COMPSCI 514: ALGORITHMS FOR DATA SCIENCE

Cameron Musco

University of Massachusetts Amherst. Fall 2021.

Lecture 4

- Problem set one posted this morning. Due next Friday 9/24, at 8pm.
- Class pacing: Just right – 53/87, A bit too fast – 29/87.
- If things feel way too fast, come to my office hours or send me an email and we can discuss how to make things more manageable.
- **Reminder:** My office hours are now Thursdays 9am on Zoom. Pratheba's in person office hours are Mondays 12pm-1pm in CS 207.

Last Class:

- 2-Level Hashing Analysis (linearity of expectation and Markov's inequality)
- 2-universal and pairwise independent hash functions
- Start on hashing for load balancing.

This Time:

- Finish hashing for load balancing. Motivating:
  - Stronger concentration inequalities: Chebyshev's inequality, exponential tail bounds, and their connections to the law of large numbers and central limit theorem.
  - The union bound to bound the probability that one of multiple possible correlated events happens.

Randomized Load Balancing:

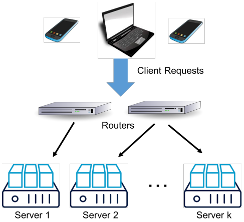

- n requests randomly assigned to k servers.
- Expected load on server i is $\mathbb{E}[\mathbf{R}_i] = \frac{n}{k}$.
- By Markov's inequality, if we provision each server to handle twice this expected load (so $\frac{2n}{k}$ requests), it will be overloaded with probability $\leq 1/2$.

CHEBYSHEV'S INEQUALITY

With a very simple twist, Markov's inequality can be made much more powerful.

For any random variable X and any value $t > 0$:

$$\Pr(|X| \geq t) = \Pr(X^2 \geq t^2).$$

X^2 is a nonnegative random variable. So can apply Markov's inequality:

Chebyshev's inequality:

$$\Pr(|X - \mathbb{E}[X]| \geq t) \quad \Pr(|X| \geq t) = \Pr(X^2 \geq t^2) \leq \frac{\mathbb{E}[X^2]}{t^2} \quad \frac{\text{Var}[X]}{t^2}.$$

(by plugging in the random variable $X - \mathbb{E}[X]$)

$$\Pr(|X - \mathbb{E}[X]| \geq t) \leq \frac{\mathrm{Var}[X]}{t^2}$$

What is the probability that X falls s standard deviations from it's mean?

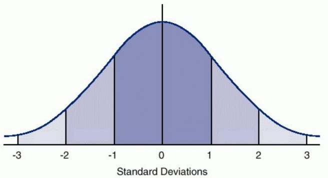

$$\Pr(|X - \mathbb{E}[X]| \geq s \cdot \sqrt{\mathrm{Var}[X]}) \leq \frac{\mathrm{Var}[X]}{s^2 \cdot \mathrm{Var}[X]} = \frac{1}{s^2}.$$

Why is this so powerful?

X: any random variable, t, s: any fixed numbers.

Consider drawing independent identically distributed (i.i.d.) random variables $X_1, \ldots, X_n$ with mean μ and variance σ^2.

How well does the sample average $S = \frac{1}{n}\sum_{i=1}^{n} X_i$ approximate the true mean μ?

$$\mathrm{Var}[S] = \mathrm{Var}\left[\frac{1}{n}\sum_{i=1}^{n} X_i\right] = \frac{1}{n^2}\sum_{i=1}^{n}\mathrm{Var}\left[X_i\right] = \frac{1}{n^2}\cdot n \cdot \sigma^2 = \frac{\sigma^2}{n}.$$

By Chebyshev's Inequality: for any fixed value $\epsilon > 0$,

$$\Pr(|S - \mathbb{E}[S]\mu| \geq \epsilon) \leq \frac{\mathrm{Var}[S]}{\epsilon^2} = \frac{\sigma^2}{n\epsilon^2}.$$

Law of Large Numbers: with enough samples n, the sample average will always concentrate to the mean.

- Cannot show from vanilla Markov's inequality.

We can write the number of requests assigned to server i, R_i as:

$$\mathsf{R}_i = \sum_{j=1}^{n} \mathsf{R}_{i,j} \quad \mathrm{Var}[\mathsf{R}_i] = \sum_{j=1}^{n} \mathrm{Var}[\mathsf{R}_{i,j}] \qquad \text{(linearity of variance)}$$

where $\mathsf{R}_{i,j}$ is 1 if request j is assigned to server i and 0 otherwise.

$$\begin{aligned}
\mathrm{Var}[\mathsf{R}_{i,j}] &= \mathbb{E}\left[\left(\mathsf{R}_{i,j} - \mathbb{E}[\mathsf{R}_{i,j}]\right)^2\right] \\
&= \Pr(\mathsf{R}_{i,j} = 1) \cdot \left(1 - \mathbb{E}[\mathsf{R}_{i,j}]\right)^2 + \Pr(\mathsf{R}_{i,j} = 0) \cdot \left(0 - \mathbb{E}[\mathsf{R}_{i,j}]\right)^2 \\
&= \frac{1}{k} \cdot \left(1 - \frac{1}{k}\right)^2 + \left(1 - \frac{1}{k}\right) \cdot \left(0 - \frac{1}{k}\right)^2 \\
&= \frac{1}{k} - \frac{1}{k^2} \leq \frac{1}{k} \implies \mathrm{Var}[\mathsf{R}_i] \leq \frac{n}{k}.
\end{aligned}$$

n: total number of requests, k: number of servers randomly assigned requests, R_i: number of requests assigned to server i.

Letting R_i be the number of requests sent to server i, $\mathbb{E}[\mathsf{R}_i] = \frac{n}{k}$ and $\text{Var}[\mathsf{R}_i] \leq \frac{n}{k}$.

Applying Chebyshev's:

$$\Pr\left(\mathsf{R}_i \geq \frac{2n}{k}\right) \leq \Pr\left(|\mathsf{R}_i - \mathbb{E}[\mathsf{R}_i]| \geq \frac{n}{k}\right) \leq \frac{n/k}{n^2/k^2} = \frac{k}{n}.$$

- Overload probability is extremely small when $k \ll n$!
- Might seem counterintuitive – bound gets worse as k grows.
- When k is large, the number of requests each server sees in expectation is very small so the law of large numbers doesn't 'kick in'.

n: total number of requests, k: number of servers randomly assigned requests, R_i: number of requests assigned to server i.

What is the probability that the maximum server load exceeds $2 \cdot \mathbb{E}[R_i] = \frac{2n}{k}$. I.e., that some server is overloaded if we give each $\frac{2n}{k}$ capacity?

$$\Pr\left(\max_i(R_i) \geq \frac{2n}{k}\right) = \Pr\left(\left[R_1 \geq \frac{2n}{k}\right] \cup \left[R_2 \geq \frac{2n}{k}\right] \cup \ldots \cup \left[R_k \geq \frac{2n}{k}\right]\right) = \Pr$$

We want to show that $\Pr\left(\bigcup_{i=1}^{k}\left[R_i \geq \frac{2n}{k}\right]\right)$ is small.

How do we do this? Note that $R_1, \ldots, R_k$ are correlated in a somewhat complex way.

n: total number of requests, k: number of servers randomly assigned requests, R_i: number of requests assigned to server i. $\mathbb{E}[R_i] = \frac{n}{k}$. $\mathrm{Var}[R_i] = \frac{n}{k}$.

Union Bound: For any random events $A_1, A_2, ..., A_k$,

$$\Pr\left(A_1 \cup A_2 \cup \ldots \cup A_k\right) \leq \Pr(A_1) + \Pr(A_2) + \ldots + Pr(A_k).$$

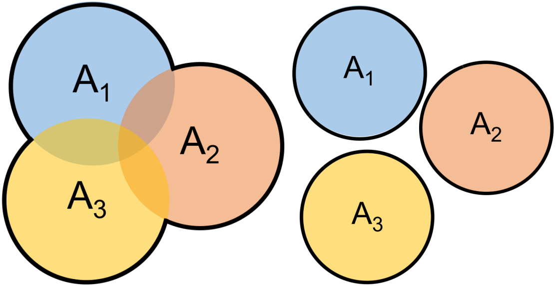

When is the union bound tight? When $A_1, ..., A_k$ are all disjoint.

On the first problem set, you will prove the union bound, as a consequence of Markov's inquality.

APPLYING THE UNION BOUND

What is the probability that the maximum server load exceeds $2 \cdot \mathbb{E}[\mathbf{R}_i] = \frac{2n}{k}$. I.e., that some server is overloaded if we give each $\frac{2n}{k}$ capacity?

$$\Pr\left(\max_i(\mathbf{R}_i) \geq \frac{2n}{k}\right) = \Pr\left(\bigcup_{i=1}^{k}\left[\mathbf{R}_i \geq \frac{2n}{k}\right]\right)$$

$$\leq \sum_{i=1}^{k} \Pr\left(\left[\mathbf{R}_i \geq \frac{2n}{k}\right]\right) \qquad \text{(Union Bound)}$$

$$\leq \sum_{i=1}^{k} \frac{k}{n} = \frac{k^2}{n} \qquad \text{(Bound from Chebyshev's)}$$

As long as $k \leq O(\sqrt{n})$, with good probability, the maximum server load will be small (compared to the expected load).

> n: total number of requests, k: number of servers randomly assigned requests, $\mathbf{R}_i$: number of requests assigned to server i. $\mathbb{E}[\mathbf{R}_i] = \frac{n}{k}$. $\text{Var}[\mathbf{R}_i] = \frac{n}{k}$.

The number of servers must be small compared to the number of requests ($k = O(\sqrt{n})$) for the maximum load to be bounded in comparison to the expected load with good probability.

- There are many requests routed to a relatively small number of servers so the load seen on each server is close to what is expected via law of large numbers.
- **A More Natural Variant:** Given n requests, and assuming all servers have fixed capacity C, how many servers should you provision so that with probability $\geq 99/100$ no server is assigned more than C requests?

n: total number of requests, k: number of servers randomly assigned requests.

Questions on union bound, Chebyshev's inequality, random hashing?

We flip $n = 100$ independent coins, each are heads with probability 1/2 and tails with probability 1/2. Let H be the number of heads.

$$\mathbb{E}[\mathrm{H}] = \frac{n}{2} = 50 \text{ and } \mathrm{Var}[\mathrm{H}] = \frac{n}{4} = 25 \rightarrow s.d. = 5$$

Markov's:	Chebyshev's:	In Reality:
$\Pr(\mathrm{H} \geq 60) \leq .833$	$\Pr(\mathrm{H} \geq 60) \leq .25$	$\Pr(\mathrm{H} \geq 60) = 0.0284$
$\Pr(\mathrm{H} \geq 70) \leq .714$	$\Pr(\mathrm{H} \geq 70) \leq .0625$	$\Pr(\mathrm{H} \geq 70) = .000039$
$\Pr(\mathrm{H} \geq 80) \leq .625$	$\Pr(\mathrm{H} \geq 80) \leq .0278$	$\Pr(\mathrm{H} \geq 80) < 10^{-9}$

H has a simple Binomial distribution, so can compute these probabilities exactly.

To be fair.... Markov and Chebyshev's inequalities apply much more generally than to Binomial random variables like coin flips.

Can we obtain tighter concentration bounds that still apply to very general distributions?

- Markov's: $\Pr(X \geq t) \leq \frac{\mathbb{E}[X]}{t}$. First Moment.
- Chebyshev's: $\Pr(|X - \mathbb{E}[X]| \geq t) = \Pr(|X - \mathbb{E}[X]|^2 \geq t^2) \leq \frac{\mathrm{Var}[X]}{t^2}$. Second Moment.
- What if we just apply Markov's inequality to even higher moments?