

Cryptography

Lecture 16

Announcements

- HW6 due 4/17 *extended*

Agenda

- Practical constructions of Collision-Resistant Hash Functions (K/L 6.3)
- New Unit: Number Theory!

Preliminaries

- How much security can we hope for from a CRHF that outputs ℓ bits?
- Discuss the “**birthday bound**”
 - No matter what function is used, collisions can be found with high probability after making $2^{\ell/2}$ queries.

Weaker Notions of Security

- Second preimage or target collision resistance: Given s and a uniform x it is infeasible for a ppt adversary to find $x' \neq x$ such that $H^s(x') = H^s(x)$.
- Preimage resistance: Given s and uniform y it is infeasible for a ppt adversary to find a value x such that $H^s(x) = y$.

Hash Functions From Block Ciphers

- Hash functions are generally constructed in two steps:
 - First, a compression function (fixed-length hash function) h is designed
 - Next, some mechanism (e.g. Merkle-Damgard) is used to extend h so as to handle arbitrary input lengths
- We will focus on the first step

Hash Functions From Block Ciphers

- Davies-Meyer construction:
 - F is a block-cipher with n -bit key and ℓ -bit block length.




FIGURE 6.10: The Davies-Meyer construction.

- Above forms a compression function from $n + \ell$ bits to n bits.

Security Analysis

- We do not know how to prove collision-resistance of the compression function based on the assumption that F is a strong PRP.
- Requires stronger assumption that F behaves like an **ideal cipher**.
 - Like a truly random permutation, except can query oracle on **different keys**.
 - Each key $k \in \{0, 1\}^n$ specifies an independent, uniform permutation $F(k, \cdot)$ on ℓ -bit strings.

Security Analysis

- Theorem: If F is modeled as an ideal cipher, then the Davies-Meyer construction yields a collision-resistant compression function. Concretely, any attacker making $q < 2^{\ell/2}$ queries to its ideal-cipher oracles finds a collision with probability at most $q^2/2^\ell$.

MD5

- 128-bit output length.
- Designed in 1991, and for several years was believed to be collision-resistant. Over a period of several years, various weaknesses began to be found in MD5 but
- these did not appear to lead to any easy way to find collisions.
- In 2004 a team of Chinese cryptanalysts presented a new method for finding
- collisions in MD5 and were able to demonstrate an explicit collision!
- Since then, the attack has been improved—collisions can be found in under a minute on a desktop PC—and extended so that even “controlled collisions” (e.g., two postscript files generating arbitrary viewable content) can be found.
- Due to these attacks, MD5 should no longer be used today for any application requiring cryptographic security.

SHA-0, SHA-1, SHA-2

- The Secure Hash Algorithm (SHA) refers to a series of cryptographic hash functions standardized by NIST.
- SHA-1, was introduced in 1995. This algorithm has a 160-bit output length, and supplanted a predecessor called SHA-0 which was withdrawn due to unspecified flaws discovered in that algorithm.
- Theoretical analysis over the past few years indicates that collisions in SHA-1 can be found using significantly fewer than the 280 hash function evaluations that would be necessary using a birthday attack.
- Recently an explicit collision has been found.
- It is therefore recommended to migrate to SHA-2, which does not currently appear to have the same weaknesses.
- SHA-2 comprises two related functions: SHA-256 and SHA-512, with 256- and 512-bit output lengths, respectively.

SHA-0, SHA-1, SHA-2

- All hash functions in the SHA family are constructed using the same basic design:
 - A compression function is first defined using the Davies-Meyer construction as applied to some block cipher
 - Extended to support arbitrary length inputs using the Merkle-Damgård transform.
- The block cipher in each case was designed specifically for building the compression function.
 - Block ciphers SHACAL-1 (for SHA-1) and SHACAL-2 (for SHA-2). Have large block lengths (160 and 256 bits respectively) and 512-bit key lengths.

SHA-3 (Keccak)

- NIST announced in late 2007 a public competition to design a new cryptographic hash function to be called SHA-3.
- Submitted algorithms were required to support both 256- and 512-bit output lengths.
- 51 first-round candidates were narrowed down to 14 in December, 2008, and these were further reduced to five finalists in 2010. The remaining candidates were subject to intense scrutiny by the cryptographic community over the next two years.
- In October, 2012, NIST announced the selection of Keccak as the winner of the competition.
- This algorithm is currently undergoing standardization as the next-generation replacement for SHA-2.

SHA-3 (Keccak)

- Keccak is unusual in several respects.
 - One of the reasons Keccak was chosen is because its structure is very different from that of SHA-1 and SHA-2.
- It is based on an **unkeyed** permutation f with a large block length of 1600 bits; this is radically different from, e.g., the Davies-Meyer construction which relies on a keyed permutation.
- Keccak does not use the Merkle-Damgard transform to handle arbitrary input lengths. Instead, it uses a newer approach called the **sponge** construction.
- Keccak—and the sponge construction more generally—can be analyzed in the random-permutation model
 - Here parties have access to an oracle for a random permutation $f : \{0, 1\}^\ell \rightarrow \{0, 1\}^\ell$ (and possibly its inverse).
 - This is weaker than the ideal-cipher model.