MIT 6.854/18.415: Advanced Algorithms Spring 2016

Lecture 2 – February 8, 2016

Prof. Ankur Moitra *Scribe: Calvin Huang, Linh V. Nguyen*

In this lecture, we analyze the problem of scheduling n equal size tasks arriving online to n different machines. This could easily be solved with a centralized server that knows which machines are occupied and assigns all n tasks to unique machines. Instead, we want to see if it's possible to distribute the tasks without without such coordination among the tasks. In this problem, the metric we care about is the maximum load on any given machine.

We will consider two such approaches: naive random assignment, and the "power of two choice".

1 Random Assignment

One approach might be that each task chooses a machine independently at random. This is the classic "balls-in-bins" (ball = task, bin = machine). Let z_i be the number of balls (tasks) assigned to bin (machine) i. There are a couple ways we might try to put a bound on the maximum load.

1.1 Markov's inequality

Theorem 1. *Markov's inequality: If X is a nonnegative random variable, then for any $a > 0$*

$$\mathbb{P}[X \geq a] \leq \frac{\mathbb{E}[X]}{a}$$

Using Markov's inequality, since $\mathbb{E}[z_i] = 1$ (on average, each machine will receive 1 task) we get, for some k, $\mathbb{P}(z_i \geq k) \leq \frac{1}{k}$, and using the union bound $\mathbb{P}(\exists i : z_i \geq k) \leq \frac{n}{k}$, which is not very useful (for all practical values of k, it tells us that the probability is at most 1).

1.2 Chebyshev's inequality

Theorem 2. *Chebyshev's inequality: Let X be a random variable with expectation μ and variance σ^2. Then*

$$\mathbb{P}[|X - \mu| \geq k\sigma] \leq \frac{1}{k^2}$$

Let $y_{i,j}$ be random indicator variable, defined so that $y_{i,j} = 1$ if the j-th ball is thrown into the i-th bin, and $y_{i,j} = 0$ otherwise. Note that $y_{i,j} = 1$ with probability of $\frac{1}{n}$, and $\text{Var}(y_{i,j}) = \frac{1}{n} - \frac{1}{n^2}$.

Then for $z_i = \sum_{j=1}^n y_{i,j}$, we get $\mathbb{E}[z_i] = 1$ and

$$\begin{aligned}
\mathrm{Var}(z_i) &= \sum_{j=1}^n \mathrm{Var}(y_{i,n}) \\
&= n(1/n - 1/n^2) \\
&= 1 - \frac{1}{n} \\
&< 1.
\end{aligned}$$

Using Chebyshev's inequality and a union bound:

$$\begin{aligned}
\mathbb{P}(z_i - 1 \geq k) \leq \mathbb{P}(z_i - \mu \geq k\sigma) &\leq \frac{1}{k^2} \\
\mathbb{P}(\exists i : z_i \geq k+1) &\leq \frac{n}{k^2}
\end{aligned} \tag{1}$$

and so we can place a constant bound on the probability that any bin has more than k balls, if $k > \sqrt{n}$.

1.3 Direct analysis

In a more direct analysis, we take a union bound over all possible subsets of k balls that might be sent to machine i:

$$\mathbb{P}(z_i \geq k) \leq \sum_{S \subset [n], \|S\|=k} \mathbb{P}(\text{all balls in } S \text{ sent to machine } i) = \binom{n}{k} \frac{1}{n^k} \leq \left(\frac{e}{k}\right)^k \tag{2}$$

Here, we used the fact that $\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$.

The bound $\left(\frac{e}{k}\right)^k$ is on the order of $\frac{1}{n}$ when $k = \Theta\left(\frac{\log n}{\log \log n}\right)$, so we can put a constant bound on the probability that no bin will have more than $\Theta\left(\frac{\log n}{\log \log n}\right)$ balls.

1.4 Chernoff bound

Theorem 3. *Upper tail of Chernoff bound: Let X_i be independent Bernoulli random variables with $\mathbb{E}[\sum X_i] = \mu$. Then*

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

In this case the X_i are whether or not each ball falls into a particular bin, so $\mu = 1, k = 1+\delta$, and so we get $\mathbb{P}(\sum X_i \geq k) \leq \frac{e^{k-1}}{k^k}$, which gives the same result as in direct analysis – $\mathbb{P}(\sum X_i \geq k) = O\left(\frac{1}{n}\right)$ when $k = \Theta\left(\frac{\log n}{\log \log n}\right)$.

2 Power of Two Choices

This is a simple variation on the fully random method: instead of picking one machine at random, each task picks *two* random machines, and assigns itself to the machine with the lower load. This simple change will drastically lower the expected max load from $\Theta\left(\frac{\log n}{\log\log n}\right)$ to $\Theta(\log\log n)$.

We'll start with a heuristic justification of this improvement, before giving a more rigorous proof.

2.1 Heuristic Analysis

Consider this nonrigorous analysis to help understand where the $\Theta(\log\log n)$ bound arises.

Let β_i be an upper bound on the number of bins with at least i balls at the end of the process. Now, consider the number of bins which will contain at least $i+1$ balls. When a ball arrives, in order for it to increase a bin to $i+1$, both random bins picked must have at least i balls. The probability of a given random choice ending up that way would be at most $\frac{\beta_i}{n}$; since that must happen twice independently of such an increase would be at most $\left(\frac{\beta_i}{n}\right)^2$. We'd then expect, over all n arrivals, at most $n\left(\frac{\beta_i}{n}\right)^2 = \frac{\beta_i^2}{n}$ bins containing at least $i+1$ balls, and thus set $\beta_{i+1} = \frac{\beta_i^2}{n}$.

Let $\beta_6 = \frac{n}{2e} < \frac{n}{6}$ (must be true by pigeonhole principle), then by induction $\beta_i \leq \frac{n}{(2e)^{2^{i-6}}}$, therefore the lowest i_c such that $\beta_{i_c} < 1$ will be $O(\log\log n)$.

Note that this heuristic is not a formal proof, since it assumes that the β_i bounds always hold (while in reality there must be some small failure probability) and since we bounded the expectation of β_{i+1} but we will want a bound that holds with high probability. However, the full proof is conceptually essentially the same.

2.2 Rigorous Proof

Theorem 4. *Suppose that n balls are distributed to n bins, according to the "power of two choices" method. Then in the end, with high probability the most loaded bin contains at most $O(\log\log n)$ balls.*

If $d \geq 2$ choices are used instead, then the load is at most $O\left(\frac{\log\log n}{\log d}\right)$ balls.

Proof: We shall give the proof for general d. First we define a few notations.

- Let t be the time right after the t-th ball is placed.
- Define $h(t)$, the height of the t-th ball be the number of balls in the same bin as the tth ball immediately after it is placed (including the tth ball itself).
- Define $\beta_6 = \frac{n}{2e}$ and $\beta_{i+1} = \frac{e\beta_i^d}{n^{d-1}}$ for $i \geq 6$. Note that we can bound $\beta_{i+1} \leq cn(\frac{\beta_i}{n})^d$ for some constant c, so with $j = O(\log\log n)$ we have $\beta_j < 1$.
- Let $B(n,p)$ be a Bernoulli random variable denoting the total number of heads resulting from flipping n coins, each flip with probability of heads p.

- Let $\nu_i(t)$ be the number of bins with load of at least i, and $\mu_i(t)$ be the number of balls of height at least i. Obviously $\nu_i(t) \leq \mu_i(t)$.
- Let $\mathcal{E}_i$ be the event that $\nu_i(n) \leq \beta_i$. Since $\beta_6 = \frac{n}{2e} < \frac{n}{6}$, the event $\mathcal{E}_6$ holds with certainty.

Fix an integer i. Let Y_i be a random binary variable such that

$$Y_t = 1 \text{ iff } h(t) \geq i+1 \text{ and } \nu_i(t-1) \leq \beta_i$$

That is, $Y_t = 1$ iff the height of the t-th ball is at least $i+1$, and at time $t-1$ there are at most β_i bins of load at least i.

Let ω_j be the bin selected by the j-th ball. Then

$$\mathbb{P}[Y_t = 1 | \omega_1, .., \omega_{t-1}] \leq \frac{\beta_i^d}{n^d} := p_i$$

The proof will also use the concept of (first-order) *stochastic dominance.* A random variable x is said to *stochastically dominate* another random variable y if for all k,

$$\mathbb{P}(x \geq k) \geq \mathbb{P}(y \geq k)$$

We use the following theorem on stochastic dominance:

Theorem 5. *Stochastic Dominance: If $X_1, X_2, \dots X_n$ are i.i.d. random variables $Y_i = f_i(X_1, X_2, \dots X_i)$, and $\mathbb{P}(Y_i = 1 | X_1, X_2, \dots X_{i-1}) \leq p$, then $\sum_i Y_i$ is stochastically dominated by $B(n,p)$.*

Using the Stochastic Dominance theorem

$$\mathbb{P}[\sum_{t=1}^{n} Y_t \geq k] \leq \mathbb{P}[B(n, p_i) \geq k]$$

With respect to $\mathcal{E}_i$ it follows $\mu_{i+1}(n) = \sum_{t=1}^{n} Y_t$. Thus

$$\mathbb{P}[\nu_{i+1} \geq k | \mathcal{E}_i] \leq \mathbb{P}[\mu_{i+1} \geq k | \mathcal{E}_i] = \mathbb{P}[\sum_{t=1}^{n} Y_t \geq k | \mathcal{E}_i] \leq \frac{\mathbb{P}[\sum_{t=1}^{n} Y_t \geq k]}{\mathbb{P}[\mathcal{E}_i]} \leq \frac{\mathbb{P}[B(n, p_i) \geq k]}{\mathbb{P}[\mathcal{E}_i]}$$

By the Chernoff bound, with $k = \beta_{i+1}$

$$\mathbb{P}[\nu_{i+1} \geq \beta_{i+1} | \mathcal{E}_i] \leq \frac{\mathbb{P}[B(n, p_i) \geq enp_i]}{\mathbb{P}[\mathcal{E}_i]} \leq \frac{1}{e^{p_i n} \, \mathbb{P}[\mathcal{E}_i]}$$

If $p_i n \geq 2 \log n$, we can rewrite the above expression as $\mathbb{P}[\neg \mathcal{E}_{i+1} | \mathcal{E}_i] \leq \frac{1}{n^2 \, \mathbb{P}[\mathcal{E}_i]}$. That implies with high probability $\nu_i(n) \leq \beta_i$ for large enough i.

However, if $p_i n \leq 2 \log n$, we show with high probability that there is no ball at height $i+2$. Let i^* be the smallest i such that $\frac{\beta_i^d}{n^d} \leq 2 \log n$. Since $\beta_{i+6} \leq \frac{n}{2^{d^i}}$, by induction $i^* \leq \frac{\log \log n}{\log d} + O(1)$.

We have

$$\mathbb{P}[\nu_{i^*+1}(n) \geq 6 \log n | \mathcal{E}_{i^*}] \leq \frac{\mathbb{P}[B(n, 2 \log n / n) \geq 6 \log n]}{\mathbb{P}[\mathcal{E}_{i^*}]} \leq \frac{1}{n^2 \, \mathbb{P}[\mathcal{E}_{i^*}]}$$

also

$$\mathbb{P}[\mu_{i^*+2} \geq 1 | \mu_{i^*+1} \leq 6 \log n] \leq \frac{\mathbb{P}[B(n, (6 \log n/n)^d) \geq 1]}{\mathbb{P}[\mu_{i^*+1} \leq 6 \log n]} \leq \frac{n(6 \log n/n)^d}{\mathbb{P}[\mu_{i^*+1} \leq 6 \log n]}$$

Apply the inequality

$$\mathbb{P}[\neg \mathcal{E}_{i+1}] \leq \mathbb{P}[\neg \mathcal{E}_{i+1} | \mathcal{E}_i] \, \mathbb{P}[\mathcal{E}_i] + \mathbb{P}[\neg \mathcal{E}_i]$$

we obtain

$$\mathbb{P}[\mu_{i^*+2} \geq 1] \leq \frac{(6 \log n)^d}{n^{d-1}} + \frac{i^* + 1}{n^2} = O\left(\frac{1}{n}\right)$$

That is, with high probability there's no ball at height $i^* + 2 = O\left(\frac{\log \log n}{\log d}\right)$. That concludes the proof.