Power of Two Choices &
Randomized Load Balancing

Hoeteck Wee

Load Balancing

PROBLEM. assign a sequence of m jobs to n machines so that every machine gets as few jobs as possible.

- job $\rightarrow$ ball; machine $\rightarrow$ bin

TRIVIAL SOLUTION. assign ball i to bin $i \pmod n$

- well-spread: every bin gets m/n balls
- requires coordination and maintaining global state

RANDOMIZED SOLUTION. assign each ball to a random bin

- problem? with prob $1/n^m$, all balls are assigned to bin 1

  CLAIM. with prob ≥ 0.9, every bin gets $\leq m/n \log m$ balls
- focus on $m = n$ for simplicity

THEOREM. with prob ≥ 0.9, every bin gets $O(\frac{\log n}{\log \log n})$ balls

Basic Tools

FACT. (union bound) Let $E_1, \ldots, E_n$ be events. Then

$$\Pr[E_1 \vee E_2 \vee \cdots \vee E_n] \leq \Pr[E_1] + \Pr[E_2] + \cdots + \Pr[E_n]$$

– $E_1, \ldots, E_n$ do not have to be independent

– example: 100 devices, each fails 0.01% prob

E_i: device i fails

$E_1 \vee E_2 \vee \cdots \vee E_{100}$: some device fails

with prob 99%, no device fails ("union bound over 100 devices")

FACT.: (bounds on binomial coefficients)

$$\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$$

Maximum Load

1. bound $\Pr[\text{bin } i \text{ gets} \geq k \text{ balls}]$

$$
\begin{aligned}
& \Pr[\exists \text{ subset of } k \text{ balls all of which fall into bin } i] \\
\leq \quad & \binom{n}{k} \cdot (1/n)^k \quad \text{(union bound over subsets of } k \text{ balls)} \\
\leq \quad & (ne/k)^k \cdot (1/n)^k = (e/k)^k \\
\leq \quad & 1/n^2 \qquad \text{for } k \geq \tfrac{3 \ln n}{\ln \ln n}
\end{aligned}
$$

2. union bound over all n bins:

$$\Pr[\text{some bin gets} \geq \tfrac{3 \ln n}{\ln \ln n} \text{ balls}] \leq 1/n$$

example: $n = 1$ million, maximum load is at most 16 w.h.p.

Power of Two Choices

TWO-CHOICE APPROACH. pick 2 bins at random, and place ball in least loaded bin

THEOREM. maximum load is $O(\log \log n)$ with prob ≥ 0.9

- arbitrary tie breaking
- sampling *with* replacement
- start by considering $n/8$ balls and n bins

Analysis (step one)

- STEP ONE: a graph representation
  - n vertices (one for each bin)
  - $n/8$ edges: for each ball, connect the two random bins with an edge
  - ignore ordering of the balls/edges for now
- Q: what is the maximum load for a star graph? ternary tree of depth 2?
- CLAIM 1: w.p. $1 - O(1/n)$, every connected component in the graph has size $O(\log n)$
- CLAIM 2: w.p. $1 - O(1/n)$, the average degree in every induced subgraph is at most 6.

Analysis (step one)

- CLAIM 1: w.p. $1 - O(1/n)$, every connected component in the graph has size $O(\log n)$
- PROOF

  connected component of size $\geq k$ means

  $\exists$ a subset of k vertices with $\geq k-1$ internal edges

$$
\begin{aligned}
& \sum_{|S|=k, |E|=k-1} \Pr[E \text{ is a set of } k-1 \text{ edges with both end-points in } S] \\
\leq & \sum_{|S|=k, |E|=k-1} ((k/n)^2)^{k-1} \\
= & \frac{n^2}{k^2}\binom{n}{k}\binom{n/8}{k-1}\left(\frac{k}{n}\right)^{2k} \leq \frac{n^2}{k^2}\left(\frac{ne}{k}\right)^k\left(\frac{ne}{8k}\right)^k\left(\frac{k}{n}\right)^{2k} \leq n^2\left(\frac{e^2}{8}\right)^k
\end{aligned}
$$

- ALTERNATIVE PROOF ... $\exists$ a path of length k

Analysis (step one)

- CLAIM 2: w.p. $1 - O(1/n)$, the average degree in every induced subgraph is at most 6.
- PROOF

  The complement is, $\exists$ a subset of k vertices with $> 3k$ internal edges.

$$\sum_{k=1}^{n} \sum_{|S|=k, |E|=3k} \Pr[E \text{ is a set of } 3k \text{ edges with both end-points in } S]$$

$$\leq \sum_{k=1}^{n} \binom{n}{k} \binom{n/8}{3k} \left(\frac{k}{n}\right)^{2\cdot 3k} \leq \sum_{k=1}^{n} \left(\frac{ne}{k}\right)^{k} \left(\frac{ne}{24k}\right)^{3k} \left(\frac{k}{n}\right)^{6k} = \sum_{k=1}^{n} \left(\frac{e^3}{24} \cdot \frac{k^2}{n^2}\right)^{k}$$

$$\leq \sum_{k=1}^{4\log n} \frac{16 \log^2 n}{n^2} + \sum_{k=4\log n}^{n} \frac{1}{n^2} = O(1/n)$$

Analysis (step one)

- STEP ONE: a graph representation
  - n vertices (one for each bin)
  - $n/8$ edges: for each ball, connect the two random bins with an edge
- CLAIM 1: every connected component has size $O(\log n)$
- CLAIM 2: the average degree in every induced subgraph is at most 6.
- NEXT, iterative vertex removal

Analysis (step two)

- STEP TWO: iterative vertex removal
- at each iteraction, remove all vertices (& incident edges) with deg ≤ 12
- OBSERVATION 1: remove at least $1/2$ the vertices in each iteration.
  - CLAIM 2 says average deg is ≤ 6, so less than $1/2$ have deg > 12.
- OBSERVATION 2: terminates in $\log\log n + O(1)$ iterations.
  - CLAIM 1 says initially, every connected component has size $O(\log n)$
  - apply previous observation to each connected component.
- CLAIM 3: The load of a bin (vertex) removed in iteration i is $\leq 13i$.
  - implies maximum load is $13\log\log n + O(1)$ w.h.p.
  - will refer to ordering of the balls/edges

Analysis (step two)

- CLAIM 3: The load of a bin (vertex) removed in iteration i is $\leq 13i$.
- PROOF by induction.
  - Base case: trivially follows from deg bound.
  - Analyze load of bin v (vertex) removed in iteration $i+1$,

    3 types of incident edges:

    I. ball does not land in v

    II. deleted in iterations $1, \ldots, i$ $\longleftarrow$ any number of such edges

    III. deleted in iteration $i+1$ $\longleftarrow$ at most 12 of these
  - Exploit ordering: examine loads right after the last type II edge (v, w).
    - w is deleted in iteration $\leq i$, so bin w has load at most $13i$.
    - Ball lands in bin v means bin v now has load $\leq 13i + 1$.
    - At most 12 type III edges after that

From $n/8$ balls to n balls

- CLAIM: Suppose when tossing $n/8$ balls into n bins, the max load is L w.p. at least $1-\delta$. Then, when tossing n balls into n bins, the max load is $8L$ w.p. at least $1-8\delta$.
- PROOF:
  - View tossing n balls as running 8 phases of tossing $n/8$ balls
  - By a union bound, prob. that max load exceeds L in any of the 8 phases is at most 8δ.

randomization can be useful!

终

THE END