CSE 521: Design and Analysis of Algorithms
Assignment #9
Due: Wednesday, March 10

Reading: from Williamson, see my email.

Problems: do *all* of the following problems.

1. Consider an arbitrary optimization problem, say a minimization problem, that can be formulated as an integer program (e.g., minimum vertex cover or travelling salesman, or any minimization problem in NP), and let I represent an instance (input) for this problem.

   Consider the value of
   $$\sup_I \frac{OPT(I)}{OPTFRACT(I)},$$
   where $OPT(I)$ is the value of the optimal solution to the problem (i.e., the optimal solution of the integer program) on instance I, and $OPTFRACT(I)$ is the value of the optimal fractional solution to the problem on instance I (i.e., the optimal solution of the linear programming relaxation). This quantity is called the *integrality gap* of the LP relaxation.

   - Show that the integrality gap for the vertex cover problem is at least 2.
   - Discuss whether or not it is possible for an approximation algorithm designed using an LP relaxation to achieve a better approximation guarantee than the integrality gap of the relaxation.

2. Prove theorem 2.3 (page 15) in Williamson's notes.

3. Consider the set multicover problem: The input is

   - a universe U of n elements and a value r_e for each element $e \in U$;
   - a collection $\mathcal{S}$ of subsets of U, and a cost $c(S)$ associated with each set $S \in \mathcal{S}$.

   The goal is to pick a subset of the sets in $\mathcal{S}$ (we allow a set to be chosen multiple times) such that for each element $e \in U$, e is covered at least r_e times and such that the sum of the costs of the selected sets is minimized.

   - Write down an integer programming formulation of this problem.
   - Write down the dual of the linear programming relaxation.

- Give an $O(\log n)$ factor randomized rounding algorithm for the set multicover problem. Your algorithm should have the property that the probability that the algorithm fails to output a valid multicover is at most n^{-c} for some constant $c > 0$. Be sure to show the details of your analysis.

  You can use the simple fact that if $p_1, p_2, \ldots, p_k$ are probabilities such that $\sum_{1 \leq i \leq k} p_i \geq 1$, then $\prod_{1 \leq i \leq k}(1 - p_i)$ is maximized when all the p_i's are equal to $1/k$. (I'm giving this to you in case you've forgotten elementary calculus.)