CS271 Randomness & Computation **Fall 2022**

Lecture 15: October 13

Instructor: Alistair Sinclair

Disclaimer: *These notes have not been subjected to the usual scrutiny accorded to formal publications. They may be distributed outside this class only with the permission of the Instructor.*

15.1 Balls and Bins

In the standard balls and bins model, we throw m balls randomly into n bins (i.e., each ball picks its bin independently and u.a.r.), and analyse the distribution of the loads (number of balls) in the bins. For instance, a typical question of interest is: What is the maximum load in any bin? Such questions have direct relevance to practical systems that assign jobs to machines or requests to web servers, because we can abstract these systems using the balls and bins model.

In our analysis, we assume m, n to be comparable i.e., $m = cn$ for a constant c. Though most of our analysis focuses on $c = 1$, it is easily extendable to $c > 1$. The load X_i in a fixed bin i is easy to analyse: X_i is distributed as $Binomial(m, 1/n)$ and hence is asymptotically Poisson with parameter $\lambda = \frac{m}{n}$ (denoted $Poisson(\lambda)$). [**Exercise:** Show $\Pr[X_i = k] \to \frac{e^{-\lambda}\lambda^k}{k!}$ as $n \to \infty$, for each fixed k.]

The joint distribution of all loads $\{X_i\}_{i=1}^n$ is multinomial, which is much harder to deal with. To get around the dependences between different bin loads, it is tempting to relate this joint distribution to that of n *independent* $Poisson(\lambda)$ loads $\{Y_i\}_{i=1}^n$. If we let $m = \sum_{i=1}^n k_i$, then

$$\Pr[X_1 = k_1, \ldots, X_n = k_n] = \frac{1}{n^m} \frac{m!}{k_1! \ldots k_n!} = \frac{\prod_{i=1}^n \frac{e^{-\lambda}\lambda^{k_i}}{k_i!}}{\frac{e^{-n\lambda}(n\lambda)^m}{m!}} = \frac{\prod_{i=1}^n \Pr[Y_i = k_i]}{\Pr[\sum Y_i = m]} \ . \tag{15.1}$$

In the denominator here, note that $\sum Y_i$ is the sum of independent Poisson r.v.'s and so is itself Poisson. Thus the real loads $\{X_i\}$ indeed behave like independent $Poisson(\lambda)$ loads $\{Y_i\}$ **but** conditioned on $\sum Y_i = m$. As the following useful theorem from [Mi96] shows, we can remove this conditioning and work more conveniently with independent Y_i's if we only need to bound the probability of rare events.

Theorem 15.1 *Let $\mathcal{E}$ be any event that depends only on the bin loads s.t. $\Pr[\mathcal{E}]$ is monotonically increasing (or decreasing) with m. Then*

$$\Pr_X[\mathcal{E}] \le 4 \Pr_Y[\mathcal{E}],$$

where $\Pr_X$ is the probability in the standard balls and bins model, and $\Pr_Y$ is the probability in the independent Poisson model.

Proof: Assume $\Pr[\mathcal{E}]$ is increasing (the decreasing case is entirely analogous). Then,

$$\begin{aligned}
\Pr_Y[\mathcal{E}] &= \sum_{k=0}^{\infty} \Pr_Y\left[\mathcal{E} \mid \sum Y_i = k\right] \Pr_Y\left[\sum Y_i = k\right] \\
&\geq \Pr_Y\left[\mathcal{E} \mid \sum Y_i = m\right] \; \Pr_Y\left[\sum Y_i \geq m\right] \\
&\geq \Pr_Y\left[\mathcal{E} \mid \sum Y_i = m\right] \times \frac{1}{4} \\
&= \Pr_X[\mathcal{E}] \times \frac{1}{4} \; .
\end{aligned}$$

In the second line here we used monotonicity, and in the third line we used the fact [**Exercise!**] that for any Poisson r.v. Y we have $\Pr[Y \geq E(Y)] \geq \frac{1}{4}$, applied to the r.v. $Y = \sum_i Y_i$ (which is the sum of independent Poissons and hence Poisson). ■

We now return to our main question about the maximum load in any bin. Working with the independent Poisson model instead of the standard model leads to a simple proof of the following fundamental theorem.

Theorem 15.2 *For the balls and bins model with $m = n$, the maximum load in any bin is $\Theta\left(\frac{\ln n}{\ln \ln n}\right)$ w.h.p.*

Proof: Actually we'll prove the stronger statement that the maximum load lies in the range $(1 \pm \varepsilon)\left(\frac{\ln n}{\ln \ln n}\right)$ w.h.p., for any constant $\varepsilon > 0$. Define $c_1 = 1 + \epsilon$ and $c_2 = 1 - \epsilon$, and consider two events:

$$\begin{aligned}
\mathcal{E}_1 &\equiv \text{“some bin contains more than } \tfrac{c_1 \ln n}{\ln \ln n} \text{ balls”}, \\
\mathcal{E}_2 &\equiv \text{“no bin contains more than } \tfrac{c_2 \ln n}{\ln \ln n} \text{ balls”}.
\end{aligned}$$

Our goal is to prove that $\Pr[\mathcal{E}_1]$, $\Pr[\mathcal{E}_2]$ are both $o(1)$. Moreover, we can work with the independent Poisson model as both $\mathcal{E}_1, \mathcal{E}_2$ satisfy the conditions of Theorem 15.1. Since the Y_i's are identically distributed, we focus on bin 1 and look at $p_k \equiv \Pr_Y[Y_1 \geq k]$. Since $Y_1 \sim \mathit{Poisson}(1)$ we have $p_k = \sum_{j=k}^{\infty} \frac{e^{-1}}{j!}$. We will need these simple bounds on p_k:

$$\frac{1}{ek!} \;\leq\; p_k \;\leq \frac{1}{ek!}\left(1 + \frac{1}{k+1} + \frac{1}{(k+1)(k+2)} + \dots\right) \;\leq\; \frac{1}{k!} \; .$$

We first show $\Pr_Y[\mathcal{E}_1] = o(1)$ via a union bound on $\Pr_Y[Y_i \geq k]$. Take $k = \frac{c_1 \ln n}{\ln \ln n}$ with $c_1 = 1 + \epsilon$. Then $\ln p_k \leq -\ln k! \sim -k \ln k = -\frac{c_1 \ln n}{\ln \ln n}(\ln \ln n + \ln c_1 - \ln \ln \ln n) \sim -c_1 \ln n$, which due to the value of c_1 implies $p_k = o(n^{-1})$. Taking a union bound over all n bins gives $\Pr_Y[\mathcal{E}_1] = o(1)$, as desired.

We next show $\Pr_Y[\mathcal{E}_2] = o(1)$ using the independence of the Y_i's, which we may do by virtue of Theorem 15.1. Take $k = \frac{c_2 \ln n}{\ln \ln n}$; then

$$\Pr_Y[\mathcal{E}_2] = (1 - p_k)^n \leq \left(1 - \frac{1}{ek!}\right)^n \leq e^{-\frac{n}{ek!}}.$$

We need to show that the exponent $\frac{n}{ek!} \to \infty$ as $n \to \infty$. Taking logs, this is equivalent to $\ln n - \ln(k!) \to \infty$. But by the same derivation as before we see that $\ln(k!) \sim c_2 \ln n$, so $\ln n - \ln(k!) \sim \varepsilon \ln n \to \infty$, as desired. ■

Exercise: Tighten up the above argument slightly to show that the maximum load is actually $(1 + o(1))\left(\frac{\ln n}{\ln \ln n}\right)$ w.h.p.

15.2 The Power of Two Choices

Suppose now that instead of placing each ball in a single random bin, we choose d random bins for each ball, place it in the one that currently has fewest balls, and proceed in this manner sequentially for each ball.

Clearly if $d = n$ the maximum load will be optimal: $\lceil m/n \rceil$. But this requires global knowledge of the state of all bins. We will show that with $d = 2$ choices, the maximum load is a.a.s. at most

$$\frac{\ln \ln n}{\ln 2} + \Theta(1)$$

That is, having only one more choice than the basic model reduces the maximum load exponentially. As a concrete example, if we take $m = n = 10^6$, we find experimentally that the maximum bin load in the basic model is almost always in the range $8 - 10$; however, with two choices it drops dramatically to 4. [**Exercise:** Write some simple code to check these numbers, and play around with other values of m, n.]

This result has been discovered multiple times; see [Mi96] for more background and various extensions. The proof technique we use here is due to [ABKU94].

15.2.1 Proof idea

Let B_i be the number of bins with load $\geq i$ at the end of the process. Suppose we could find upper bounds β_i so that $B_i \leq \beta_i$ w.h.p. Then

$$\Pr[\text{given ball is placed in a bin with load} \geq i] \quad \leq \quad \left(\frac{\beta_i}{n}\right)^2,$$

because both of the ball's choices must land in a bin with load $\geq i$. This gives us a crude upper bound on the number of bins with load $\geq i + 1$: the distribution of B_{i+1} is dominated by the binomial distribution $\text{Bin}\left(n, (\beta_i/n)^2\right)$. The mean of this distribution is β_i^2/n, so by a Chernoff bound we can take $\beta_{i+1} = c\beta_i^2/n$ for some constant c, and we have

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

Thus $\frac{\beta_i}{n}$ decreases quadratically, so for $i \approx \frac{\ln \ln n}{\ln 2}$ we will have $\beta_i < 1$, which implies that the maximum load is $\frac{\ln \ln n}{\ln 2}$ w.h.p.

15.2.2 Full proof

For algebraic convenience, set $\beta_6 = \frac{n}{2e}$. Note that $B_6 \leq \beta_6$ is trivial since there can be at most $\frac{n}{6} < \frac{n}{2e}$ bins with ≥ 6 balls in them. For $i > 6$, let

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

Now define the event $\mathcal{E}_i = \{B_i \leq \beta_i\}$ (so $\Pr[\mathcal{E}_6] = 1$). We have

$$\Pr[\neg\mathcal{E}_{i+1} | \mathcal{E}_i] \quad = \quad \Pr[B_{i+1} > \beta_{i+1} | \mathcal{E}_i] \quad \leq \quad \frac{\Pr[\text{Bin}(n, (\beta_i/n)^2) \;\geq\; \beta_{i+1}]}{\Pr[\mathcal{E}_i]}.$$

Note that the denominator is necessary here: we cannot claim that the numerator bounds the conditional probability, because once we condition on $\mathcal{E}_i$ the bin choices are no longer independent.

We now apply a Chernoff bound of the form $\Pr[X \geq e\mu] \leq e^{-\mu}$, which follows directly from a more general form discussed in Lecture 13. (Specifically, plug $\beta = e - 1$ into bound $(*)$ of Corollary 13.3.) Thus,

$$\Pr[\neg\mathcal{E}_{i+1} | \mathcal{E}_i] \quad \leq \quad \frac{e^{-\beta_i^2/n}}{\Pr[\mathcal{E}_i]} \quad \leq \quad \frac{1/n^2}{\Pr[\mathcal{E}_i]},$$

provided $\beta_i^2/n \geq 2\ln n$. Now to remove the conditioning we prove by induction on i that $\Pr[\neg\mathcal{E}_i] \leq \frac{i}{n^2}$. In the base case, $\Pr[\neg\mathcal{E}_6] = 0$. For the inductive step,

$$\Pr[\neg\mathcal{E}_{i+1}] \leq \Pr[\neg\mathcal{E}_{i+1}|\mathcal{E}_i]\Pr[\mathcal{E}_i] + \Pr[\neg\mathcal{E}_i] \leq \frac{1/n^2}{\Pr[\mathcal{E}_i]} \cdot \Pr[\mathcal{E}_i] + \frac{i}{n^2} \leq \frac{i+1}{n^2}.$$

So $\Pr[\neg\mathcal{E}_i] \leq i/n^2 \leq 1/n$ for all i such that $\beta_{i-1}^2/n \geq 2\ln n$.

Now let i^* be the minimum i for which $\beta_i^2 < 2n\ln n$. Then $i^* = \frac{\ln\ln n}{\ln 2} + O(1)$. [**Exercise:** prove this.] To get a feel for what's going on at this point, note that w.h.p. there are $\leq \sqrt{2n\ln n}$ bins with load $\geq i^*$, so the expected number of balls falling in bins with load $\geq i^*+1$ is at most $2\ln n$. The following claim finishes the proof:

Claim 15.3 $\Pr[B_{i^*+2} \geq 1] \leq O\left(\frac{\log^2 n}{n}\right)$.

Proof: Define $\mathcal{E}_{i^*+1} = \{B_{i^*+1} \leq 6\ln n\}$. We have

$$\begin{aligned}
\Pr[\neg\mathcal{E}_{i^*+1}] &\leq \Pr[B_{i^*+1} \geq 6\ln n|\mathcal{E}_{i^*}] \cdot \Pr[\mathcal{E}_{i^*}] + \Pr[\neg\mathcal{E}_{i^*}] \\
&\leq \frac{\Pr\left[\mathrm{Bin}(n, 2\ln n/n) \geq 6\ln n\right]}{\Pr[\mathcal{E}_{i^*}]} \cdot \Pr[\mathcal{E}_{i^*}] + \frac{1}{n} \\
&\leq \frac{1}{n^2} + \frac{1}{n} = O\left(\frac{1}{n}\right),
\end{aligned}$$

where the bound on $\Pr[\neg\mathcal{E}_{i^*}]$ comes from our previous calculation (using the fact that $\beta_{i^*-1}^2/n \geq 2n\ln n$), and we have again used a Chernoff bound on the Binomial distribution. Finally,

$$\begin{aligned}
\Pr[B_{i^*+2} \geq 1] &\leq \Pr[B_{i^*+2} \geq 1 \,|\, \mathcal{E}_{i^*+1}] \cdot \Pr[\mathcal{E}_{i^*+1}] + \Pr[\neg\mathcal{E}_{i^*+1}] \\
&\leq \frac{\Pr\left[\mathrm{Bin}\left(n, (6\ln n/n)^2\right) \geq 1\right]}{\Pr[\mathcal{E}_{i^*+1}]} \cdot \Pr[\mathcal{E}_{i^*+1}] + O\left(\frac{1}{n}\right) \\
&\leq \left(\frac{6\ln n}{n}\right)^2 \cdot n + O\left(\frac{1}{n}\right) = O\left(\frac{(\ln n)^2}{n}\right),
\end{aligned}$$

where we have used a simple union bound on the probability that the Binomial is nonzero. ■

Exercise: Prove that the maximum load is $\Omega(\ln\ln n)$ w.h.p.

Exercise: Extend the above analysis to show that if each ball has d choices, the maximum load is $\frac{\ln\ln n}{\ln d} + O(1)$ w.h.p. Thus, additional choices beyond 2 affect only the constant.

References

[ABKU94] Y. Azar, A. Broder, A. Karlin and E. Upfal, "Balanced allocations," in *Proceedings of the 26th ACM Symposium on Theory of Computing*, 1994, pp. 593–602.

[Mi96] M. Mitzenmacher, *The Power of Two Choices in Randomized Load Balancing*, PhD Thesis, UC Berkeley, 1996.