Randomized Algorithms

CME309/CS365, Winter 2012-2013, Stanford University
Instructor: Ashish Goel
Final Exam. Due 3/21/2013 @ 5:00pm.

Collaboration policy: No collaboration allowed. You may use the textbook and any posted references.

Non-letter grade students: please do any two problems. If you do more, we will grade any two.

1. $G(n,p)$ is a graph with n vertices that is generated by taking each pair of vertices i, j and adding the edge $e = (i,j)$ independently with probability p.

   (a) What is the expected number of distinct spanning trees of $G(n,p)$?
   Hint: The total number of distinct spanning trees of a complete graph is n^{n-2}.

   (b) Find an upper bound for the probability that $G(n,p)$ is not connected for $p = \frac{k \ln n}{n}$ for some large enough constant k. You may assume that $(1 - \frac{x}{n})^n \approx e^{-x}$. Use this to argue that $G(n,p)$ is connected with high probability when $p = \frac{k \ln n}{n}$.
   Hint: $\binom{n}{i} \leq \left(\frac{ne}{i}\right)^i$.

2. Let G be a 3-colorable graph. Consider the following algorithm for coloring the vertices of G with 2 colors so that no triangle of G is monochromatic. The algorithm begins with an arbitrary 2-coloring of G. While there is a monochromatic triangle in G, it chooses one such triangle, and changes the color of a randomly chosen vertex of that triangle. Derive an upper bound on the expected number of such recoloring steps before the algorithm finds a 2-coloring with the desired property.

3. Consider a long stretch of road on a highway and cameras that are placed at m different locations. Whenever a car passes by a camera, it takes a picture of the license plate of the car along with the direction of the car. Unfortunately, each camera only has a small amount of memory which stores k of these (license plate, direction) pairs. Our goal is to design an algorithm for the camera's storing logic. Data should be stored in a way such that for any pair of cameras i, j, we can estimate the number of cars that either entered or exited between these two cameras using only the stored data in time $O(k)$. Specifically, let $N(i,j)$ denote the number of cars that either entered or exited between i and j and let $\hat{N}(i,j)$ be your estimate. Then find an algorithm such that $\mathbf{Pr}[|\hat{N}(i,j) - N(i,j)| < \delta N(i,j)] > 1 - \epsilon$ for k only depending on δ and ϵ regardless of how large m is or how much time has passed.

4. You are given an $n \times n$ matrix A and a non-negative $n \times 1$ vector x such that $Ax = b$. Assume that every entry of b is at least λ and every $A_{i,j} \in [0,1]$. An integer vector r is a γ-approximation if $b \leq Ar \leq \gamma b$. Give a randomized rounding algorithm to compute r and use Chernoff bounds to obtain bounds on γ when

   (a) $\lambda = 1$

   (b) $\lambda = \log n$

   (c) $\lambda = \sqrt{n}$