Lecture 06: The power of two choices

The Power of 2-Choices

- We throw $m = n$ balls into n bins in the following manner
  1. Each ball chooses two bins (both) uniformly and independently at random
  2. The ball is put into the bin that has lower load (at that time). If both the bins have an identical number of balls, then put the ball in either of the bins (i.e., break the tie arbitrarily)
- We are interested in studying the Max Load of this experiment
- In a seminal paper Azar, Broder, Karlin and Upfal showed that, with high probability, the Max Load is at most $\log \log n + O(1)$
- Note that this is exponentially better than random allocation, in which case the max load is $\approx \log n / \log \log n$ as discussed in the previous lecture
- If d > 2 choices are used to place each ball, then there is not much improvement. The Max Load is at most $\log \log n / \log d + O(1)$

Objective of the Lecture

- The objective of this lecture note is to assist the students read the proof presented in Section 1.2 of the Ph.D. thesis of Michael Mitzenmacher
- At this point of time, we still do not know two key concepts to completely write down the proof of this theorem
  - Coupling Argument: We shall not see this formally introduced in this course. Please refer to online resources to read this.
  - Chernoff Bound: We shall see this formally introduced in the next lecture. It is highly recommended that students revisit this lecture after the next lecture.
- The lecture note will introduce the main idea of the proof. A small extremal case in the analysis will be left and students are recommended to look it up from Section 1.2 of Michael Mitzenmacher?s Ph.D. thesis. The main reason is that, we want to sift the "key-ideas" from the mechanical step of "plugging outlier cases"

Intuitive Overview of the Proof I

- Let $\mathbb{X}_t$ be the bin where the t-th ball lands
- Note that $(\mathbb{X}_1, \dots, \mathbb{X}_n)$ defines exactly where each ball went and defines the entire state of the experiment
- Terminology for today's lecture
  - "At time t" of the experiment represents when the t-th ball is thrown.
  - "Just before time t" refers to the state of the experiment just before the t-th ball is thrown.
  - "Just after time t" refers to the state of the experiment just after the t-th ball is thrown

Intuitive Overview of the Proof II

- Let $(\beta_1, \beta_2, \ldots, \beta_n)$ be the thresholds defined as follows.

  "We expect that the number of bins with load $\geqslant i$ at the end of time n (i.e., the end of the experiment) to be less than β_i"

  Formally,

$$\beta_i \geqslant \mathbb{E}\left[\sum_{j=1}^{n} \mathbf{1}_{\{\mathbb{L}_j \geqslant i\}}\right]$$

- **A new set of random variables.** The random variable $\#\mathrm{Bins}_{\geqslant i}(t)$ represents the number of bins with load $\geqslant i$ at the end of time t

How to find the thresholds?

- Suppose we are already given that $\#\text{Bins}_{\geqslant i}(n) \leqslant \beta_i$ is true
- Conditioned on this fact, we want to compute a likely upper-bound on $\#\text{Bins}_{\geqslant i+1}$. This upper-bound we shall define to be β_{i+1}

Intuitive Overview of the Proof IV

- Height of a Ball. Imagine the bins to be narrow, and balls stack up on each other as they are allocated to a bin. The height of a ball is the "number of balls below it plus one." For example, the first ball in a bin has height 1, the second ball in that bin has height 2, and so on. Note that future ball allocations do not change the heights of the balls that have already been assigned.
- **A set of new random variables.** The random variable $\#\mathrm{Balls}_{\geqslant i}(t)$ represents the number of balls that have height $\geqslant i$ at the end of time t

Intuitive Overview of the Proof V

Observation

We always have

$$\#\text{Balls}_{\geqslant i}(t) \geqslant \#\text{Bins}_{\geqslant i}(t)$$

Proof.

At the end of time t, any bin with load $\geqslant i$ also has at least one ball with height $\geqslant i$. $\square$

Intuitive Overview of the Proof VI

The beginning of Inductive Step.

- Note that for a ball to land at height $\geqslant (i+1)$, it must be the case that both the bins chosen to allocate it already has $\geqslant i$ balls

The Hindsight Argument.

- Conditioned on $\#\mathrm{Bins}_{\geqslant i}(n)$ being $\leqslant \beta_i$, the probability of choosing both bins that have height $\geqslant i$ (at any time) is at most
$$p_i = \left(\frac{\beta_i}{n}\right)^2$$
- At each time $t \in \{1, 2, \dots, n\}$, the t-th ball has height $\geqslant (i+1)$ with probability $\leqslant p_i$.
- Therefore, the expected number of balls that have height $\geqslant (i+1)$ is at most np_i.
- So, we set $\beta_{i+1} = np_i = \frac{\beta_i^2}{n}$

Intuitive Overview of the Proof VII

Finishing up the proof.

- So, conditioned on the fact that

$$\#\mathrm{Bins}_{\geqslant i}(n) \leqslant \beta_i$$

  we shall show the following happens with high probability

$$\#\mathrm{Bins}_{\geqslant (i+1)}(n) \leqslant \beta_{i+1}$$

- Note that we can set $\beta_2 = n/2$ (because, the number of bins with two-or-more balls is obviously less than $n/2$)
- Therefore, we get $\beta_{i+2} n/2^{2^i}$ as the solution of the recursion
- For $i > i^* = \log\log n$, we have $\beta_{i+2} < 1$. This means that no bins have load $\geqslant i^* + 2$ at the end of time n. This proves the upper-bound on the max-load of the experiment

Formal Proof Outline I

- Is the actual proof, we shall use β_is that have a slight "slack" in-built
- We shall start with $\beta_6 = n/2\mathrm{e}$ (note that the number of bins that have 6 or more balls is at most $n/6 < n/2\mathrm{e}$.
- We shall recursively define

$$\beta_{i+1} = \mathrm{e}\frac{\beta_i^2}{n}$$

- We define

$$p_i = \left(\frac{\beta_i}{n}\right)^2$$

Formal Proof Outline II

- The event $\mathbb{G}_i$ represents the "good event" that $\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i$
- We shall show that

$$\mathbb{P}\left[\mathbb{G}_6, \mathbb{G}_7, \dots, \mathbb{G}_n\right] \approx 1$$

- To show this, we shall show that

$$\mathbb{P}\left[\overline{\mathbb{G}_6} \text{ or } \overline{\mathbb{G}_7} \text{ or } \cdots \text{ or } \overline{\mathbb{G}_n}\right] \leqslant \frac{1}{n}$$

- To prove this, we shall show that

$$\mathbb{P}\left[\overline{\mathbb{G}_6}\right] = 0$$

$$\mathbb{P}\left[\overline{\mathbb{G}_{i+1}}, \mathbb{G}_i\right] \leqslant \frac{1}{n^2} \qquad \text{For } i \in \{6, \dots, n-1\}$$

The statement above will follow by a "Union Bound" kind of argument. Think how to prove it.

Formal Proof Outline III

- Let us begin our formal analysis

$$
\begin{aligned}
&\mathbb{P}\left[\overline{\mathbb{G}_{i+1}} | \mathbb{G}_i\right] \\
=&\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant(i+1)} > \beta_{i+1} | \#\mathsf{Bins}_{\geqslant i} \leqslant \beta_i\right] \\
\leqslant&\mathbb{P}\left[\#\mathsf{Ball}_{\geqslant(i+1)} > \beta_{i+1} | \#\mathsf{Bins}_{\geqslant i} \leqslant \beta_i\right]
\end{aligned}
$$

  The last inequality is due to the observation that every bin with load $\geqslant (i+1)$ has at least one ball with height $\geqslant (i+1)$

- For $t \in \{1, \ldots, n\}$, let $\mathbb{Y}_t$ be the indicator variable for the event that the t-th ball throw chose both bins with load $\geqslant i$

Formal Proof Outline IV

- Note that

$$\mathbb{P}\left[\mathbb{Y}_t = 1 | \mathbb{X}_1, \dots, \mathbb{X}_{t-1}, \mathbb{G}_i\right] \leqslant p_i = \frac{\beta_i^2}{n^2}$$

  This is because, we have $\#\text{Bins}_{\geqslant i}(n) \leqslant \beta_i$. So, we have $\#\text{Bins}_{\geqslant i}(t) \leqslant \beta_i$, for all $t \in \{1, \dots, n\}$

- Note that

$$\#\text{Balls}_{\geqslant(i+1)}(n) = \sum_{t=1}^{n} \mathbb{Y}_t$$

Formal Proof Outline V

- So, we have

$$
\begin{aligned}
&\mathbb{P}\left[\#\text{Balls}_{\geqslant(i+1)}(n) > \beta_{i+1} | \#\text{Bins}_{\geqslant i}(n) \leqslant \beta_i\right] \\
=&\mathbb{P}\left[\sum_{t=1}^{n} \mathbb{Y}_t > \beta_{i+1} | \#\text{Bins}_{\geqslant i}(n) \leqslant \beta_i\right] \\
\leqslant&\mathbb{P}\left[B(n, p_i) > \beta_{i+1} | \#\text{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]
\end{aligned}
$$

Here $B(n, p_i)$ is the sum of n i.i.d. Bernoulli trials, each of which are 1 with probability p_i. The last inequality is due to a "Coupling Argument." In this course, we shall not see this topic. However, students are encouraged to read it online

Formal Proof Outline VI

- We continue our analysis

$$\begin{aligned}
&\mathbb{P}\left[B(n,p_i) > \beta_{i+1} \middle| \#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right] \\
=&\frac{\mathbb{P}\left[B(n,p_i) > \beta_{i+1}, \#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]} \\
\leqslant&\frac{\mathbb{P}\left[B(n,p_i) > \beta_{i+1}\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]}
\end{aligned}$$

  The last inequality is due to the fact that $\mathbb{P}[\mathbb{A},\mathbb{B}] \leqslant \mathbb{P}[\mathbb{A}]$

- Note that $\beta_{i+1} = \mathrm{e} \cdot np_i$
- Continuing our analysis

$$\begin{aligned}
&\frac{\mathbb{P}\left[B(n,p_i) > \beta_{i+1}\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]} \\
=&\frac{\mathbb{P}\left[B(n,p_i) > \mathrm{e} \cdot np_i\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]}
\end{aligned}$$

Formal Proof Outline VII

- By "Chernoff Bound," (a topic that we shall cover in the next lecture) we know that it is unlikely that $B(n, p_i)$ shall exceed its expected value by e times. This shall be covered in the next lecture. Continuing our expansion

$$\frac{\mathbb{P}\left[B(n, p_i) > \mathrm{e} \cdot np_i\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant i}(n) \leqslant \beta_i\right]}$$

$$\leqslant \exp(-np_i) \cdot \frac{1}{\mathbb{P}\left[\mathbb{G}_i\right]}$$

- When $np_i \geqslant 2\log n$, we have $\exp(-np_i) \leqslant 1/n^2$. The technique of dealing with the remaining case of $np_i < 2\log n$ is left as a reading exercise for the students
- So, we have obtained the following

$$\mathbb{P}\left[\overline{\mathbb{G}_{i+1}} | \mathbb{G}_i\right] \leqslant \frac{1}{n^2} \frac{1}{\mathbb{P}\left[\mathbb{G}_i\right]}$$

Formal Proof Outline VIII

- Cross-multiplying, we get

$$\mathbb{P}\left[\overline{\mathbb{G}_{i+1}}, \mathbb{G}_i\right] \leqslant \frac{1}{n^2}$$

- This completes the proof