


→ HMAC (most widely used MAC)

So how do we use hash functions to obtain a secure MAC? Will revisit after studying constructions of CRHFs.

Many cryptographic hash functions (e.g., MD5, SHA-1, SHA-256) follow the Merkle-Damgård paradigm: start from hash function on short messages and use it to build a collision-resistant hash function on a long message:

1. Split message into blocks
2. Iteratively apply compression function (hash function on short inputs) to message blocks



h : compression function

$t_0, \dots, t_\ell$: chaining variables

padding introduced so last block is multiple of block size

→ must also include an encoding of the message

length: typically of the form $100 \dots 0 \parallel \langle s \rangle$

where $\langle s \rangle$ is a fixed-length binary representation of message length in blocks

Recall: $100 \dots 0$ padding was used in the ANSI standard

if not enough space to include the length, then extra block is added (similar to CBC encryption)

Hash functions are deterministic, so IV is a fixed string (defined in the specification) — can be taken to be all-zeros string, but usually set to a custom value in constructions

for SHA-256:
 $X = \{0, 1\}^{256} = \mathbb{Y}$

Theorem. Suppose $h: X \times Y \rightarrow X$ be a compression function. Let $H: Y^{\leq \ell} \rightarrow X$ be the Merkle-Damgård hash function constructed from h . Then, if h is collision-resistant, H is also collision-resistant.

Proof. Suppose we have a collision-finding algorithm A for H . We use A to build a collision-finding algorithm for h :

1. Run A to obtain a collision M and M' ($H(M) = H(M')$ and $M \neq M'$).
2. Let $M = m_1, m_2, \dots, m_u$ and $M' = m'_1, m'_2, \dots, m'_v$ be the blocks of M and M' , respectively. Let $t_0, t_1, \dots, t_u$ and $t'_0, t'_1, \dots, t'_v$ be the corresponding chaining variables.
3. Since $H(M) = H(M')$, it must be the case that

$$H(M) = h(t_{u-1}, m_u) = h(t'_{v-1}, m'_v) = H(M')$$

If either $t_{u-1} \neq t'_{v-1}$ or $m_u \neq m'_v$, then we have a collision for h .

Otherwise, $m_u = m'_v$ and $t_{u-1} = t'_{v-1}$. Since m_u and m'_v include an encoding of the length of M and M' , it must be the case that $u = v$. Now, consider the second-to-last block in the construction (with output $t_{u-1} = t'_{u-1}$):

$$t_{u-1} = h(t_{u-2}, m_{u-1}) = h(t'_{u-2}, m'_{u-1}) = t'_{u-1}$$


Either we have a collision or $t_{u-2} = t'_{u-2}$ and $m_{u-1} = m'_{u-1}$. Repeat down the chain until we have collision or we have concluded that $m_i = m'_i$ for all i , and so $M = M'$, which is a contradiction.

Note: Above construction is sequential. Easy to adapt construction (using a tree) to obtain a parallelizable construction.

Sufficient now to construct a compression function.

Typical approach is to use a block cipher.

Davies-Meyer: Let $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{X}$ be a block cipher. The Davies-Meyer compression function $h: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{X}$ is then



$$h(k, x) := F(k, x) \oplus x$$

Many other variants also possible: $h(k, x) = F(k, x) \oplus k \oplus x$
[used in Whirlpool hash family]

Need to be careful with design!

- $h(k, x) = F(k, x)$ is not collision-resistant: $h(k, x) = h(k', F^{-1}(k', F(k, x)))$
- $h(k, x) = F(k, x) \oplus k$ is not collision-resistant: $h(k, x) = h(k', F^{-1}(k', F(k, x) \oplus k \oplus k'))$

Theorem. If we model F as an ideal block cipher (ie, a truly random permutation for every choice of key), then Davies-Meyer is collision-resistant.

Conclusion: Block cipher + Davies-Meyer + Merkle-Damgård $\Rightarrow$ CRHFs

Examples: SHA-1: SHACAL-1 block cipher with Davies-Meyer + Merkle-Damgård

SHA-256: SHACAL-2 block cipher with Davies-Meyer + Merkle-Damgård

Why not use AES?

- Block size too small! AES outputs are 128-bits, not 256 bits (so birthday attack finds collision in 2^{64} time)
- Short keys means small number of message bits processed per iteration.
- Typically, block cipher designed to be fast when using same key to encrypt many messages
 - ↳ In Merkle-Damgård, different keys are used, so alternate design preferred (AES key schedule is expensive)

→ birthday attack run-time: $\sim 2^{80}$
attack ran in time $\sim 2^{64}$ (100,000x faster)
January, 2020: chosen-prefix collision in $\sim 2^{63.4}$ time!

← no longer secure [first collision found in 2017!]

- SHA-1 extensively used (eg, git, svn, software updates, PGP/GPG signatures, certificates) → attacks show need to transition to SHA-2 or SHA-3

Recently: SHA-3 family of hash functions standardized (2015)

- ↳ Relies on different underlying structure ("sponge" function)
- ↳ Both SHA-2 and SHA-3 are believed to be secure (most systems use SHA-2 - typically much faster)

↙ or even better, a large-domain PRF

Back to building a secure MAC from a CRHF - can we do it more directly than using CRHF + small-domain MAC?

- ↳ Main difficulty seems to be that CRHFs are keyless but MACs are keyed

Idea: include the key as part of the hashed input

By itself, collision-resistance does not provide any "randomness" guarantees on the output

- ↳ For instance, if H is collision-resistant, then $H'(m) = m_{10} \parallel \dots \parallel m_{10} \parallel H(m)$ is also collision-resistant even though H' also leaks the first 10 bits/blocks of m

- ↳ Constructing a PRF/MAC from a hash function will require more than just collision resistance

- Option 1: Model hash function as an "ideal hash function" that behaves like a fixed truly random function (modeling heuristic called the random oracle model - will encounter later in this course)

- Option 2: Start with a concrete construction of a CRHF (eg, Merkle-Damgård or the sponge construction) and reason about its properties

- ↳ We will take this approach

Suppose H is a Merkle-Damgård hash function built from a secure compression function

Several ways to build a keyed function:

1. Prepend key: $F(k, m) := H(k \| m)$

↳ Insecure due to structure of Merkle-Damgård: can mount an "extension attack": given $H(k \| m)$, can compute $H(k \| m \| m')$ by extending Merkle-Damgård chain

2. Append key: $F(k, m) := H(m \| k)$

↳ Similar to hash-then-MAC construction and vulnerable to same offline attack: adversary finds a collision in the Merkle-Damgård prefix and uses that to construct a forgery

↳ Structure exploited in SHA-1 collision demonstration (can generate arbitrary collisions once prefix matches)

3. Envelope method: $F(k, m) := H(k \| m \| k)$

4. Two-key nest: $F(k_1, k_2, m) := H(k_2 \| H(k_1, m))$

} for reasonable pseudorandomness assumptions on h (e.g., both $F_1(k, m) := h(k, m)$ and $F_2(k, m) := h(m, k)$ is a PRF), both of these constructions are secure PRFs on a variable-size domain

hash-based MAC

HMAC is a PRF/MAC based on the two-key nest (though with correlated keys):

$$\text{HMAC}(k, m) := H(k_1 \| H(k_2, m))$$

where $k_1 \leftarrow k \oplus \text{ipad}$ and $k_2 \leftarrow k \oplus \text{opad}$

and ipad and opad are fixed strings (specified in the HMAC standard)

↑
0x36 repeated

↑
0x5C repeated

Security: Since k_1 and k_2 are correlated, need to make stronger assumption on security (e.g., h remains pseudorandom under a related-key attack)

Instantiations: Typically, denoted $\text{HMAC-}H$ where H is the hash function

e.g., HMAC-SHA1

HMAC-SHA256 — one of the most widely-used MAC on the web (used in SSL/TLS, IPsec, SSH, and more)

HMAC for key-derivation: Recall that under reasonable assumptions, HMAC is a secure PRF

In many protocols, we need to derive multiple keys from a single master key (e.g., derived from a password)

↳ To derive multiple independent cryptographic keys, a PRF is a natural primitive:

$$k_{\text{enc}} \leftarrow \text{HMAC}(k_{\text{master}}, \text{"enc"})$$

$$k_{\text{mac}} \leftarrow \text{HMAC}(k_{\text{master}}, \text{"mac"})$$

↑
derived keys

↑
master key

↑
tag (just has to be unique)

} PRF security says derived keys are computationally indistinguishable from uniform

This approach is used in TLS and IPsec to derive session keys during session setup

↳ General paradigm is the "expand" step in hash-based key-derivation (HKDF — RFC 5869)

↳ Consists of two procedures:

- Extract: derive a master key from entropy source (e.g., a user password)

- Expand: derive sub-keys from the master key

Both steps rely on HMAC