CO 750 Randomized Algorithms (Winter 2011)
Assignment 2

Due: Thursday, March 10th.

Question 1: Consider adapting the randomized minimum cut algorithm to the problem of finding a minimum s-t cut in an undirected graph. In this problem, we are given an undirected graph G together with two distinguished vertices s and t. An s-t cut is a set of edges whose removal from G disconnects s and t; we seek an s-t cut of minimum cardinality. As the algorithm proceeds, the vertex s may get amalgamated into a new vertex as a result of an edge being contracted; we call this vertex the s-vertex. (Initially the s-vertex is s itself.) Similarly, we have a t-vertex. As we run the contraction algorithm, we ensure that we never contract an edge between the s-vertex and the t-vertex.

Show that there are graphs (or multigraphs) in which the probability that this algorithm finds a minimum s-t cut is exponentially small.

Question 2: When you buy a box of cereal it comes with a randomly chosen toy. Suppose there are m different types of toys, and the toys in each cereal box are independently and uniformly chosen. In vague terms, the question we'd like to answer is: how many boxes of cereal must you buy in order to collect at least one toy of each type?

More formally, let $X_1, X_2, ...$ be independent random variables, each uniformly distributed in $\{1, ..., m\}$. We would like to prove that, for some constants c_1 and c_2, if $t \geq c_1\, m \log m$ then

$$\Pr\left[\, \{X_1, ..., X_t\} = \{1, ..., m\} \,\right] \;\geq\; 1 - m^{-c_2}.$$

Prove this using the **Ahlswede-Winter** (or **Rudelson**'s) inequality.

Question 3: Suppose we have an algorithm $\mathsf{Test}(x, r)$ for deciding whether x belongs to a language L. (Here r is an additional input string which contains "advice", or simply random bits.) If $x \in L$ then $\mathsf{Test}(x, r) = 1$ for at least half of the possible values of r; a value of r such that $\mathsf{Test}(x, r) = 1$ is called a *witness* for x. If $x \notin L$ then $\mathsf{Test}(x, r) = 0$ always.

If we run the algorithm Test twice on an input $x \in L$ by choosing two numbers r_1 and r_2 independently and uniformly from S and evaluating $\mathsf{Test}(x, r_1)$ and $\mathsf{Test}(x, r_2)$, then we find a witness with probability at least $3/4$. Argue that we can obtain a witness with probability at least $1 - 1/t$ using the same amount of randomness by letting $s_i = r_1 i + r_2 \bmod p$ and evaluating $\mathsf{Test}(x, s_i)$ for values $0 \leq i \leq t < p$.

Question 4: Let $V = \{v_1, ..., v_m\}$ be a set of vectors in $\mathbb{R}^n$. We generate a random subset U of these vectors by (independently) adding each v_i to U with some probability p_i. We are interested in $\operatorname{rank} U = \dim(\operatorname{span}(U))$. Let $\mu = \mathrm{E}\left[\, \operatorname{rank} U \,\right]$. Prove that

$$\Pr\left[\, |\operatorname{rank} U - \mu| \geq \lambda \,\right] \;\leq\; 2\exp(-2\lambda^2/m).$$

You may use any theorems from the textbook.