Randomness in Computing

CS
537

LECTURE 11

Last time

- Hoeffding Bounds
- Applications of Chernoff-Hoeffding Bounds
  - Estimating a Parameter

Today

- Applications of Chernoff-Hoeffding Bounds
  - Set Balancing
  - Routing on the hypercube

Sums of Independent Bounded RVs

Hoeffding Bound

Let $X_1, \ldots, X_n$ be independent random variables with $\Pr[a \le X_i \le b] = 1$.

Let $X = X_1 + \cdots + X_n$ and $\mu = \mathbb{E}[X]$.

Then for all $\epsilon > 0$,

- (upper tail) $\Pr[X \ge \mu + \epsilon n] \le e^{-2n\epsilon^2/(b-a)^2}$
- (lower tail) $\Pr[X \le \mu - \epsilon n] \le e^{-2n\epsilon^2/(b-a)^2}$

Application: Set Balancing

- **Given:** an $n \times m$ matrix A with 0-1 entries
- **Definition:** $||(x_1, \ldots, x_n)||_\infty = \max_{i \in [n]} |x_i|$
- **Find:** $b \in \{-1,1\}^m$ minimizing $\left||Ab|\right|_\infty$

***n* features**

$$\begin{pmatrix} a_{11} & a_{12} & \ldots & a_{1m} \\ a_{21} & a_{22} & \ldots & a_{2m} \\ & & \ldots & \\ a_{n1} & a_{n2} & \ldots & a_{nm} \end{pmatrix} \quad \begin{pmatrix} b_1 \\ b_2 \\ \ldots \\ b_m \end{pmatrix}$$

***m* subjects**

Partition subjects into two groups, so that each feature is balanced.

Application: Set Balancing

Algorithm: Choose each b_i independently from $\{-1,1\}$.

Theorem. $\Pr\left[||Ab||_\infty \geq \sqrt{4m \ln n}\right] \leq 2/n$.

Proof: Let $\overline{a_i}$ be the vector in row i of A for all $i \in [n]$.

- Let B_i be the (bad) event that $|\overline{a_i} \cdot b| \geq \sqrt{4m \ln n}$.
- For $i \in [n]$, let Z_i be the random variable $\sum_{j\in[m]} \underbrace{a_{ij} b_j}$.
- $\mathbb{E}[Z_i] =$
- By Hoeffding,

$$\Pr[B_i] = \Pr[|Z_i| \geq \epsilon m] \leq 2e^{-\frac{2\epsilon^2 m}{4}} = 2e^{-2\ln n} \leq 2n^{-2}$$

- By a union bound over all n rows,

$$\Pr\left[||Ab||_\infty \geq \sqrt{4m \ln n}\right] = \Pr\left[\bigcup_{i\in[n]} B_i\right]$$

***m* independent RVs**

Each is either +1 or -1 w.p. 1/2 (if $a_{ij} = 1$) or always 0 (if $a_{ij} = 0$)

$$\epsilon m = \sqrt{4m \ln n}$$
$$\epsilon^2 = \frac{4 \ln n}{m}$$

Application: Routing on Hypercube

An n-dimensional *hypercube* is a directed graph with

- $N = 2^n$ nodes, each indexed by n-bit integer
- containing the directed edge (x, y) iff x and y differ in exactly one bit

x	001001
y	011001

How many edges?

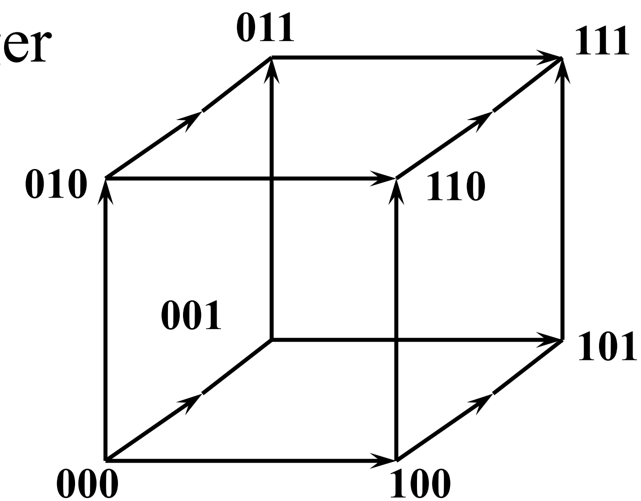

For each depicted edge, there is also the edge in the opposite direction.

- **Routing.** Each node is a routing switch.

  Edges = communication channels

  An edge can carry one packet in one step.

- A ***routing algorithm*** specifies a path from s to t for each pair of nodes (s, t) and a ***queuing policy*** for ordering packets that are waiting for the same link (e.g., FIFO – First In First Out or FTG – Furthest to Go)

CS 537 Permutation Routing Problem

- Each node is the source of one packet
- Each node is the destination of one packet

➢ E.g., on a complete graph, it can be solved in one time step.

➢ On sparse graphs?

Hypercube: N nodes, $N \log N$ edges.

Bit-Fixing Routing Algorithm for the Hypercube

1. Let x be the current node and y be the destination of a packet
2. Find smallest $i \in [n]$ such that $x_i \neq y_i$
3. Traverse the edge $(x, x_1 \ldots x_{i-1}\overline{x_i}x_{i+1} \ldots x_n)$

How long does it take to route a packet if there are no delays?

Bad example for congestion

- Transpose permutation (n is even)

  From each x, send a packet to $(x_{n/2+1}, \ldots, x_n, x_1, \ldots, x_{n/2})$

  Exercise: Show that Bit-Fixing Algorithm takes $\Omega(\sqrt{N})$ steps on transpose permutation.

- Known:

Any deterministic *oblivious* algorithm on a network with N nodes, each of outdegree d, takes $\Omega\left(\sqrt{N/d}\right)$ steps on some permutation.

Randomization is essential!

Randomized Routing

Randomized Routing Algorithm (RRA)

0. For each packet going from x to y, pick a node z uniformly and independently at random.
1. Use Bit-Fixing to route the packet from x to z.
2. Use Bit-Fixing to route the packet from z to y.

Theorem

$\log N = \boldsymbol{n}$

For every permutation, RRA takes $O(\log N)$ steps w.p. $1 - O\left(\frac{1}{N}\right)$.

Proof:

Idea 1: Analyze Phase 1. (Phase 2 is ``symmetric'' going backwards; first think about waiting with Phase 2 until all packets are done with Phase 1.)

Idea 2: In an intermediate destination, each bit z_i is 0/1 uniformly & independently

Idea 3: # steps in Phase1 is (# bits to fix) + (waiting time in queues in Phase 1)

$\leq \quad n \quad + \quad$ delay

Homework Lemma

Lemma

Let p_i be the path of some packet i in Phase 1.

Let S be the set of packets (other than i) whose routes pass through at least one edge of p_i. Then the delay of i is at most $|S|$.

Analysis of RRA

Main Lemma

Consider any packet i. It fails to reach its destination in phase 1 within $3n$ steps with probability at most $\frac{1}{N^2}$.

Proof: Let X be the number of packets (other than i) that use at least one edge from the path of i.

We would like to find an upper bound on $\mathbb{E}[X]$.

Analysis of RRA

Main Lemma

Consider any packet i. It fails to reach its destination in phase 1 within $3n$ steps with probability at most $\frac{1}{N^2}$.

Proof: Let X be the number of packets (other than i) that use at least one edge from the path of i.

Find an upper bound on $\mathbb{E}[X]$

- For any edge e, let $Y_e = \#$ routes that pass via e.
- Let $p_i = (e_1, \ldots, e_K)$ be the path of packet i. Then $K \leq n$.
- Then $X \leq Y_{e_1} + Y_{e_2} + \cdots + Y_{e_K}$.
- By symmetry of the hypercube, $\mathbb{E}[Y_e]$ is the same for all edges.
- By linearity of expectation,

  $\mathbb{E}[X|K = k] \leq \mathbb{E}[Y_{e_1}] + \cdots + \mathbb{E}[Y_{e_k}] = k \cdot \mathbb{E}[Y_e]$

Calculating $\mathbb{E}[Y_e]$

- For any edge e, let $Y_e = \#$ routes that pass via e.
- Suppose e is an edge in dimension $d \in [n]$, that is, $e = (x_1 \dots x_n, x_1 \dots x_{d-1}\overline{x_d} \dots x_n)$
- Only packets with source $* \cdots * x_d \dots x_n$ can traverse e
- To traverse e, such a packet must have destination which happens with probability
- Thus, $\mathbb{E}[Y_e] =$

Analysis of RRA

Main Lemma

Consider any packet i. It fails to reach its destination in phase 1 within $3n$ steps with probability at most $\frac{1}{N^2}$.

Proof: Let X be the number of packets (other than i) that use at least one edge from the path of i.

Find an upper bound on $\mathbb{E}[X]$

- For any edge e, let $Y_e = \#$ routes that pass via e.
- Let $p_i = (e_1, \ldots, e_K)$ be the path of packet i. Then $K \leq n$.
- By linearity of expectation,

  $\mathbb{E}[X|K = k] \leq \mathbb{E}\left[Y_{e_1}\right] + \cdots + \mathbb{E}\left[Y_{e_k}\right] = k \cdot \mathbb{E}[Y_e] =$
- By compact law of total expectation, $\mathbb{E}[X]$

$$= \mathbb{E}\big[\mathbb{E}[X|K]\big] \leq \mathbb{E}[K/2] =$$

Analysis of RRA

Main Lemma

Consider any packet i. It fails to reach its destination in phase 1 within $3n$ steps with probability at most $\frac{1}{N^2}$.

Proof: Let X be the number of packets (other than i) that use at least one edge from the path of i.

- So far: $\mathbb{E}[X] \leq \frac{n}{4}$ and travel time for packet i is $\leq n + X$
- By Hoeffding bound, $\Pr[X \geq 2n] =$

Analysis of RRA

Main Lemma

Consider any packet i. It fails to reach its destination in phase 1 within $3n$ steps with probability at most $\frac{1}{N^2}$.

Using the Main Lemma to complete the analysis:

- By a union bound over N packets, the probability that at least one packet fails to reach its destination in phase 1 within $3n$ steps is