COMS 4995 (Randomized Algorithms): Exercise Set #9

For the week of November 11–15, 2019

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 in the course discussion forum.

(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 40

Let $X_1, \ldots, X_n$ be i.i.d. geometric random variables with parameter $p \in (0, 1)$. Prove that

$$\mathbf{E}\left[\max_{i=1}^{n} X_i\right] = O\left(\frac{1}{p}\ln n\right).$$

Exercise 41

Recall the coupon collecting problem from Lecture #17, and let Y denote the number of samples (with replacement) from $\{1, 2, \ldots, n\}$ needed to see each element at least once. In lecture we used linearity of expectation and the expectation of geometric random variables to show that $\mathbf{E}[Y] \approx n \ln n$. Prove that with high probability (approaching 1 as $n \to \infty$), Y is at most $2n \ln n$.

[Hint: One approach is to compute (an upper bound on) the variance of Y and use Chebyshev's inequality. Another is to analyze directly the probability of missing a fixed coupon, and then take a Union Bound.]

Exercise 42

Recall the online learning setup from Lecture #17. Recall that the adversary is responsible for picking the reward vector $r^t : A \to [-1, 1]$ at each time step (where A is the finite set of actions). In lecture, we glossed over the distinction between two types of adversaries, *oblivious* and *adaptive*. An oblivious adversary commits up front to a sequence $r^1, \ldots, r^T$ of reward vectors (with knowledge of the learning algorithm but not any of its coin flips). An adaptive adversary commits to a reward vector r^t only after seeing everything that happened in the first $t - 1$ time steps (including the algorithm's past coin flips), in addition to the probability distribution p^t over actions chosen by the algorithm at time t (but not the coin flips at time t).

(a) Explain why the version of the FTPL algorithm given in lecture enjoys the stated regret guarantee (of $O(\sqrt{T \ln n})$) only for oblivious adversaries, and not for adaptive adversaries.

(b) Suppose we modify the algorithm so that, instead of sampling fictitious bonuses $\{X_a\}_{a \in A}$ once and for all at the beginning of the algorithm, we instead sample a fresh set $\{X_a^t\}_{a \in A}$ of bonuses at each time step (distributed as before, as twice a geometric random variable with parameter ϵ). At step t, the algorithm chooses the action that maximizes the sum of the current fictitious bonus X_a^t of the action and the cumulative reward so-far of the action. Verify that the regret bound from lecture applies to this modified algorithm, even with an adaptive adversary.

Exercise 43

Prove that an irreducible (finite) Markov chain has a unique stationary distribution.

[Hint: Assume that π, π' are both stationary distributions and consider a state j that minimizes π_j/π'_j; let c denote this minimum. Prove that every state i with positive transition probability P_{ij} to j must then also satisfy $\pi_i/\pi'_1 = c$. Use irreducibility to prove that $\pi = \pi'$.]

Exercise 44

Let $G = (V, E)$ be an undirected graph and $M \subseteq E$ a matching (i.e., a subset of edges, no two of which share a vertex). Recall that an *augmenting path* with respect to M is a path in G that begins and ends at unmatched vertices, and alternates edges not in M with edges in M. Prove that M is a maximum-cardinality matching of G if and only if there is no augmenting path with respect to M.

[Hint: Let M^* be a maximum-cardinality matching and consider the symmetric difference $M \triangle M^*$.]

Exercise 45

Let $G = (V, E)$ be a directed Eulerian graph, meaning that for every vertex $v \in V$, the in-degree $deg^-(v)$ equals the out-degree $deg^+(v)$. Assume also that G is strongly connected. Prove that for a random walk on G (where at each step, an outgoing edge from the current vertex is chosen uniformly at random), in the (unique) stationary distribution $\pi \in \mathbb{R}^V$, for every $v \in V$,

$$\pi_v = \frac{deg^+(v)}{m}.$$