CS 4540, Advanced Algorithms
Homework 1

Mon, Aug 30, 2010 Due Fri, Sept 10, 2010

Problem 1

Exercise 3, pages 782(bottom)-784(top) of Chapter 13 "Randomized Algorithms" by Kleinberg and Tardos (posted in class lecture notes and assinged reading 08-23-10). Recommendation: In addition to lecture notes and assigned reading of 08-23-10, it will be quite useful to go over Solved Exercise 1, pages 776(bottom)-779(top) of Chapter 13 "Randomized Algorithms" by Kleinberg and Tardos.

Problem 2

Suppose that n distributed non-communicating processes want to access a machine M where, if more than two processes attempt to access M at the same time unit, then M rejects all requests, but if at most two processes attempt to acess M at the same time unit, then M satisfies their requets. (a) Suppose that, at each time unit, each process attempts to access M with probability $\frac{1}{n}$. Determine a value t as a function of n and c such that, the probability that all processes n have succeeded in accessing M at least once after t time units is at least $1 - \frac{1}{n^c}$. (b) Suppose that, at each time unit, each process attempts to access M with probability $\frac{2}{n}$. Determine a value t as a function of n and c such that, the probability that all processes n have succeeded in accessing M at least once after t time units is at least $1 - \frac{1}{n^c}$. (Of course, in both cases, aim for the smallest value t for which you can prove the required high probability statement). Hint: Just like Problem 1 above, this is a relatively straightforward modification of the contention resolution problem (lecture 08-23-10).

Problem 3

Suppose that n distributed non-communicating processes want to access a machine M where, if more than one processes attempt to access M at the same time unit, then M rejects all requests. In class (lecture 08-23-10) we proved the following "high probability" statement: If at each time unit each process attempts to access M with probability $\frac{1}{n}$, then, for $t = (c+1)\,\mathrm{e}\,n \ln n$ time units, the probability that all processes n have succeeded in accessing M is at least once is at least $1 - \frac{1}{n^c}$. Alternatively, the probability that there exists at least one process that has failed to access M after t time units is at most $\frac{1}{n^c}$. Suppose that one wants to compare the above value of t to the expected number of time units by which all processes have accessed M at least once. So let T be a random variable denoting the number of time units by which all processes have accessed M at least once, and we are interested in $\mathrm{E}[T]$. Since $\mathrm{E}[T]$ does not give any immediate "high probability" guarantee, one would expect that $\mathrm{E}[T] < t$, and would want to quantify what is the additional number of rounds required to get the "high probability" statement. (a) Show that $\mathrm{E}[T] \leq \mathrm{e}\,n\,H(n)$, where $H(n)$ is the n-th harmonic number. (b) Compare and comment on the bounds for t and $\mathrm{E}[T]$. Hint: For part (a), you may write $T = X_1 + \ldots + X_n$, for suitably defined random variables X_i, and proceed in the spirit of the analysis of coupon collection (lecture 08-27-10).

Problem 4

Let X be a non-negative random variable (ie $\Pr[X < 0] = 0$) and let $\mu = \mathrm{E}[X]$. Recall Markov's inequality: For any $\lambda > 1$

$$\Pr[X > \lambda\,\mu] < \frac{1}{\lambda}\,. \tag{1}$$

Recall also the proof of Markov's inequality:

$$
\begin{aligned}
\mu &= \sum_{x \geq 0} x \Pr[X = x] \\
&= \sum_{0 \leq x \leq \lambda\mu} x \Pr[X = x] + \sum_{x > \lambda\mu} x \Pr[X = x] \\
&> \sum_{x > \lambda\mu} x \Pr[X = x] \\
&\geq \sum_{x > \lambda\mu} \lambda \mu \Pr[X = x] \\
&= \lambda \mu \sum_{x > \lambda \mu} \Pr[X = x] \\
&= \lambda \mu \Pr[X > \lambda \mu] \quad . \\
\text{Thus} \qquad \mu &> \lambda \mu \Pr[X > \lambda \mu] \quad , \qquad \text{implying} \\
\Pr[X > \lambda \mu] &< \frac{1}{\lambda} \quad .
\end{aligned}
$$

Now let Y be a non-negative random variable (ie $\Pr[Y < 0] = 0$), let $\mu = \mathrm{E}[Y]$, and suppose that we have the additional information that $\Pr\left[Y < \frac{\mu}{2}\right] = 0$. Intuitively, with this additional information, one would expect to be able to get a stronger bound than (1). Indeed, prove that, for any $\lambda > 1$

$$\Pr[X > \lambda \mu] < \frac{1}{2\lambda - 1} \ , \tag{2}$$

and argue that (2) is a stronger bound than (1).