file *x* (which could be a virus) acts as a fingerprint/identifier of the file
- **Deduplication**: Particularly important in cloud storage. You send a hash of the file you want to store to the service (e.g. DropBox); they check if the file already exists, in that case they don't need to store it again, a pointer to it would be enough.

- **Merkle Trees**: Suppose you have *n* files x_1, \dots, x_n , where *n* is a power of 2. Instead of hashing them all, i.e. $H(x_1, \dots, x_n)$, Ralph Merkle proposed a solution that works as follows:
 - Compute $h_{1,2} \leftarrow H(x_1, x_2), \cdots, h_{n-1,n} \leftarrow H(x_{n-1}, x_n)$.
 - Compute
 - $h_{1,2,3,4} \leftarrow H(h_{1,2},h_{3,4}), \cdots, h_{n-3,n-2,n-1,n} \leftarrow H(h_{n-3,n-2},h_{n-1,n})$
 - The process is iterated, until the root $h_{1,...,n}$ is computed.
- Merkle Tree can be thought of as an alternative to the Merkle Damgård transform to extend the domain of fixed-length collision-resistant hash functions.
- Its drawback: it is not collision-resistant if *n* is not fixed!

- **Password Hashing**: a hash of the password is usually stored instead of the password itself.
- What if the password is chosen from a small space?
- Is it enough to have a preimage resistance hash function H?
- ONLY if you are sampling your password *uniformly* from a large space, i.e. $\{0,1\}^n$ with suitable *n*.
- In practice: if your password is a random combination of 8 alphanumeric characters, the space is $S = 62^8 \approx 2^{47.6}$.
- There is an attack (that requires some preprocessing) which only uses time and space $N^{2/3} \approx 2^{32}$.
- There are mechanisms that can be used to mitigate this threat (adding a long random *salt*, etc.).

Commitment Schemes

- A commitment scheme allows a party to commit to a value *v* by producing a commitment on it.
- The commitment keeps the value *v* hidden, i.e. it reveals nothing about it. This property is called *hiding*.
- The party cannot change it later on, i.e. it cannot open to two different values *v*₁, *v*₂. This property is called *binding*.
- Think of it as a sealed envelope!
- It is a very important cryptographic tool, which can be built using hash functions!

11 of 34

Commitments Schemes

Definition

A commitment scheme consists of two algorithms, KeyGen and Commit, defined as follows

- KeyGen(*n*) : *it outputs public parameters* p;
- Commit(p, m ∈ {0, 1}ⁿ, r ∈ {0, 1}ⁿ) : it takes the public parameters, a message m and a random value r, and it outputs com_(m)

The sender can at anytime reveal the message *m* to the receiver by sending (m, r). The receiver can easily verify the correctness of the sender's claim by testing $\text{Commit}(\mathbf{p}, m, r) \stackrel{?}{=} \text{com}_{(m)}$

Informally speaking, a commitment scheme is secure if it is both binding and hiding.

12 of 34

Commitments Schemes

- Suppose that we have a hash function that is modelled as a random oracle. We can define a commitment scheme where Commit ← H(m||r).
- Binding: follows from the fact that the hash function is collision-resistant.
- Hiding: follows from the fact that r is chosen uniformly from $\{0,1\}^n$.
- There are other commitment schemes that don't assume the existence of a random oracle, i.e. they are proven secure in the standard model.

Outline



2 Hash Functions: Additional Applications

Bash Functions: Constructions

Hash functions are commonly constructed in two steps:

- a fixed-length collision-resistant hash function *h* is constructed;
- to allow for arbitrary-length inputs, some techniques may be applied to extend *h* (e.g. Merkle-Damgård transform).

Hash Functions From Block Ciphers

- We can use a *special* block cipher to build a fixed-length collision-resistant hash function *h*.
- There are different ways to do this.
- Davies-Meyer method is the most common one.
- Given a block cipher with *n* as key-length and *ℓ* as block-length,
 h is defined as follows:

$$\begin{split} h: \{0,1\}^{n+\ell} \to \{0,1\}^\ell \\ (k,x) &\mapsto h(k,x) = F_k(x) \oplus x. \end{split}$$

Hash Functions From Block Ciphers

- The assumption that *F* is a strong pseudo-random permutation is NOT enough to prove collision resistance of *h*.
- We need to rely on something similar to the random oracle model's idea.
- *F* is modeled as an ideal cipher.
- This means having a public oracle for evaluating a random keyed permutation *F* : {0,1}ⁿ × {0,1}^ℓ → {0,1}^ℓ, and its inverse *F*⁻¹.
- Each party has to query the oracle to compute F(k, x) or $F^{-1}(k, y)$ (similar to the ROM).

- Designed in 1991. It has output length equal to 128.
- It is totally broken, collisions can be found in less than a minute on a PC!
- Given a message:
 - the bit 1 is appended,
 - the bit 0 is appended until the bit string has a length congruent to 448 modulo 512,
 - the original length of the message is appended, encoded as a 64-bit string.
- The padded message is divided into blocks of length 512. Each block is divided into 16 chunks *M_j* of 32 bits each.

- In processing a block, 64 operations are executed.
- They are grouped in four rounds, each of 16 operations.
- We have 4 non-linear functions, *F*, *G*, *H*, *I*, one for each round.
- *M_g* denotes a 32-bit chunk, where *g* is a function of the operation's index *i*.
- *K_i* denotes a 32-bit constant, different for each operation.
- ««, denotes a left bit rotation by s places; s varies for each operation.
- Addition is done modulo 2³².

19 of 34

MD5



Figure: One MD5 operation in the first round (From wikipedia)

MD5

It uses 4 functions. Each of them takes as input three 32-bit string and generate as output one 32-bit string:

• $\mathbf{F}(B, C, D) = (B \land C) \lor (\neg B \land D)$

•
$$\mathbf{G}(B,C,D) = (B \wedge D) \vee (C \wedge \neg D)$$

•
$$\mathbf{H}(B,C,D) = B \oplus C \oplus D$$

•
$$\mathbf{I}(B, C, D) = C \oplus (B \lor \neg D)$$

If *i* is the operations' counter, then g = i when $i \in \{0, ..., 15\}$, $g = 5i + 1 \pmod{16}$ when $i \in \{16, ..., 31\}$, $g = 3i + 1 \pmod{16}$ when $i \in \{32, ..., 47\}$, $g = 7i + 1 \pmod{16}$ when $i \in \{48, ..., 63\}$.

Performing compression step



From Cryptool software.

Secure Hash Algorithms: SHA-1 and SHA-2

- A family of cryptographic hash functions standardized by NIST.
- First, they all use Davies-Meyer construction to build a fixed-length collision-resistant hash function from a block cipher.
- The block ciphers were specifically designed for this purpose.
- The block cipher SHACAL-1 with 160-bit block length for SHA1. The block cipher SHACAL-2 with 256-bit block length for SHA2. The key length is 512-bit in both of them.
- Second, they are extended to handle arbitrary-length inputs using Merkle-Damgård transform.

SHA-1

- SHA-1 was introduced in 1995. It has 160-bit output length.
- After the padding, each 512-bit block is extended into eighty 32-bit chunks. For each chunk *W*_{*t*}, an operation is executed.
- In theory, collisions can be found significantly better that the birthday attack, i.e. much less 2⁸⁰ hash function evaluations.
- In practice, no collisions of this type. But highly recommended to move to SHA-2 (or perhaps to SHA-3). SHAttered- Move now to SHA-2!
- Very recent attack (see references at the last slide).
- Example that shows the steps of SHA-1: http://www.metamorphosite.com/ one-way-hash-encryption-shal-data-software

SHA-1



Figure: From Wikipedia

- Similar to MD5 and SHA-1. It has 256-bit output length.
- After the padding, each of the *N* 512-bit blocks $M^{(1)}, \ldots, M^{(N)}$ is extended into sixty-four 32-bit chunks. For each chunk W_t , an operation is executed.
- Fix the initial hash values $A^{(0)}, \dots, H^{(0)}$ with the fractional parts of the square roots of the first eight primes.
- Compute $Hash^{(i)} = Hash^{(i-1)} + C(M^{(i)}, Hash^{(i-1)})$ where *C* is the fixed-length collision-resistant hash function
- Output $Hash^{(N)}$ as the hash of the message *M*.
- For detailed description see: http://www.iwar.org.uk/ comsec/resources/cipher/sha256-384-512.pdf



27 of 34

The logical functions are as follows:

- $Ch(E, F, G) = (E \land F) \oplus (\neg E \land G)$
- $Ma(A, B, C) = (A \land B) \oplus (A \land C) \oplus (B \land C)$
- $\Sigma_0(A) = (A \gg 2) \oplus (A \gg 13) \oplus (A \gg 22)$
- $\Sigma_1(E) = (E \ggg 6) \oplus (E \ggg 11) \oplus (E \ggg 25)$

The constant words, K_0, \dots, K_{63} are the first 32 bits of the fractional parts of the cube roots of the first sixty-four primes.

SHA-3 (Keccak)

- In 2012, Keccak was announced as the winner of the NIST competition (was called SHA-3) to design a new cryptographic hash function.
- All candidates were of 256- and 512-bit output length.
- Its structure is different from SHA-1 and SHA-2.
- it uses an *unkeyed* permutation with 1600-bit block length!
- For instance, Davies-Meyer construction uses a keyed permutation
- it doesn't use Merkle-Damgård to extend the fixed-length collision-resistant hash function to deal with arbitrary-length inputs.
- *Sponge construction* is the new approach that it uses instead of Merkle-Damgård.

29 of 34

SHA-3 (Keccak)

- The permutation *f* operates on blocks of length *b*.
- A rate *r* is fixed.
- The output length is *d*.
- The initial message is padded using pad, which returns a string of length *n* · *r*. We divide the string into *n* blocks *M*₀, ..., *M*_{*n*-1} of length *r*.
- Initialise the state *S* as a string of *b* zeros.
- For each M_i, extend it at the end with c = b − r zeros; xor the obtained string with S; apply the permutation f.
- Initialise *Z* to be the empty string.
- While the length of *Z* is less than *d* append the first *r* bits of *S* to *Z* and apply *f* to *S* to obtain a new state.
- Truncate Z to d bits.

Keccak- Sponge Function



sponge

Complete description:

http://sponge.noekeon.org/CSF-0.1.pdf 31 of 34

Further Reading (1)

Mihir Bellare and Phillip Rogaway.

Random oracles are practical: A paradigm for designing efficient protocols.

In *Proceedings of the 1st ACM conference on Computer and communications security*, pages 62–73. ACM, 1993.

 Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche.
 Keccak sponge function family main document. Submission to NIST (Round 2), 3:30, 2009.

Further Reading (2)

 Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya.
 Merkle-damgård revisited: How to construct a hash function.

In *Advances in Cryptology–CRYPTO 2005*, pages 430–448. Springer, 2005.

- Pierre Karpman, Thomas Peyrin, and Marc Stevens.
 Practical free-start collision attacks on 76-step sha-1.
 In Advances in Cryptology–CRYPTO 2015, pages 623–642.
 Springer, 2015.
- Neal Koblitz and Alfred J Menezes.
 The random oracle model: a twenty-year retrospective.
 Designs, Codes and Cryptography, pages 1–24, 2015.

- Alfred J Menezes, Paul C Van Oorschot, and Scott A Vanstone. Handbook of applied cryptography. CRC press, 1996.
- Marc Stevens.

New collision attacks on sha-1 based on optimal joint local-collision analysis.

In *Advances in Cryptology–EUROCRYPT 2013*, pages 245–261. Springer, 2013.