Lecture 04: The Power of 2-Choices

The Power of 2-Choices

- n balls are thrown into n bins in the following manner
  - Each ball chooses two bins (both) uniformly at random
  - The ball is put into the bin that has lesser number of balls (at that time)
  - If both bins have 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
- 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, where max load is $\approx \log n / \log \log n$
- If $d > 2$ choices are used to place each ball, then there is not much improvement. The Max Load is at most $\frac{\log \log n}{\log d} + O(1)$

Aim of this Lecture Note

- The aim of this lecture note is to assist student read the proof presented in Section 1.2 of the Ph.D. thesis of Michael Mitzenmacher
- We still do *not* know two key concepts to completely write down the proof of this theorem
  - Coupling Argument: We will not see this formally introduced in this course. Please refer to online resources to read this.
  - Chernoff Bound: We will 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 typical step of "plugging outlier cases"

Intuitive Overview of the Proof

- Let ω_t be the bin where the t-th ball lands
- Note that $(\omega_1, \ldots, \omega_n)$ defines exactly where each ball went and defines the entire state of the experiment
- In our terminology, "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 $(\beta_1, \beta_2, \dots)$ be thresholds that we define such that: 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
- Random Variable: $\#\text{Bins}_{\geqslant i}(t)$ represents the number of bins with load $\geqslant i$ at the end of time t
- 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 $\#\text{Bins}_{\geqslant i+1}(n)$

- Height of a Ball: We imagine the bins to be narrow and balls to stack up on them as them are allocated to a bin. The height of a ball is the "number of balls below it plus one." For example, the first ball into a bin has height 1, the second ball into that bin has height 2, and so on. Note that future ball allocations does not change the heights of the balls that have already been assigned.
- Random Variable: $\#\mathrm{Balls}_{\geqslant i}(t)$ represents the number of balls that have height $\geqslant i$ at the end of time t
- Observation: For all ball allocations and time t, the quantity $\#\mathrm{Balls}_{\geqslant i}(t)$ is always larger than $\#\mathrm{Bins}_{\geqslant i}(t)$. Because each bin that has $\geqslant i$ balls has at least one ball with height $\geqslant i$.

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 $\#\text{Bins}_{\geqslant i}(n)$ being $\leqslant \beta_i$, the probability of choosing both bins that have height $\geqslant i$ is at most

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

- In each time step $t \in \{1, \ldots, n\}$, the t-th ball has height $\geqslant (i+1)$ with probability at most p_i
- Therefore, the number of balls that have height $\geqslant (i+1)$ is expected to be at most np_i
- So, we set $\beta_{i+1} = np_i = n\left(\frac{\beta_i}{n}\right)^2$

- So, conditioned on the fact that

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

We shall show that the following happens with high probability

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

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

Formal Proof Outline

- In the actual proof, we will use β_is that has a slight slack in built
- We shall use $\beta_6 = n/2\mathrm{e}$ as the base case (note that the number of bins that have 6 or more balls is at most $n/6 \leqslant n/2\mathrm{e}$)
- We shall recursively define

$$\beta_{i+1} = \mathrm{e} \cdot n \left(\frac{\beta_i}{n}\right)^2$$

- We still define

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

Formal Proof Outline II

- Event: $\mathcal{E}_i$ is the event that $\#\mathrm{Bins}_{\geqslant i}(n) \leqslant \beta_i$
- We will show that

$$\mathbb{P}\left[\mathcal{E}_6, \ldots, \mathcal{E}_n\right] \approx 1$$

- To show this, we will show that

$$\mathbb{P}\left[\overline{\mathcal{E}_6} \vee \cdots \vee \overline{\mathcal{E}_n}\right] \leqslant 1/n$$

- To prove this, we will show that:

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

$$\mathbb{P}\left[\overline{\mathcal{E}_{i+1}}, \mathcal{E}_i\right] \leqslant 1/n^2 \qquad \text{For } i \geqslant 6$$

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

$$
\begin{aligned}
&\mathbb{P}\left[\overline{\mathcal{E}_{i+1}}|\mathcal{E}_i\right] \\
=&\mathbb{P}\left[\#\mathrm{Bins}_{\geqslant \mathrm{i+1}}(n) > \beta_{i+1} \mid \#\mathrm{Bins}_{\geqslant \mathrm{i}}(\mathrm{n}) \leqslant \beta_i\right] \\
\leqslant&\mathbb{P}\left[\#\mathrm{Balls}_{\geqslant \mathrm{i+1}}(n) > \beta_{i+1} \mid \#\mathrm{Bins}_{\geqslant \mathrm{i}}(n)\right]
\end{aligned}
$$

The last inequality is due to the fact that every ball 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 two bins with load $\geqslant i$.
- Note that

$$
\mathbb{P}\left[\mathbb{Y}_t = 1|\omega_1, \ldots, \omega_{t-1}\right] \leqslant p_i = \left(\frac{\beta_i}{n}\right)^2
$$

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

- Note that

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

- So

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

- 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 will *not* see this topic. But students are encouraged to read it online.

- Now, we continue our analysis

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

- The last inequality is due to the fact that $\mathbb{P}[A, B] \leqslant \mathbb{P}[A]$
- Note that $\beta_{i+1} = \mathrm{e} \cdot np_i$
- Continuing our analysis

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

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

$$\frac{\mathbb{P}\left[B(n, p_i) > \mathrm{e} \cdot np_i\right]}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant \mathrm{i}}(n)\right]} \leqslant \frac{1}{\exp\left(np_i\right)} \cdot \frac{1}{\mathbb{P}\left[\#\mathsf{Bins}_{\geqslant \mathrm{i}}(n)\right]}$$

- When $np_i \geqslant 2\log n$, we have $\frac{1}{\exp(np_i)} \leqslant \frac{1}{n^2}$
- How to deal with the remaining case of $np_i < 2\log n$ is left as a reading exercise for the students.

- So, we have obtained:

$$\mathbb{P}\left[\overline{\mathcal{E}_{i+1}} \mid \mathcal{E}_i\right] \leqslant \frac{1}{n^2} \cdot \frac{1}{\mathbb{P}\left[\mathcal{E}_i\right]}$$

- Cross-multiplying, we get

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

- This completes the proof