

6.854 / 18.415: Advanced Algorithms

Last time: Load balancing and Power of two choices

what is max load of n balls in n bins? $O\left(\frac{\log n}{\log \log n}\right)$

what if each chooses least loaded of two bins? $O(\log \log n)$

Today: Consistent hashing, random trees and distinct elements

remind notation to be consistent with first lecture

Setup: m items to store on n distributed webcaches

Recall our favorite 2-universal hash family

$$h_{a,b}(x) = (ax + b \pmod{p}) \pmod{n}$$

But what if n changes?

(1) $\pmod{n} \rightarrow \pmod{n+1}$

Almost all items need to move so we are searching in right place

(2) keep n

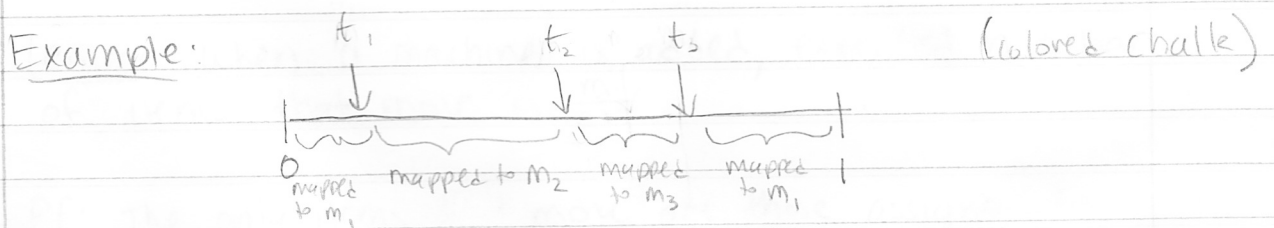
New cache gets none of load

Consistent Hashing: [Karger, Lehman, Leighton, Levine, Lewin, Panigrahy]

(1) map each machine randomly to interval $[0, 1]$

(2) map each item

(3) assign each item to first machine on right (with wrap around)



Implementation:

(1) maintain binary search tree whose keys are machine values

To insert item, find predecessor ^{successor} add item to machine


To insert machine, find predecessor ^{successor} and steal items

To delete machine, find predecessor ^{successor} and push items

only jobs that need to move, mac

lemma: w.h.p. no machine owns more than $O\left(\frac{\log n}{n}\right)$ of interval

Pf: Consider intervals of length $\frac{2 \log n}{n}$



Probability no machine in a fixed interval

$$\left(1 - \frac{2 \ln n}{n}\right)^n = \frac{1}{n^2}$$

gives $\frac{4 \log n}{n}$

why?

By union bound, probability each interval has a machine is at least $1 - \frac{1}{n}$ \square

Exercise: Use Chernoff to prove max load (if $m \geq n$) is at most $O\left(\frac{m}{n} \log n\right)$ (don't lose extra log from balls in bins)

Lemma: When a machine is added, expected number of items that move is $\frac{m}{n+1}$

Pf: The only items that move are those assigned to machine $n+1$. By symmetry, expected number of ^{such} items is $\frac{m}{n+1}$ \square

How small will the smallest load be? By birthday paradox, expect some machine to own $\frac{1}{n^2}$ of interval

spread out load

Refinement: Each machine is mapped to $\log n$ points in $[0, 1]$ interval, owns union of left segments

Relieving Hot Spots on Web

examples from paper: IBM website during Deep Blue - Kasparov


examples from real life: Victoria Secret, Star Wars EP1

setup: ^{root}single server, n proxy caches C

If all requests go to root it can't handle them, and web could crash due to traffic. Can we share load?

Random Trees

(1) choose random d -ary tree with n virtual nodes V



hash: $h_i: V \rightarrow C_i$

why hash? many pages = many roots

(2) each request. (if view C_i)

(a) choose random leaf, consider leaf-to-root path p

(b) apply h_i to leaf, ask machine

(c) If storing page, return it. otherwise increment page's counter, ask next machine on path. If counter reaches q , store page.

lemma: Each cache not mapped to a leaf is asked for same page at most

$O(d \log n)$ times

Pf. Each cache is responsible for at most $O\left(\frac{\log n}{\log t \log n}\right)$ virtual nodes. (balls in bins)
(non leaf)

Each virtual node can be asked at most dq times before its descendants all store the page \square

Inconsistent world. What if everyone has a different view of set of live caches?

$$C_1, C_2, \dots, C_\ell \subseteq C \text{ with } \frac{|C_j|}{|C|} \geq \frac{1}{\ell} \quad \forall j$$

Now where a virtual node is mapped to depends on the particular view


$$h_j: V \rightarrow C_j$$

Lemma. If there are at most n total requests, each cache is asked for a page at most

$$O(dq t^2 \log^2 n \ell)$$
 times

Part 2:
how many virtual nodes can a cache be responsible for overall?

Pf. Fix cache c_i , and let it be mapped to t_i . Then consider the interval



of length $\frac{10t \log n \ell}{n}$. Fix a view C_j .

The probability no cache in C_j is mapped to I

is at most $(1 - 10 \frac{t \log n e}{n})^{|C_j|} \leq (1 - 10 \frac{t \log n e}{n})^{\frac{n}{e}} \leq \frac{1}{(en)^{10}}$

Now union bound over all l views, and with probability at least $1 - \frac{1}{n^{10}}$ we have that each view C_j maps a cache to \underline{I}

on circle

This cache v is not responsible for any virtual node mapped outside of \underline{I} . Now we can use Chernoff to conclude v is responsible for at most

$O(t \log n e)$ virtual nodes, whp

hash virtual nodes to caches

Then union bound over v , and each virtual node gets at most $d e$ requests as before. \square


Next time: Distinct elements and Count min

setup: router wants to estimate number of distinct IP addresses

We could hash m items into m^2 bins, and count. In fact we can do it with $\log m$ bits of storage.

part 2:
How many caches can be mapped to a single virtual node over all views?

Now fix a virtual node j , and let it be mapped to v_j



and with high probability every view maps

a cache to I' . But by Chernoff, at most $O(t \log n e)$ caches are mapped to I' in all of C .

In any view, virtual node j can only be mapped to caches ^{that land} in I . Thus each virtual node can be responsible for at most $O(c t \log n e)$ requests. ~~\square~~

↓
Now proceed as before