

Lecture 3 - Collision Resistance

G.1600 - fall 2022

C-G, Kalai, Zeldovich

MIT

_____ 

Collision - Resistant Hash Functions

Plan


- * Intuition & motivation
- * Defn of CRHF
- * Constructions
- * Attacks
- * If time: applications.

Logistics

- * Lab 0 code & Lab 0 theory due tomorrow 10pm ET via Gradescope
↳ Latex for written parts
- * Lab 1 out on Friday.

Last time...

authenticating PEOPLE



Passwords, Pass storage, MACs,
biometrics, ...

↓
yes!

Today...

authenticating FILES / CODE / DATA


Main new tool:

Collision-resistant hash functions. (CRHF)

Intuition behind CRHF....

Compressing!

A hash fn $H: \{0,1\}^* \rightarrow \{0,1\}^{256}$
↳ In practice SHA2/SHA256, SHA3/Keccak, ... (Broken: MD5, SHA1)




Security goal: It is "computationally infeasible" to find two distinct files that have the same hash value (a "collision")

(There are other security goals as well.... pset 1)

Application I: Secure mirroring

1. Get hash from trustworthy source



2. Fetch large file from untrustworthy source




If Hash is CRHF, then sketchy mirror will not be able to find a file $\hat{f} \neq f$ that client will accept.

Used in subresource integrity (SRI)!


```
<script src="sketchy.mit.edu/code.js" integrity="sha256-ogAB....." >
```

Application II: Outsourced File Storage

1.



2.



IF hash is CDF, then Google can't track you into accepting incorrect photos/files.

More generally, CRHFs let you authenticate a LONG message by authenticating only a SHORT string.

We will see more applications...
"Hash and sign", ...

Adversary's goal in breaking CRHF.




m_0, m_1

Distinct msgs s.t.

$$H(m_0) = H(m_1)$$

Observe: There are lots of collisions!

All bitstrings



256-bit strings

Mega Pigeonhole Principle!

Jamming infinitely many pigeons in finite holes

- IS CRHF is good/secure, these collisions will be hard to find.

↳ How do we formalize this?

Definition: Collision-Resistant Hash Function

A function $H: \{0,1\}^* \rightarrow \{0,1\}^\lambda$ is collision resistant if for all "efficient" adversaries A

Adv is randomized

$$\Pr [H(m_0) = H(m_1) : (m_0, m_1) \leftarrow A(\cdot)] \leq \text{"negligible"}$$

(To be useful, H must also be efficiently computable.)

In theory: λ = "security parameter" (\approx key length)

"efficient" = randomized alg runs in time $\text{poly}(\lambda)$

"negl" = $O(\frac{1}{\lambda^c}) \quad \forall c \in \mathbb{N}$

(e.g. $\frac{1}{2^\lambda}, \frac{1}{2^{\sqrt{\lambda}}}, \frac{1}{\lambda^{\log \lambda}}, \dots$)

In practice:

$\lambda = 128, 256, 384$

"efficient" adversary \approx runs in time $\leq 2^{128}$

"negl" \approx prob $\leq 2^{-128}$

In practice, aim to defend against advs running in time $< 2^{28}$.

Time

2^{30} ops/sec on your laptop

2^{58} ops/sec on Fugaku supercomputer ($\approx \$1$ billion)

2^{61} hashes/sec computed by Bitcoin miners

2^{92} hashes/year " " "

2^{114} hashes requires enough energy to bail all water on the planet

2^{140} hashes requires one year of Sun's energy

Lenstra
Kleinjung
Thome

Probability

2^{-1} fair coin lands heads

2^{-8} tax returns audited by IRS

2^{-13} probability that randomly sampled MIT grad is Nobel prize winner

2^{-11} struck by lightning in next year

2^{-28} probability of winning Mega Millions jackpot

2^{-69} probability of all happening (assuming independence)
 2^{-128} ... a billion billion times less likely than that.

How to construct CRHFs.

Two steps:

1 Construct CRHF $H_{\text{small}}: \{0,1\}^{2^n} \rightarrow \{0,1\}^n$

↳ More art than science.

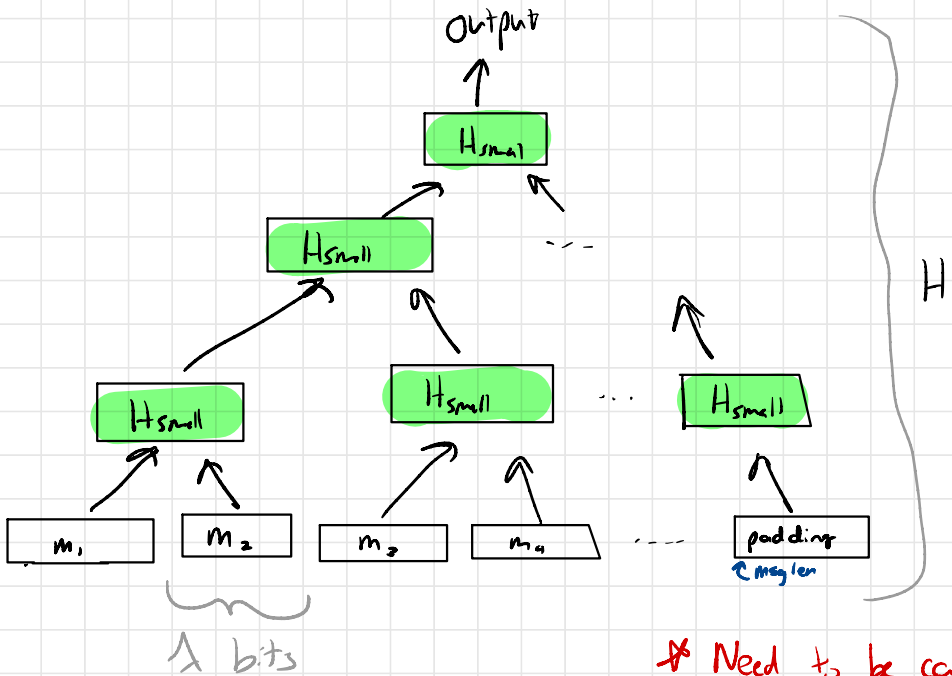
Come up with candidate, try to break it using known techniques, assume it's CRHF

↳ Current standard is SHA256, designed by NSA, published 2001

↳ Can also build from number theory (Factoring, etc)
...but too slow


[Aside: If $P=NP$, CRHFs don't exist.
So security of CRHFs relies on unproven assumptions... $P \neq NP$ & more]

2. Use H_{small} to construct $H: \{0,1\}^* \rightarrow \{0,1\}^\lambda$
 "Merkle-Damgard"
 $\hookrightarrow H$ is CRHF if H_{small} is
 (No need for extra assumption)





* Need to be careful about padding... don't implement yourself!

Another variant (not parallelizable)



Why Merkle-Damgard works

Consider hash fn $H: \{0,1\}^{3\lambda} \rightarrow \{0,1\}^\lambda$



Claim: If \exists eff adv A that finds collisions in H_{big}
 \exists eff adv B that finds collisions in H_{small}

Run $A() \rightarrow [(m_0, m_1, m_2), (m'_0, m'_1, m'_2)]$

$$(x, m_2) \stackrel{?}{=} (x, m'_2)$$

NO

$(x, m_2), (x, m'_2)$
is a collision
for H_{small} !

YES

$$(m_0, m_1) \stackrel{?}{=} (m'_0, m'_1)$$

NO

$(m_0, m_1), (m'_0, m'_1)$
is a collision
for H_{small} !

YES

Contradiction!
 $(m_0, m_1, m_2) = (m'_0, m'_1, m'_2)$

Given hash f_n with n -bit output, can find collision in time $O(\sqrt{2^n}) = O(2^{n/2})$

versus 2^n for brute-force search

Random pigeonhole principle (a.k.a. birthday paradox)

If you throw \sqrt{N} pigeons into N pigeonholes independently & uniformly at random, then with prob $\approx 1/4$, \exists two pigeons share same hole.

\Rightarrow To find collision, hash $2^{n/2}$ random strings. By an argument involving R.P.P./Birthday, \exists colliding pair often

\Rightarrow If you want adv to do 2^{128} work to find collision, need to have 256-bit output.

\rightarrow In practice, we use SHA256 (or SHA3)
(on my laptop, get ≈ 1 GB/s)
openssl speed sha256

Historical Note:

- * For many years, MD5 (designed by Ron Rivest) was the standard CRHF - 128-bit output
 - * 2004 Weng et al find collision - time is now $\approx 2^{24}$
 - * We used to use SHA1 (160-bit output)
 - * In 2017 researchers at CWI AMs & Google found a collision in SHA1 using 2^{63} hashes
 - * Attack cost \approx \$100k - \$500k
- \Rightarrow SHA1 deprecated
- $\approx 100,000\times$ faster than brute force.

Domain Separation

Given **one-input** CRHF $H: \{0,1\}^* \rightarrow \{0,1\}^{256}$,

often want to build two-input CRHF:

$$H_2(x, y).$$

BAD IDEA:

$$H_2(x, y) := H(x \parallel y)$$

↖ string concat

Notice: $H_2(\text{"key"}, \text{"value"})$
 $= H_2(\text{"ke"}, \text{"yvalue"})$

→ Even though H is CRHF, H_2 is not!

action=create & user = nickol


Some hash value! $\left[\begin{array}{l} \text{action=create \& user = action delete user nickola:} \\ \text{a=ctioncreateuser \& action=delete \& user = nickola:} \end{array} \right.$

Flickr and Amazon EC2 APIs were vulnerable to this attack!

Better idea: Unambiguous encoding (e.g. length, val, length, val, ...)

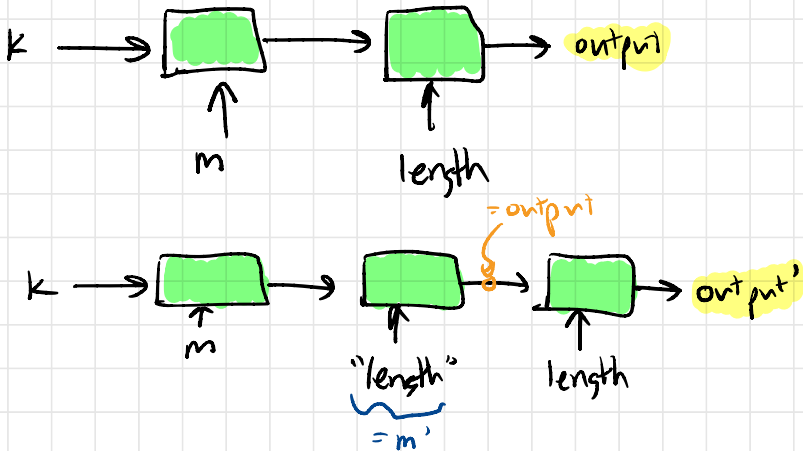
Length-extension attack

Recall message authentication codes (MACs) from last lecture.



BAD IDEA: $\text{MAC}(k, m) := H(k \| m)$

Given $\text{MAC}(k, m)$ can compute $\text{MAC}(k, m \| m')$ without knowing key k .




$$H(k \| m) = H(k, m \| m')$$

Flicker & Amazon APIs also vulnerable to this attack?

Better idea: Use MAC for a MAC — not CRHF.

Application: Merkle trees (Authenticating many files with a single digest)

A variant on our secure mirroring application...




$\hat{f}_i \stackrel{?}{=} f_i$

Option: Source sends N hashes
↳ a lot of communication over wide-area net


Option: Client downloads all N files
↳ even more communication!

Better idea: Use the Merkle construction



⇒ Mirror sends one full file + $O(\log N)$ hashes
 ≪ than N hashes!
 ≪ than N files!

⇒ CP property ensures that mirror can't cheat



Used in certificate transparency, ...

Thc

Ekd

COL. FINDER

(See Bellare textbook appendix)

Given: $H: \{0,1\}^n \rightarrow \{0,1\}^n$ [Model H as a random function]

Find: $m_0, m_1 \in \{0,1\}^{2n}$ st. $m_0 \neq m_1$
 $H(m_0) = H(m_1)$.

Let $T = 2^{n/2}$

Choose distinct $r_1, r_2, r_3, \dots, r_T \xleftarrow{R} \{0,1\}^{2n}$

Compute $H(r_1), H(r_2), \dots, H(r_T)$.

↳ Likely to find a collision!

B_i = event that \nexists collision after computing i th hash

$$\Pr[B_i | B_{i-1}] = 1 - \frac{i}{2^n}$$

$$\Pr[\text{no collision}] = \Pr[B_T]$$

$$= \Pr[B_T | B_{T-1}] \cdot \Pr[B_{T-1}]$$

$$\dots$$
$$= \prod_{i=1}^T \Pr[B_i | B_{i-1}]$$

$$= \prod_{i=1}^T \left(1 - \frac{i}{2^n}\right)$$

$$\leq \prod_{i=1}^T e^{-i/2^n}$$

$$\leq \exp\left(-\sum_{i=1}^T \frac{i}{2^n}\right) \leq \exp\left(-\Omega\left(\frac{T^2}{2^n}\right)\right)$$

$$\Pr[\text{collision}] \geq 1 - \text{constant.}$$

↪ repeat a few times

Useful life fact.
 $1+x \leq e^x$


Application: Commitments

- * "Sealed envelope" with cryptography.
- * Just a small tweak to the earlier applications
- * Requires a bit more than plain CRHF, but any CRHF can be made suitable ...

[Halverson, Micali '96]

"Coin Flipping"

1. Hiding: $H(m, r)$ "hides" Alice's bit b_A *
2. Binding: Alice can't change bit after sending to Bob.



Modulo Alice refusing to open, neither party can control bit b .

↳ Distributed randomness used for protocols that require good randomness w/o trustworthy dealer (e.g. lots, ...)