CS261: Exercise Set #7

For the week of May 11–15, 2015

Instructions:

(1) *Do not turn anything in.*

(2) The course staff is happy to discuss the solutions of these exercises with you in office hours or on Piazza.

(3) While these exercises are certainly not trivial, you should be able to complete them on your own (perhaps after consulting with the course staff or a friend for hints).

Exercise 31

In Lecture #13 we rounded a linear programming relaxation to obtain a 2-approximation algorithm for the Vertex Cover problem, where the vertices can have arbitrary nonnegative costs. In this exercise we'll consider the special case where every vertex cost c_v is 1. Prove that the following is a 2-approximation algorithm for this case:

1. Given the input graph $G = (V, E)$, compute a maximal matching M of G.[1]
2. Let S denote the $2|M|$ endpoints of the edges of M, and return S.

Exercise 32

Recall from Lecture #13 our linear programming relaxation of the Vertex Cover problem (with nonnegative edge costs):

$$\min \sum_{v \in V} c_v x_v$$

subject to

$$x_v + x_w \geq 1 \qquad \text{for all edges } e = (v, w) \in E$$

and

$$x_v \geq 0 \qquad \text{for all vertices } v \in V.$$

Prove that there is always a *half-integral* optimal solution $\mathbf{x}^*$ of this linear program, meaning that $x_v^* \in \{0, \frac{1}{2}, 1\}$ for every $v \in V$.

[Hint: start from an arbitrary feasible solution and show how to make it "closer to half-integral" while only improving the objective function value.]

Exercise 33

Prove *Markov's inequality*: if X is a non-negative random variable with finite expectation and $c > 1$, then

$$\mathbf{Pr}[X \geq c \cdot \mathbf{E}[X]] \leq \frac{1}{c}.$$

[1] A matching M is maximal if no strict superset of M is a matching. A maximum matching is always maximal, but a maximal matching need not be maximum.

Exercise 34

Let X be a random variable with finite expectation and variance; recall that $\text{Var}[X] = \mathbf{E}\left[(X - \mathbf{E}[X])^2\right]$ and $\text{StdDev}[X] = \sqrt{\text{Var}[X]}$. Prove *Chebyshev's inequality*: for every $c > 1$,

$$\mathbf{Pr}[|X - \mathbf{E}[X]\,| \geq c \cdot \text{StdDev}[X]] \leq \frac{1}{c^2}.$$

[Hint: apply Markov's inequality to the (non-negative!) random variable $(X - \mathbf{E}[X])^2$.]

Exercise 35

There are n identical bins. Consider the following random process (which is relevant for the analysis of hashing, among other things):

1. Repeat n independent times:

   (a) Throw a ball into a bin chosen uniformly at random.

At the conclusion of this process, the expected number of balls in a given bin is 1 (why?). Prove that, with probability at least $1 - \frac{1}{n}$, no bin has more than $\frac{3 \ln n}{\ln \ln n}$ balls in it.

[Hint: this is a special case of the randomized rounding analysis we did in Lecture #14.]