(D) Load Balancing (again) & Power of 2 choices

Know: n balls each ball into $\underbrace{\text{uniformly at random}}_{\text{u.a.r.}}$ bin independently

n bins

$$\Rightarrow \text{Max load} \approx \frac{\log n}{\log\log n}.$$

Of course: can do "round robin"

Send ball i to bin i $\Rightarrow$ max load $= 1$.

But needs state! what if dynamic load? (arrivals, departures)

distributed?

Lovely Idea: [Azar Broder Karlin Upfal '94]

For each ball, probe 2 bins u.a.r. & independently.

Send to min loaded bin among these two!

Power of 2 choices

Theorem

The max load for n balls and n bins is now.

$$\ln\ln n + O(1) \quad \text{whp.}$$

no constant factors, mind you!

for Power-of-2-choices.

For d choices, answer is

$$\frac{\ln\ln n}{\ln d} + \Theta(1) \quad \text{whp.}$$

And these results are essentially optimal !

What do we mean ?

Well, it depends on what the algo is allowed.

If this particular algo, these are the tight bounds.

For any "local" "symmetric" algo, also optimal

Can do a bit better if not "symmetric" (better constants)

— more later

— X —

OK : what's the intuition ?

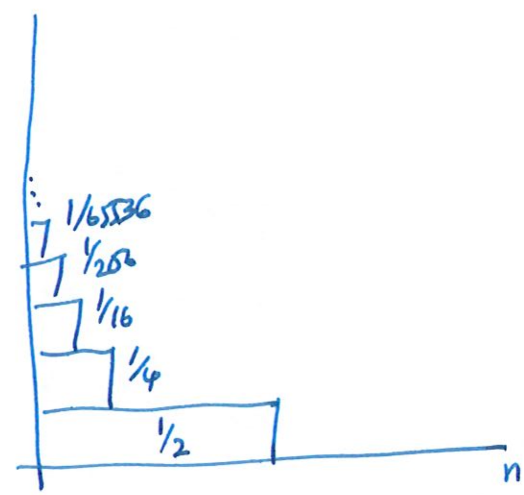

Def : bin has height h
~~if~~ final load $\geq h$.

Def 2: ball has height h if it's bin had
$h-1$ balls just before it arrived.

Fact 1 : if ball has height $h \Leftrightarrow$ both probes to bins of height $h-1$.

Fact 2 : if bin has height h, contains ball of height h.

Fact 3 : (fraction of bins of height 2) $\leq$ 1/2

Let f_h = fraction of bins of height h. $(\Rightarrow f_2 \leq 1/2)$.

Fact 4 : $\Pr$(ball has height h) = $\Pr$(both probes to $h-1$ height) (by Fact 1)
$= (f_{h-1})^2$.

Fact 5: $E[\text{fraction \# balls of height } h] = (f_{h-1})^2$.

$\Rightarrow E[\text{fraction of bins of height } h] \leq \quad = (f_{h-1})^2$
$= f_h$ (by def)

$f_2 = 1/2$

$f_3 \leq (1/2)^2 = 1/4$

$f_4 \leq 1/16$

$\vdots$

$\Rightarrow$ (induction)

$f_i \leq 2^{-2^{i-2}}$

$\Rightarrow f_{i+1} \leq 2^{-2^{i-2}} \cdot 2^{-2^{i-2}}$

$= 2^{-2^{i-1}}$ as claimed

$\Rightarrow f_{\log\log n} \leq 1/n \Rightarrow$ no bins with height $\log\log n + O(1)$.

☺

—— x ——

Aha! But all in expectation. ☹

How to do it formally?

Chernoff bounds. (Induct on heights)

- Show that #balls / bins of the correct heights is close to what you expect.

- Requires careful calculations, skipping, see paper / notes.

Let's see a simpler proof using random graphs / basic calculation.

☺

Power of 2 choices via basic Calculations

[Karp
Luby
Meyer-auf der Heide
'92]

Weaker bound:

throw $m = \frac{n}{\text{constant}}$ balls into n bins

lose constant

then the maximum load is $O(\log\log n)$

lose constants

Proof uses 2 facts about random graphs

Build a graph on n bins, where each edge chooses its two endpoints indep. and uniformly at random (so self-loops possible).

m edges/balls

Let this graph be G.

Fact 1: whp, the largest connected component in $G = O(\log n)$ sized

Fact 2: For all subsets $S \subseteq V$, the number of edges that fall into $G[S]$ is at most $6|S|$.

graph induced by S, i.e. both endpoints of edge belong to S

— x —

OK. given these facts, how to show the weaker theorem above?

See the Latex writeup (sorry – out of time).