10 | 10 | 10 | 6 | 8

6.856 **Dan Ports**

Randomized Algorithms drkp@mit.edu

2005/02/23 Collaborators: {sarahl,gpickard}

Problem Set 3

Problem 1:

a) Let X_i be independent geometrically distributed random variables with mean 2. Thus X_i is the number of flips of an unbiased coin until the first heads. We can view these sequences of coin flips as being performed in sequence, and since the n geometric random variables are independent and each performs coin flips until it reaches a heads, we can consider this to be a sequence of flips of a single coin, and X_i to be the number of coin flips between the ith and $i-1$th head. Thus $X = \sum_{i=1}^{n} X_i$ is the number of coin flips performed before the nth head. We can express this as sum of Bernoulli random variables Y_i, where Y_i is simply the outcome from the ith single coin flip. The probability that X is greater than some x is the probability that more than x flips are required to find n heads, which is the probability that $\sum_{i=1}^{x} Y_i$ is less than n. In particular, we want to consider $x = (1+\epsilon)\mu_X = (1+\epsilon)2n$. Write $Y = \sum_{i=1}^{(1+\epsilon)2n} Y_i$. Then $\mathrm{E}[Y] = (1+\epsilon)n$, and

$$\begin{aligned}
\Pr[X > (1+\epsilon)2n] &= \Pr[Y < n] \\
&= \Pr\left[Y < \frac{n}{1+\epsilon}\mathrm{E}[Y]\right] \\
&= \Pr\left[Y < \left(1 + \frac{-\epsilon}{1+\epsilon}\right)\mathrm{E}[Y]\right] \\
&\leq e^{-\frac{\epsilon^2(1+\epsilon)n}{2(1+\epsilon)^2}} = e^{-\frac{\epsilon^2 n}{2(1+\epsilon)}}
\end{aligned}$$

since Y is a sum of Bernoulli random variables and hence the Chernoff bound applies.

b) As in the Chernoff bound analysis, we note that the desired probability is equivalent to $\Pr\left[e^{tX} > e^{t(1+\epsilon)2n}\right]$. By the Markov bound, this is less than $\frac{\mathrm{E}[e^{tX}]}{e^{t(1+\epsilon)2n}}$. Because X is the sum of independent X_is, the expectation is $\prod_i \mathrm{E}\left[e^{tX_i}\right]$. Now note that

$$\begin{aligned}
\mathrm{E}\left[e^{tX_i}\right] &= \sum_{j=1}^{\infty} e^{tj}\frac{1}{2}^j = \sum_{j=1}^{\infty}\left(\frac{e^t}{2}\right)^j \\
&= \frac{\frac{e^t}{2}}{1-\frac{e^t}{2}} = \frac{2e^t}{2-e^t}
\end{aligned}$$

assuming $t < \ln 2$ so that the series converges. So

$$\Pr[X > (1+\epsilon)2n] \leq \frac{\left(\frac{2e^t}{2-e^t}\right)^n}{e^{t(1+\epsilon)2n}} = \frac{1}{((2-e^t)e^{t+2t\epsilon})^n}$$

The TI-89 does not flinch in fear at this equation (like I do) and tells us that the derivative is equal to zero at $t = \ln \frac{2\epsilon+1}{\epsilon+1}$. Thus this is the minimum. So

$$\begin{aligned}
\Pr\left[X > (1+\epsilon)2n\right] &\leq \frac{1}{\left(\left(2 - \frac{2\epsilon+1}{\epsilon+1}\right)\left(\frac{2\epsilon+1}{\epsilon+1}\right)^{(1+2\epsilon)}\right)^{n}} \\
&= \frac{(\epsilon+1)^n}{\left(\frac{2\epsilon+1}{\epsilon+1}\right)^{(1+2\epsilon)n}} \\
&= \frac{(\epsilon+1)^{2n\epsilon}}{(2\epsilon+1)^{(1+2\epsilon)n}}
\end{aligned}$$

By choosing ϵ sufficiently large, we can make this quantity arbitrarily small.

Problem 2:

Define a pivoting round to be "good" for a given item x if x is in a subproblem with size at most $9/10$ of the problem from which it was partitioned. So each "good" round reduces the size of the problem by $9/10$. Thus for any x, x will not be involved in more than $\log_{10/9} n$ good rounds, because this would reduce the size of a subproblem below 1.

For any subproblem of size m and any element x, the partition is a good pivoting round for x with probability at least $\frac{1}{2}$, since the pivot is chosen uniformly at random. Suppose x has rank r among the m elements. Assume without loss of generality that $r \leq \frac{m}{2}$ (if not, simply reverse the order of the ranking). Let p be the rank of the pivot. If $r < p \leq 9/10$ or $1/10 \leq p < r$, then this is a good pivoting round. This occurs with probability at least $1/2$. Moreover, the probability of any given pivoting round being good is independently above $\frac{1}{2}$ since they depend only on the relative rankings of the elements in the subproblem.

Let $f(n) = k \log_{10/9} n$. If $f(n)$ partitions are performed, then the expected number of good pivots is $\frac{k}{2} \log_{10/9} n$. Moreover, by the independence, we can apply the Chernoff bound and find that the number of good pivots is less than $(1-\epsilon)\frac{k}{2}\log_{10/9}$ with probability $e^{\frac{-\epsilon^2 \frac{k}{2} \log_{10/9} n}{2}} = e^{\frac{-\epsilon^2 k \log_{10/9} n}{4}}$. So the probability of having more than $\frac{k}{4} \log_{10/9} n$ good pivots ($\epsilon = 1/2$) is at least $e^{\frac{-k \log_{10/9} n}{16}}$. If we choose $k = 16$, then the probability of having at least $4 \log_{10/9} n$ is at least $e^{-\log_{10/9} n} = \left(\frac{1}{n}\right)^{\frac{1}{\log \frac{10}{9}}} \approx \frac{1}{x^{9.49}}$.

Thus, with high probability, a sequence of $16 \log_{10/9} n = \Theta(\log n)$ pivots gives at least $4 \log_{10/9} n$ good pivots for any x, which we showed was enough to reduce the size of a subproblem below 1, and end the algorithm. So the algorithm has runtime $O(n \log n)$ with high probability.

Problem 3:

a) Consider the transpose permutation, mapping $a_i b_i$ to $b_i a_i$. We assume that the bit-fixing algorithm fixes bits from left to right. So if we fix some b, for all a_i, $a_i b$ will be converted to bb at some point in the process of being routed to ba_i. If the leftmost bit of a_i (call it a_1^0) differs from the corresponding bit of b (call it b^0 and the remainder $b^{1..\frac{n}{2}}$), then the $n/2 + 1$th bit will be fixed and it will be routed through the link from bb to $ba^0 b^{1..\frac{n}{2}}$. There are $2^{n/2}$ choices for a_i in $a_i b$, and exactly half of these will have bit a_i^0 differing from bit b^0. So $\frac{2^{n/2}}{2} = \Omega\left(\sqrt{N}\right)$ packets are routed through this edge, and so $\Omega\left(\sqrt{N}\right)$ routing steps are required.

b) Consider the transpose permutation again, mapping $a_i b_i$ to $b_i a_i$, and again fix some b. Now, however, the bits are fixed in random order, so it is not guaranteed that $a_i b$ will be converted to bb.

Suppose a_i differs from b at k bits. Then it will be converted to bb if all of the k bits in the a_i half are fixed before the k bits in the b half. The probability of this is the probability that the first k bits chosen to be fixed from the $2k$ bits that differ are the k in the first half, i.e. $\frac{1}{\binom{2k}{k}}$. The number of bit strings that end in b and begin with some a_i that has k bits differing from b is $\binom{n/2}{k}$ since we choose the k bits that differ from the $n/2$ bits in a_i. Thus, combining these and applying the inequality, the expected number of packets that route through bb is

$$\sum_{k=1}^{n/2} \frac{\binom{\frac{n}{2}}{k}}{\binom{2k}{k}} \geq \sum_{k=1}^{n/2} \frac{\left(\frac{\frac{n}{2}}{k}\right)^k}{\left(\frac{2ek}{k}\right)^k} = \sum_{k=1}^{n/2} \left(\frac{n}{4ek}\right)^k$$

The value of the sum is certainly bounded below by any of its terms, so consider the term with $k = \lceil \frac{n}{8e} \rceil$. This reduces to $2^{\lceil \frac{n}{8e} \rceil}$ (to a constant factor), which is $2^{\Omega(n)}$. Since this is the expected number of packets that pass through bb, the routing algorithm must require this many steps in expectation.

Since the expectation is $2^{\Omega(n)}$, and this can be represented as the sum of Bernoulli random variables that indicate whether the ith packet has its bits flipped in the order that forces it to go through bb, we can apply the Chernoff bound. This tells us that the probability that the number of steps required deviates from its mean is small, and gives us a bound of $2^{\Omega(n)}$ with high probability.

Problem 4:

a) Let X_k that bin 1 contains k balls. It is given by the binomial distribution:

$$\Pr[X_k] = \binom{n}{k}\left(\frac{1}{n}\right)^k\left(1-\frac{1}{n}\right)^{n-k}$$

which is bounded below by

$$\left(\frac{n-1}{n}\right) \geq \frac{1}{2} \quad \forall n > 1$$

$$\Pr[X_k] \geq \left(\frac{n}{k}\right)^k\left(\frac{1}{n}\right)^k\left(\frac{n-1}{n}\right)^n$$
$$\geq \frac{1}{k^k}\left(\frac{1}{2}\right)^n$$

WHERE DID THIS COME FROM?

We want this probability to be at least $\frac{1}{\sqrt{n}}$. This gives us the constraint:

$$\Pr\left[X^k\right] \geq \frac{1}{k^k}\left(\frac{1}{2}\right)^n \geq \frac{1}{\sqrt{n}}$$
$$k^k \leq \sqrt{n}$$
$$2k \log k \log n \leq \frac{1}{2} \log n$$
$$k \log k \leq \frac{3}{4} \log n$$

Setting $k = \alpha \frac{\log n}{\log \log n}$,

$$\alpha \frac{\log n}{\log \log n} \log \left(\alpha \frac{\log n}{\log \log n} \right) \leq \frac{3}{4} \log n$$

$$\alpha \frac{1}{\log \log n} (\log \alpha + \log \log n - \log \log \log n) \leq \frac{3}{4}$$

$$\frac{\alpha \log \alpha}{\log \log n} + \alpha \leq \frac{3}{4}$$

For large n (greater than e^e), $\log \log n$ is greater than 1, so ignoring the denominator only strengthens the constraint. Thus we simply need to choose α small enough that $\alpha \log \alpha + \alpha \leq 3/4$. Note that k is still $\Omega\left(\frac{\log n}{\log \log n}\right)$. We have shown that the probability that bin 1 contains exactly k balls is at least $\frac{1}{\sqrt{n}}$, so this is certainly a bound on the probability of it having at least k balls.

The probability that bin 1 contains at least k balls is at least $\frac{1}{\sqrt{n}}$, and for each i, the probability that bin i contains at least k balls given that bins 1 through $i-1$ did not is the same. So the total probability is bounded below by the finite geometric series

$$\sum_{i=0}^{n-1} \frac{1}{\sqrt{n}} \left(1 - \frac{1}{\sqrt{n}}\right)^i = \frac{1}{\sqrt{n}} \frac{1 - \left(1 - \frac{1}{\sqrt{n}}\right)^n}{\frac{1}{\sqrt{n}}} = \frac{1}{n} = 1 - \left(1 - \frac{1}{\sqrt{n}}\right)^n$$

Bad

which can be made arbitrarily close to 1 for large n.

b) The previous analysis for bin 1 does not depend on the number of balls in the other bins. If it is known that bin 1 contains less than k balls, then this increases the probability that balls will be assigned to some other bin x. If $k-1$ balls are already in bin 1, then no subsequent balls will be placed in bin 1, so the probability that any particular one will be placed in bin x increases from $\frac{1}{n}$ to $\frac{1}{n-1}$. This only increases the probability that bin x will contain more balls, which improves the probability that it will contain at least k balls. So any bin contains at least k balls with probability at least $\frac{1}{\sqrt{n}}$. ?

Problem 5:

a) If there are αn balls at the beginning of a round, this is the problem of placing αn balls randomly in n bins. As the ith ball is added, there are at most $i-1$ non-empty bins (if all balls previously went to empty bins; if there were earlier collisions than there are less). So the probability of the ith ball landing in a previously occupied bin is bounded above by $\frac{i-1}{n}$. By linearity of expectation, the expected number of balls landing in previously occupied bins is

$$\sum_{i=1}^{\alpha n} \frac{i-1}{n} = \frac{1}{n} \left(\sum_{i=1}^{\alpha n} i - \sum_{i=1}^{\alpha n} 1 \right) = \frac{1}{n} \left(\frac{(\alpha n)(\alpha n + 1)}{2} - \alpha n \right) = \frac{\alpha n(\alpha n - 1)}{2n} = \frac{\alpha^2 n - \alpha}{2} < \frac{\alpha^2 n}{2}$$

This probability is a sum of indicator random variables indicating whether the ith ball is placed in a previously occupied bin; thus, the Chernoff bound applies. By the Chernoff bound, the probability that the number of balls remaining is more than $\frac{2}{1.9}$ times the expectation, or $\frac{\alpha^2}{1.9}$, is $e^{-\frac{\frac{\alpha^2 n}{2}\left(\frac{2}{1.9}\right)^2}{2}}$, so at most $\frac{\alpha^2}{1.9}$ balls remain with high probability.

handle conditioning

b) We begin with n balls, i.e. $\alpha = 1$. At each stage, we have the recurrence

$$T(\alpha n) = 1 + T\left(\frac{\alpha^2 n}{1.9}\right)$$

Define α_1 to be 1 and $\alpha_n = \frac{\alpha_{n-1}^2}{2}$. Then the recurrence is

$$T(\alpha_i n) = 1 + T(\alpha_{i+1} n)$$

This recurrence continues until an i such that $\alpha_i < \frac{1}{n}$, at which point with high probability no balls remain. Taking the log of the recurrence for α_n, we find that $\log \alpha_n = \log \frac{\alpha_{n-1}^2}{2} = 2 \log \alpha_{n-1} - \log 2 \leq 2 \log \alpha_{n-1}$. The recurrence until $\log \alpha_i \leq -\log n$. Changing variables to $\beta_i = \log \alpha_i$ and $m = \log n$, we have $\beta_i \leq 2\beta_{i-1}$, a simple exponential recurrence that terminates once β_i reaches m. So the solution in terms of m is $O(\log m) = O(\log \log n)$.

whp?

03/06/97