Lecture 8

Thursday, 29 September 2022 15:24

Hash functions

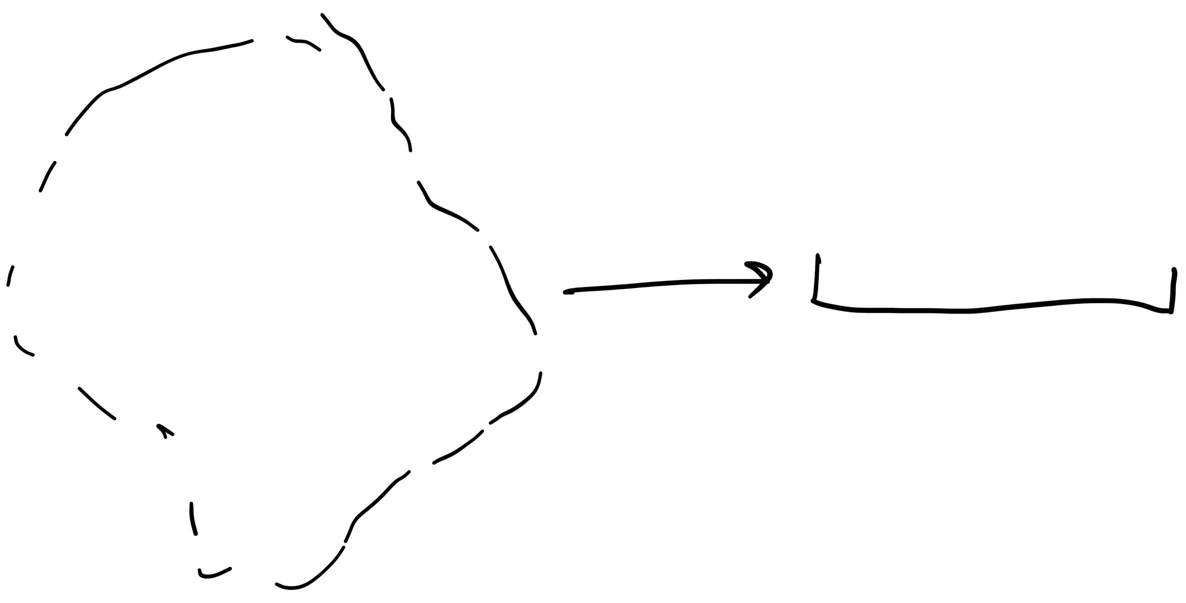

hash functions maps

$h : \{0,1\}^* \longrightarrow \{0,1\}^n$ — arbitrary length; strings of length n

$h : M \subseteq \{0,1\}^* \longrightarrow H = \{0,1\}^n$

message space; hash space

note: not all messages are valid (as e.g. PDF files)

use $h(m)$ as fingerprint of m.

Any change in m changes $h(m)$

we Need more properties from a cryptographic hash function because the adversary controls things

preimage resistance: given $y \in H$ (hash space $= \{0,1\}^n$) it is hard to find $x \in M$ with $y = h(x)$

This is no harder than $O(2^n) = O(|H|)$ (size of hash space)

by just a brute-force attack: pick $x \in M$, compute $h(x)$ and compare to y. Repeat if $h(x) \neq y$.

Second-preimage resistance: given $x \in M$, it is hard to find $x' \in M$ with $x \neq x'$ and $h(x) = h(x')$.

Same attack finds x' in $O(2^n)$.

Collision resistance: it is hard to find $x \neq x' \in M$ with $h(x) = h(x')$.

Same attack (random picking & comparing) takes $O(2^{n/2})$ by birthday paradox.

Want to use Pollard rho to save memory. Need $f(w_i) \to w_{i+1}$

Using h that can be iterated

we must be able to get back to the message space, i.e.

$M \xrightarrow{h} H \xrightarrow{\phi} M$

ϕ: could be identity if $H \subseteq M$; needs to be deterministic

E.g. PDF has some flexible part at end (which does not impact what is displayed); ϕ varies that

Once I have a collision at DP (distinguished point), how do I get $x, x' \in M$ for the first collision?

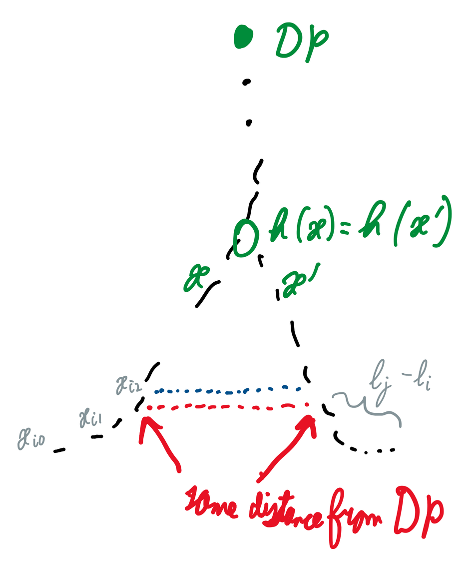

When DP is reached, report DP, starting point x_{j0} and number of steps, say l_j, to DP.

If x_{i0} and x_{j0} lead to same DP, with lengths l_i and l_j with $l_i \leq l_j$ then compare $h(x_k)$ to $h(x_{k+l_j-l_i})$ for $x \in M$, update includes ϕ.