




Lecture 1: Modern Hashing.

Syllabus recap:


- Preps: CSE 311/312
(can msg me about add codes but they're pretty strict).
- One topic (related) per week.
- One HW per topic:
 - 100% of grade!
 - collaboration policy
 - no late days but lowest dropped
 - GPT allowed but maybe not so useful.
 - Submit written on Gradescope, code on canvas.
- OH / TAs

Consistent Hashing (Karger, Lehman, Leighton, Levine, Lewin, Panigrahy).

Q: How to design cache for internet?



Problem: This cache usually will not fit on 1 machine. Must distribute across m machines



Idea 1: Brute force search. \rightarrow too slow!

Idea 2: Hashing

Recall: $h: \Sigma^* \rightarrow [m]$ (e.g. $h(x) = \text{SHA}(x) \bmod m$).

x_1, \dots, x_n a set of URLs.

assign x_i to $h(x_i)$.

If h is "pseudorandom" then #i assigned to any bucket $\approx \frac{1}{m}$.

To lookup x in cache: check machine $h(x)$.

done...?

Problem: what if a machine fails / is added?

$\text{SHA}(x) \bmod 100 \neq \text{SHA}(x) \bmod 99 \neq \text{SHA}(x) \bmod 101$ (usually).

so we'd need to redistribute (almost) all the keys.
(all but a $\sim \frac{1}{m}$ fraction) BAD!

Consistent Hashing: Take $h: \Sigma^* \rightarrow \mathbb{Z}^{32}-1$ (some large number),

and treat this as indexing a position on a circle.

hash each server


s_1, \dots, s_m

For data x ,

store it on


first server

clockwise of $h(x)$



If h is "uniform-like", then # of x in $\boxed{S_1 \dots S_F} \approx \frac{n}{m}$.

If server is added, what needs to be reindexed?




If server is removed, similar.

How to perform lookup?


given x with hash $h(x)$, find s with smallest $h(s)$ satisfying $h(s) \geq h(x)$ (ignoring the wrap around).

Naively: need to compute m hash functions
(slow).

Ideas: use a balanced search tree (BST).



Problem: Variance. In expectation, each server will see $\approx \frac{1}{m}$ traffic. But this is not with high probability






Solution : Variance reduction via virtual copies.

Idea: use k hash functions

$$h_1, \dots, h_k: \Sigma^* \rightarrow \mathbb{Z}^{32} - 1.$$

Associate each server with k intervals, given by the positions of these k hash functions.



The lengths of the intervals "average out". Can show that if $k = \log_2(n)$, then all servers have load $O(\frac{n}{m})$ whp.

Also useful for heterogeneous servers: If one server is bigger, add more virtual copies!