Neal Young

Problem Set 5

1. Recall the Chernoff bound we proved in class: the probability that any sum of independent random variables in $[0,1]$ exceeds its expected value μ by more than a factor of $1+\epsilon$ is less than $\alpha(\epsilon)^{\mu}$, where

$$\alpha(\epsilon) = \frac{e^{\epsilon}}{(1+\epsilon)^{1+\epsilon}} < 1.$$

   Let $\beta(\epsilon) = \ln 1/\alpha(\epsilon)$ so that $\alpha(\epsilon) = e^{-\beta(\epsilon)}$.

   Using whatever means you like (a calculator, Mathematica, Maple, or calculus (e.g. Taylor series)), try to understand and describe the function β as best you can. In particular, try to get a reasonable approximation for β (say, within constant factors) in terms of ϵ and ϵ^2.

   It may help to consider the ranges $\epsilon \leq 1$ and $\epsilon \geq 1$ separately.

   Proofs of the quality of your approximation would be nice but are not required.

2. Consider the following problem:

   The input is a bipartite graph $G = (V, W, E)$ and a number ℓ. The vertices in V are ordered $v_1, v_2, \ldots, v_{2n}$, and each odd vertex forms a *couple* with the next vertex. That is, v_1 and v_2 form a couple, v_3 and v_4 form a couple, and so on.

   The goal is to choose one vertex from each couple (n vertices in total), in such a way that no vertex in W has more than ℓ of its neighbors chosen. (This may or may not be possible for any given input.)

   No polynomial-time algorithm is known for this problem.

   You are to develop and analyze a randomized polynomial-time algorithm for finding an *approximate* solution.

   To start, consider the *relaxed* problem. The input is the same, but the goal is to assign non-negative weights to the vertices of V so that, for each couple, the sum of the weights on the two vertices is 1, while the total weight on edges incident to any vertex in W is at most ℓ.

   (a) Show how this relaxed version can be formulated as a linear program.

   (b) Suppose that the relaxed problem has a solution p.
   Consider the following random experiment. For each couple v, v', choose *one* of the two vertices at random: choose v with probability $p(v)$ and choose v' with probability $p(v') = 1 - p(v)$.
   Argue that for any vertex in W, the *expected* number of chosen neighbors is at most ℓ.

   (c) Figure out the smallest $\epsilon > 0$ you can so that (using the Chernoff bound) the probability that any given vertex in W has more than $(1+\epsilon)\ell$ chosen neighbors is less than $1/2|W|$. The value for ϵ will depend on ℓ.

   (d) Argue that the expected number of vertices in W with more than $(1+\epsilon)\ell$ chosen neighbors is at most $1/2$. Why does this mean that with probability at least $1/2$, all vertices in W will have at most $(1+\epsilon)\ell$ neighbors chosen?

   (e) Combine the above arguments to argue that there is a randomized algorithm for the original problem that runs in expected polynomial time and, if a solution exists, finds an *approximate* solution in that the number of chosen neighbors of any vertex $w \in W$ is at most $\ell + \epsilon\ell$.
   Describe how the quantity $\epsilon\ell$ varies as a function of ℓ.