cs473 Algorithms

Problem Set #7

Prof. Michael A. Forbes — Due: *Fri., 2026-04-10 17:00*

All problems are of equal value.

1. (Multiplicative Chernoff Bound). Let $\mathsf{X}_1, \ldots, \mathsf{X}_n$ be independent random variables taking values over the continuous real interval $[0,1]$, where we do not assume they are identically distributed. Let $\mathsf{X} = \sum_i \mathsf{X}_i$. Show the following.

   (a) For $r \in (-\infty, \ln 2]$, prove that $\mathbb{E}[\mathrm{e}^{r\mathsf{X}}] \le \mathrm{e}^{r\mathbb{E}[\mathsf{X}] + r^2\mathbb{E}[\mathsf{X}]}$, where you may use (without proof) that $1 + z \le \mathrm{e}^z$ for all $z \in \mathbb{R}$, and $\mathrm{e}^z \le 1 + z + z^2$ for $z \le \ln 2$.

   (b) Explain how the above used the independence of the X_i.

   (c) Apply Markov's inequality ($\Pr[\mathsf{Y} \ge a] \le \mathbb{E}[\mathsf{Y}]/a$) to $\mathrm{e}^{r\mathsf{X}}$, and optimize over r, to conclude that:

      i. For $0 \le \epsilon \le \ln 4$, $\Pr[\mathsf{X} \ge (1+\epsilon)\mathbb{E}[\mathsf{X}]] \le \mathrm{e}^{-\epsilon^2\mathbb{E}[\mathsf{X}]/4}$.

      ii. For $\epsilon \ge \ln 4$, $\Pr[\mathsf{X} \ge (1+\epsilon)\mathbb{E}[\mathsf{X}]] \le 2^{-\epsilon\mathbb{E}[\mathsf{X}]/2}$.

      iii. For $0 \le \epsilon \le 1$, $\Pr[\mathsf{X} \le (1-\epsilon)\mathbb{E}[\mathsf{X}]] \le \mathrm{e}^{-\epsilon^2\mathbb{E}[\mathsf{X}]/4}$.

      iv. (Additive Chernoff Bound) For $\epsilon \ge 0$, $\Pr[|\mathsf{X} - \mathbb{E}[\mathsf{X}]| \ge \epsilon \cdot n] \le 2\mathrm{e}^{-\epsilon^2 n/4}$.

   (d) Suppose now you have m independent identically random variables $\mathsf{Y}_1, \ldots, \mathsf{Y}_m$ over $\{0,1\}$, where each for each Y_i, $\Pr[\mathsf{Y}_i = 1] = \frac{\lg m}{m}$. Define $\mathsf{Y} = \sum_i \mathsf{Y}_i$. Bound the tail probability $\Pr[\mathsf{Y} \ge 2\mathbb{E}[\mathsf{Y}]]$, using both (i) and (iv) above, and compare these bounds.

   Note: The above omits one range of parameters, where one can show that $\Pr[\mathsf{X} \ge (1+\epsilon)\mathbb{E}[\mathsf{X}]] \le \mathrm{e}^{-(1+\epsilon)\ln(1+\epsilon)\mathbb{E}[\mathsf{X}]/4}$ if $\epsilon \ge 1$.

2. Balls and Bins. Kleinberg-Tardos Chapter 13, Problem #13.

3. In lecture we saw the random contraction algorithm for computing a global min-cut for undirected (and unweighted) graphs. In particular, we saw that if C is a minimum cut then the probability one round of random contraction will output C is $\ge 1/\binom{n}{2}$. This problem will explore the ability of the algorithm to output *almost* minimum cuts.

   Let $k \ge 1$ be an integer, and suppose $\hat{C}$ is a k-min-cut, in that $\hat{C}$ has value $|\hat{C}| \le k \min_C |C|$.

   (a) Prove that after $n - 2k$ randomly contracted edges, the probability the cut $\hat{C}$ has survived (that is, none of its edges have been contracted) is at least $\frac{1}{\binom{n}{2k}}$.

   (b) Suppose we modify the algorithm, so that after the first $n - 2k$ contractions, we output a random cut from the contracted graph. Prove that $\hat{C}$ is output with probability $\ge \frac{1}{2^{2k}\binom{n}{2k}}$.

   Hint: what is an upper bound for the number of cuts in the contracted graph?

   (c) Conclude that there are at most $O(n^{2k})$ k-minimum cuts.