Due by the end of October 25, 2013

Problem 1
Let $X_1, X_2, \ldots, X_n$ be n independent random variables such that X_i takes value $1/p_i$ with probability p_i and 0 otherwise. Then, show that for any p such that $p \leq p_i$ for each i, any $\epsilon \in (0, 1)$, and any $N \geq n$, the following bound holds:

$$P\left(\left|\sum_{i=1}^{n} X_i - n\right| > \epsilon N\right) < 2e^{-0.38\epsilon^2 pN}$$

Problem 2
Given an undirected, unit-capacity graph G, consider the randomized edge contraction algorithm for the global min-cut problem.

(i) Show that in any run of the edge contraction algorithm, the edges contracted form a spanning tree of G.

(ii) Let $\mathcal{T}$ denote all the spanning trees in G. If we run the contraction algorithm, we will get a random spanning tree in $\mathcal{T}$ formed by the contracted edges, and we denote this distribution of spanning trees by D_1. On the other hand, if we assign a random weight in $(0, 1)$ to each edge and compute a minimum spanning tree using Kruskal's algorithm, then we obtain another distribution D_2 over $\mathcal{T}$. Show that these two distributions are identical.

Problem 3
In this problem, we have an undirected, unit-capacity graph G with n vertices and m edges.

(i) Given a constant $\alpha \geq 1$, run the edge contraction algorithm on G until we have 2α vertices and then output a cut uniformly at random from the cuts in the resulting contracted graph. Define an α-min cut to be a cut with size at most $\alpha\lambda$, where λ is the size of the min-cut. Show that any fixed α-min cut is output by the above randomized procedure with probability at least $\frac{1}{n^{2\alpha}}$. Conclude that the number of α-min cuts is at most $n^{2\alpha}$.

(ii) Suppose every edge in G fails with probability $p \in (0, 1)$ and let p_0 denote the probability that the graph G is disconnected as a result of the edge failures. Show that if $p^\lambda > 1/n^2$, then p_0 can be estimated up to $(1 \pm \epsilon)$ error (for any fixed ϵ) using a polynomial number of samples in the Monte Carlo method.

(iii) If $p^\lambda \leq 1/n^2$, show that $p_0 \in (1 \pm \epsilon)p_1$, where p_1 is the probability that some α^*-min cut fails for a fixed constant $\alpha^* \geq 1$. Write a DNF formula on all the α^*-min cuts such that p_1 can be estimated using the number of satisfying assignments of the DNF formula.

(iv) Use importance sampling to conclude that there exists a polynomial time algorithm to estimate p_0 up to a factor of $(1 \pm \epsilon)$.