Lecture 7: January 28, 2003

Lecturer: Eric Vigoda *Scribe: Tom Hayes and Özgür Sümer*

Balls and Bins

Basic Scheme

Suppose we place n balls randomly into n bins by putting each ball into a randomly chosen bin. Its not hard to prove that the maximum load of the bin ensemble is $(1 + o(1) \ln n/ \ln \ln n$. More precisely, we have the following lemma.

Lemma 7.1 *With high probability $(1 - n^{-c})$, the fullest bin has $(1 + o(1))\frac{\ln n}{\ln \ln n}$ balls in it.*

Proof: If we don't care about the constant the proof is easy.

$$\begin{aligned} \Pr(\text{ bin } j \text{ has more than } k \text{ balls }) &\leq \binom{n}{k} n^{-k} \\ &\leq \left(\frac{en}{k}\right)^k \leq n^{-2}. \end{aligned}$$

Solving this for k, we get $k \geq \frac{2e \ln n}{\ln \ln n}$. ■

This is a natural problem with applications such as load balancing, resource allocation, and hashing. Hence it is desirable to have a "scheme" which minimizes the maximum load. We have to define what we mean by a scheme since otherwise we can achieve a max load of 1 deterministically. The strategy for each ball should be identical and should not depend on the choices of the other ball. In addition, we assume that the balls are numbered 1 through n and the i^{th} ball makes its choice at time i.

Power of multiple choices

Now, consider the following scheme. Ball i, chooses d bins uniformly at random, and places itself into the least loaded bin of these. This apparent minor change results in a maximum load of $\frac{\ln \ln n}{\ln d} + \Theta(1)$ with high probability.

Theorem 7.2 *[ABKU99] Suppose that m balls are sequentially placed into n boxes. Each ball is placed in the least full box, at the time of the placement, among d boxes, $d \geq 2$, chosen independently and uniformly at random (where ties are broken arbitrarily). Then after all the balls are placed, with high probability, as $n \to \infty$ and $m \geq n$, the number of balls in the fullest box is $(1 + o(1))\frac{\ln \ln n}{\ln d} + \Theta(\frac{m}{n})$.*

For clarity we prove the theorem for $m = n$. Before we get on to the proof we introduce some notation and provides some intuition. We need the following notation.

Notation 7.3

$$\begin{aligned} {N^{\text{bins}}}_k(t) &:= \text{\# bins with load} \geq k \text{ at time } t \\ {N^{\text{balls}}}_k(t) &:= \text{\# balls with height} \geq k \text{ at time } t \\ B(n,p) &:= \text{Binomial distribution with } n \text{ Bernoulli trials with probability } p \end{aligned}$$

Intuition: Clearly, ${N^{\text{bins}}}_k(t) \leq {N^{\text{balls}}}_k(t)$. The following is also obvious:

$${N^{\text{bins}}}_k(1) \leq {N^{\text{bins}}}_k(2) \leq \cdots \leq {N^{\text{bins}}}_k(n)$$

Suppose for some B_k we have ${N^{\text{bins}}}_k(n) \leq B_k$. Then we have the following bound

$$\Pr\left(\text{ ball } i \text{ has height} \geq k+1 \;\middle|\; {N^{\text{bins}}}_k(i) \leq B_k \right) \leq \left(\frac{B_k}{n}\right)^d.$$

So,

$$\Pr\left({N^{\text{balls}}}_{k+1}(n) \geq j \;\middle|\; {N^{\text{bins}}}_k(i) \leq B_k \right) \leq \Pr\left(B(n, \left(\frac{B_k}{n}\right)^d) \geq j \right)$$

The following lemma gives a bound on the number of successes of n Bernoulli trials. Roughly speaking, if the success probability of a trial is less than p, the distribution of total number of successes will be bounded above by $B(n,p)$.

Lemma 7.4 *(Basic Lemma) Let $X_1, \ldots, X_n$ be arbitrary random variables. Let $Y_1, \ldots, Y_n$ be $0-1$ random variables, where $Y_i = Y_i(X_1, \ldots, X_i)$ (note given $X_1, \ldots, X_i$ then Y_i is a deterministic function). If* $\mathrm{E}\left(Y_i \mid X_1, \ldots, X_{i-1} \right) \leq p$, *then* $\Pr\left(\sum_i Y_i \geq j \right) \leq \Pr\left(B(n,p) \geq j \right)$. *Similarly, If* $\mathrm{E}\left(Y_i \mid X_1, \ldots, X_{i-1} \right) \geq p$, *then* $\Pr\left(\sum_i Y_i \geq j \right) \leq \Pr\left(B(n,p) \leq j \right)$.

Remark: The stronger version which only assumes $\mathrm{E}\left(Y_i \mid Y_1, \ldots, Y_{i-1} \right) \leq p$ (and doesn't need to mention X) is also true. The proof is an easy induction.

Proof:(Main Theorem)

Now, let $B_6 = n/2e$. Let $B_{i+1} = \frac{eB_i^d}{n^{d-1}}$ for $6 \leq i \leq i^*$. $B_{i+6} = \frac{ne^{(d^i-1)/(d-1)}}{(2e)^{d^i}} \leq \frac{n}{2^{d^i}}$ for $0 \leq i \leq i_* - 6$ Define the event E_i by $\{{N^{\text{bins}}}_i(n) \leq B_i\}$. Observe that E_6 holds with probability one since at most $n/6$ bins can have load at least 6 and $B_6 > n/6$.

Our plan is now:

1. for $6 \leq i \leq i^*$, we show $\Pr\left(\overline{E_{i+1}} \mid E_i \right) \leq n^{-2}$.
2. And $\Pr\left({N^{\text{bins}}}_{i^*+2} \geq 1 \mid E_{i^*} \right) \leq n^{-2}$.

Chaining these together proves the result.

Fix $i \geq 1$. Consider a series of $0-1$ indicator variables Y_t for $2 \leq t \leq n$ where Y_t indicates the event that ball t has height $\geq i+1$ and that ${N^{\text{bins}}}_i(t-1) \leq B_i$. Note that Y_t depends on i. Intuitively Y_t indicates that the t^{th} ball was placed in a bad bin even though there are enough good bins. Being bad or good for a bin means having load more than i or less than i,respectively and enough means more than $n - B_i$. So,

$$\Pr\left(Y_t = 1 \mid \text{any placement of the first } t-1 \text{ balls} \right) \leq \left(\frac{B_i}{n}\right)^d.$$

Let us denote $p_i := \left(\frac{B_i}{n}\right)^d$. By the basic lemma we have

$$\Pr\left(\sum_t Y_t \geq k\right) \leq \Pr(\ B(n, p_i) \geq k\ )$$

Observe that if E_i holds, then ${N^{\text{balls}}}_{i+1} = \sum_{t\in\{2,\ldots,n\}} Y_t$ because the event indicated by Y_t is reduced to the event that ball t has height at least $i+1$

$$\sum_t Y_t = {N^{\text{balls}}}_{i+1}(n) \geq {N^{\text{bins}}}_{i+1}(n)$$

$$\begin{aligned}
\Pr(\ \overline{E_{i+1}} \mid E_i\ ) &= \Pr(\ {N^{\text{bins}}}_{i+1}(n) > k \mid E_i\ ) \text{ where } k = B_{i+1} = ep_i n. \\
&\leq \Pr\left(\sum Y_t > k \mid E_i\right) \\
&\leq \frac{\Pr(\ \sum Y_t > k\ )}{\Pr(\ E_i\ )} \\
&\leq \frac{\Pr(\ B(n,p_i) > k\ )}{\Pr(\ E_i\ )} \\
&\leq \frac{1}{e^{p_i n}\Pr(\ E_i\ )}.
\end{aligned}$$

Last inequality was obtained by Chernoff because $\mathrm{E}(\ B(n,p)\ ) = np$.

$$\begin{aligned}
\Pr(\ \overline{E_{i+1}}\ ) &\leq \Pr(\ \overline{E_{i+1}} \mid E_i\ )\Pr(\ E_i\ ) + \Pr(\ \overline{E_i}\ ) \\
&\leq \frac{1}{e^{p_i n}} + \Pr(\ \overline{E_i}\ ) \\
&\leq \frac{1}{n^2} + \Pr(\ \overline{E_i}\ ) \text{ for } p_i \geq \frac{2\log n}{n}
\end{aligned}$$

Since $p_i = \left(\frac{B_i}{n}\right)^d$ and $\{B_i, i \in N\}$ is a decreasing sequence, p_is are also decreasing. let i^* be the first time it becomes less than $\frac{2\log n}{n}$. so $p_{i^*} = \left(\frac{B_{i*}}{n}\right)^d \leq 2\ln(n)/n$. Up to i^*, $\Pr(\ \overline{E_i}\ )$ increases not more than $\frac{1}{n^2}$. So we have

$$\Pr(\ \overline{E_{i^*}}\ ) \leq \frac{i^*}{n^2}$$

Notice that because of $B_{i+6} \leq \frac{n}{2^{i^*}}$, $i^* \leq \frac{\log\log n}{\log d} + O(1)$

$$\Pr(\ {N^{\text{bins}}}_{i^*+1}(n) \geq 6\log n \mid E_{i^*}\ ) \leq \frac{\Pr(\ B(n, 2\log n/n) \geq 6\log n\ )}{\Pr(\ E_{i^*}\ )} \leq \frac{1}{n^2\Pr(\ E_{i^*}\ )}$$

If we drop the condition we get

$$\Pr(\ {N^{\text{bins}}}_{i^*+1} \geq 6\log n\ ) \leq \frac{1}{n^2} + \Pr(\ \overline{E_{i^*}}\ )$$

Thus,

$$\Pr(\ {N^{\text{balls}}}_{i^*+2} \geq 1 \mid {N^{\text{bins}}}_{i^*+1} \leq 6\log n\ ) \leq \frac{\Pr\left(\ B(n, (6\log n/n)^d) \geq 1\ \right)}{\Pr(\ {N^{\text{bins}}}_{i^*+1 \leq 6\log n}\ )} \leq \frac{n(6\log n/n)^d}{\Pr(\ {N^{\text{bins}}}_{i^*+1 \leq 6\log n}\ )}$$

Markov inequality was used at last inequality.

$$\begin{aligned}\Pr\left(\,{N^{\text{balls}}}_{i^*+2} \geq 1\,\right) &\leq \frac{(6\log n)^d}{n^{d-1}} + \Pr\left(\,{N^{\text{bins}}}_{i^*+1} \geq 6\log n\,\right) \\ &\leq \frac{(6\log n)^d}{n^{d-1}} + \frac{i^*+1}{n^2} \\ &= o(1).\end{aligned}$$

as desired.

■

Further schemes

An improved scheme was recently presented by Vöcking [V99]. Initially partition the bins into d groups of equal size. Assign an arbitrary ordering on the groups. When considering ball j, choose exactly one random bin from each group. Place the ball into the least loaded bin. If there's a tie then choose the lowest group according to our pre-determined ordering. Note there is now an asymmetry between the bins. This asymmetry results in a maximum load of roughly $\log\log(n)/d$, which improves the denominator.

References

[ABKU99] Azar, Y. Broder, A. Z. Karlin A. R. Upfal E. *Balanced Allocations*, Siam J. Comput. (1999) Vol. 29, No. 1: 180-200.

[V99] Vöcking, Berthold. *How asymmetry helps load balancing*, 40th Annual Symposium on Foundations of Computer Science, New York, 1999, 131–141.