Lecture 3: Concentration, Randomized Paradigms, Hypercube Routing

February 4, 2026 Lecturer: D Ellis Hershkowitz Scribe: Na Nguyen, Nuttanon Tangsunthornwiwat

1 Motivating Scenarios: Balls into Bins

Setup. We throw n balls independently and uniformly at random into m bins. Let X_i denote the number of balls in bin i.

Questions.

1. If $m = n$, what is the maximum load $\max_i X_i$?
2. How large must m be so that every bin is nonempty with high probability?

The first question concerns bounding the maximum of many random variables, while the second is the classical *coupon collector problem*. Both require tools that show random variables concentrate near their expectations.

2 The Concentration Framework

We follow a general three-step strategy for proving concentration results.

1. **Local concentration:** Show a single random variable is tightly concentrated around its expectation.
2. **Union bound:** Extend the bound to many random variables simultaneously.
3. **Conclusion:** Deduce a global high-probability statement.

Reminder (Union Bound). For any events $B_1, \dots, B_k$,

$$\mathbb{P}\left(\bigcup_{i=1}^{k} B_i\right) \le \sum_{i=1}^{k} \mathbb{P}(B_i).$$

Let

$$X_{ij} := \mathbb{1}\{\text{ball } j \text{ lands in bin } i\}, \qquad X_i := \sum_{j=1}^{n} X_{ij}.$$

Then X_i counts the number of balls in bin i.

3 Markov's Inequality

Theorem 1 (Markov)**.** *Let $X \geq 0$ be a random variable with $\mathbb{E}[X] > 0$. Then for any $a > 0$,*

$$\mathbb{P}(X \geq a\mathbb{E}[X]) \leq \frac{1}{a}.$$

Proof. We write

$$\mathbb{E}[X] = \sum_x x\mathbb{P}(X = x) \geq \sum_{x \geq a\mathbb{E}[X]} x\mathbb{P}(X = x) \geq a\mathbb{E}[X]\mathbb{P}(X \geq a\mathbb{E}[X]).$$

Dividing both sides by $a\mathbb{E}[X]$ gives the result. □

Markov's inequality requires only non-negativity but yields weak, linear decay.

4 Chebyshev's Inequality

[Chebyshev] For any random variable X with variance $\mathrm{Var}(X)$,

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

Proof. Let $Y := (X - \mathbb{E}[X])^2$, so $\mathbb{E}[Y] = \mathrm{Var}(X)$. Applying Markov's inequality,

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq a) = \mathbb{P}(Y \geq a^2) \leq \frac{\mathbb{E}[Y]}{a^2}.$$

□

Chebyshev gives quadratic decay but depends heavily on variance bounds.

5 Chernoff Bounds

Chernoff bounds give *exponential* concentration for sums of independent indicator random variables.

Theorem 2 (Simple Chernoff)**.** *Let $X = \sum_{i=1}^n X_i$ where $X_i \in \{0, 1\}$ are independent and let $\mu = \mathbb{E}[X]$. Then for any $\delta > 0$,*

$$\mathbb{P}(X \geq (1 + \delta)\mu) \leq \exp\left(-\frac{\delta^2\mu}{2 + \delta}\right),$$

and for $\delta \in (0, 1)$,

$$\mathbb{P}(X \leq (1 - \delta)\mu) \leq \exp\left(-\frac{\delta^2\mu}{2}\right).$$

Chernoff bounds yield exponentially small failure probabilities, which enables "with high probability" guarantees when $\mu = \Omega(\log n)$.

6 Applying the Framework to Balls into Bins

6.1 Case $m = n$

Fix a bin i. Then $\mathbb{E}[X_i] = 1$. By Chernoff,

$$\mathbb{P}(X_i \geq 2 + 10\log n) \leq n^{-2}.w$$

Applying a union bound over all n bins,

$$\mathbb{P}(\exists i : X_i \geq 2 + 10\log n) \leq \frac{1}{n}.$$

Thus $\max_i X_i = O(\log n)$ with high probability.

6.2 Coupon Collector

We have n bins and throw m balls independently and uniformly at random. Let X_i denote the number of balls in bin i.

Then

$$\mathbb{E}[X_i] = \frac{m}{n}.$$

The coupon collector problem asks how large must m be so that every bin is nonempty with high probability.

To ensure $\Pr(X_i = 0)$ is small for a fixed bin, we require

$$\mathbb{E}[X_i] = \frac{m}{n} = \Theta(\log n),$$

which implies

$$m = \Theta(n \log n).$$

For a fixed bin i,

$$\mathbb{P}(X_i = 0) = \left(1 - \frac{1}{n}\right)^m \leq \exp\left(-\frac{m}{n}\right).$$

Setting $m = cn\log n$ for a sufficiently large constant c gives

$$\Pr(X_i = 0) \leq n^{-c}.$$

Applying a union bound over all bins,

$$\mathbb{P}(\exists i \text{ such that } X_i = 0) \leq n \cdot n^{-c} = n^{1-c}.$$

Choosing $c > 2$ implies that with high probability, every bin is nonempty.

With n bins, $\Theta(n \log n)$ balls are both necessary and sufficient to ensure all bins are nonempty with high probability.

7 Routing in the Hypercube

Let $G = (V, E)$ be the n-node hypercube with $V = \{0,1\}^{\log n}$.

A *unit demand* is a function

$$D : V \times V \to \{0, 1\}$$

such that each node has at most one incoming and one outgoing demand.

Definition 1. *The congestion of an edge e under a set of paths $\mathcal{P}$ is*

$$\mathrm{con}_{\mathcal{P}}(e) = |\{p \in \mathcal{P} : e \in p\}|\,.$$

The congestion of $\mathcal{P}$ is

$$\mathrm{con}_{\mathcal{P}} = \max_{e \in E} \mathrm{con}_{\mathcal{P}}(e).$$

7.1 Valiant's Randomized Routing

Any unit demand in the n-node hypercube can be routed with $O(\log n)$ congestion.

Valiant's Algorithm

For each (u, v) with $D(u, v) = 1$:

- Pick $w_{u,v} \in V$ uniformly at random.
- Route $u \to w_{u,v} \to v$ by fixing bits left-to-right.

7.2 Congestion Analysis

Let $\mathcal{P}_1$ be the first-leg paths and $\mathcal{P}_2$ the second. Then

$$\mathrm{con}(\mathcal{P}) \le \mathrm{con}(\mathcal{P}_1) + \mathrm{con}(\mathcal{P}_2).$$

Fix an edge e at bit position i. Let X_e be the number of paths in $\mathcal{P}_1$ using e. We show $\mathbb{E}[X_e] \le 2$.

Applying Chernoff,

$$\Pr(X_e \ge 2 + 10 \log n) \le n^{-5}.$$

Taking a union bound over all $O(n^2)$ edges,

$$\Pr(\mathrm{con}(\mathcal{P}_1) > 2 + 10 \log n) \le n^{-3}.$$

By symmetry, the same holds for $\mathcal{P}_2$.

Conclusion. With probability at least $1 - 2n^{-3}$,

$$\mathrm{con}(\mathcal{P}) = O(\log n).$$

8 Another Section

Given a graph $G = (V, E)$, a demand is a function

$$D : V \times V \to \{0, 1\}.$$

A demand D is unit iff

$$\sum_{v \in V} D(v, u) \leq 1 \quad \text{and} \quad \sum_{u \in V} D(v, u) \leq 1 \qquad \forall v \in V.$$

(Directed) paths $\mathcal{P}$ route D; if $D(u, v) = 1$, there is a path from u to v.

The congestion of edge e is

$$\text{con}_{\mathcal{P}}(e) = |\{ p \in \mathcal{P} : e \in p \}| .$$

The congestion of paths $\mathcal{P}$ is

$$\text{con}_{\mathcal{P}} = \max_{e \in E} \text{con}_{\mathcal{P}}(e).$$

The n-node hypercube is the graph with $V = \{0, 1\}^{\log n}$ and

$\{u, v\} \in E \quad \text{iff} \quad u, v$ differ in exactly one bit.

Claim. Can route any unit demand in the n-node hypercube with $O(\log n)$ congestion.

Let $p(u, v)$ be a path from u to v that fixes bits from left to right.

Valiant's Algorithm

$\mathcal{P} = \varnothing$

For each (u, v) s.t. $D(u, v) = 1$:

- Let $w_{u,v}$ be a uniformly random node.
- Add $p(u, w_{u,v})$ and $p(w_{u,v}, v)$ to $\mathcal{P}$

Return $\mathcal{P}$.

$\mathcal{P}$ routes D.

Let

$$\mathcal{P}_1 = \{ p(u, w_{u,v}) : D(u, v) = 1 \} \quad \text{and} \quad \mathcal{P}_2 = \{ p(w_{u,v}, v) : D(u, v) = 1 \}.$$

$$\text{con}(\mathcal{P}) \leq \text{con}(\mathcal{P}_1) + \text{con}(\mathcal{P}_2).$$

Firstly, bound $\text{con}(\mathcal{P}_1)$.

Fix $e = \{x, y\} \in E$ where

$$x = a_1, a_2, \ldots, a_i, \ldots, a_{\log n}$$

and

$$y = a_1, a_2, \ldots, \bar{a}_i, \ldots, a_{\log n}.$$

Consider $u, w_{u,v}, v$ s.t. $D(u, v) = 1$.

$$u = u_1, u_2, \ldots, u_i, \ldots, u_{\log n}$$

$$w_{u,v} = w_1, w_2, \ldots, w_i, \ldots, w_{\log n}.$$

$p(u, w_{u,v})$ only uses e if:

$$1)\ w_1 = a_1,\ w_2 = a_2,\ \ldots,\ w_{i-1} = a_{i-1}$$

$$2)\ u_{i+1} = a_{i+1},\ u_{i+2} = a_{i+2},\ \ldots,\ w_{\log n} = a_{\log n}.$$

$X_u := \mathbb{1}(w_{u,v}, u \text{ satisfies (1) and (2)})$.

$X_e := \sum_u X_u$.

Then,

$$\begin{aligned}
\mathbb{E}[X_e] &= \mathbb{E}\left[\sum_u X_u\right] \\
&= (\#\, u \text{ that satisfies (2)}) \cdot \mathbb{P}(w_{u,v} \text{ satisfies (1)}) \\
&\leq 2^i \cdot \frac{1}{2^{\,i-1}} \\
&\leq 2.
\end{aligned}$$

Let

$$B_e := \{X_e \geq 2 + 10 \log n\}.$$

Then,

$$\begin{aligned}
\mathbb{P}(B_e) &= \mathbb{P}(X_e \geq 2 + 10 \log n) \\
&\leq \exp\left(-\frac{2 \cdot 25 \log^2 n}{2 + 5 \log n}\right) \qquad \text{(by Chernoff bound)} \\
&\leq \exp(-5 \log n) \\
&\leq n^{-5}.
\end{aligned}$$

Then,

$$\begin{aligned}\mathbb{P}\left(\bigcup_e B_e\right) &\leq \sum_e \mathbb{P}(B_e) \\ &\leq n^2 \cdot n^{-5} \\ &\leq n^{-3}.\end{aligned}$$

$$\therefore \; \mathbb{P}(\mathrm{con}(\mathcal{P}_1) > 2 + 10\log n) \leq n^{-3}.$$

By symmetry,

$$\mathbb{P}(\mathrm{con}(\mathcal{P}_2) > 2 + 10\log n) \leq n^{-3}.$$

$$\therefore \; \mathbb{P}(\mathrm{con}(\mathcal{P}) > 4 + 20\log n) \leq 2\,n^{-3}.$$