CS264: Homework #8

Due by midnight on Wednesday, November 19, 2014

Instructions:

(1) Students taking the course pass-fail should complete the exercises. Students taking the course for a letter grade should also complete some of the problems — we'll grade your problem solutions out of a total of 40 points (with any additional points counting as extra credit).

(2) All other instructions are the same as in previous problem sets.

Lecture 15 Exercises

Exercise 47

A *fully polynomial-time approximation scheme (FPTAS)* for a maximization problem takes as input a problem instance and a parameter ϵ, and returns a feasible solution with objective function value at least $(1-\epsilon)$ times the maximum possible, in time polynomial in the size of the instance and in $\frac{1}{\epsilon}$. Prove that if a binary optimization problem with a maximization objective admits an FPTAS, then it also admits a pseudopolynomial-time algorithm (i.e., an exact algorithm that, when the objective function coefficients v_i are integers, runs in time polynomial in the instance size and $\max_{i=1}^n v_i$).

[Comment: there are also converses to this statement, but they are harder to prove.]

Exercise 48

Look up the terms "strongly NP-hard" and "ZPP." Recall from lecture that if a binary optimization problem admits an algorithm with polynomial smoothed complexity, then it also admits a randomized pseudopolynomial-time algorithm. Explain why this result implies that strongly NP-hard binary optimization problems do not have polynomial smoothed complexity, unless $NP \subseteq ZPP$.

Exercise 49

Here is the original form of the Isolation Lemma. Consider a binary optimization problem with a maximization objective, as defined in lecture. Let each v_i be a positive integer chosen uniformly at random from $\{1, 2, 3, \ldots, n^2\}$. Prove that with high probability (i.e., approaching 1 as $n \to \infty$), the resulting instance has a unique optimal solution.

Exercise 50

The point of this exercise is to investigate the independence assumptions we've been making in our smoothed analyses. Recall the following:

(i) In Lecture #13, our perturbation model assumed that each point in the TSP instance was drawn independently from a distribution with density function bounded above by $1/\sigma$.

(ii) In Lecture #14, our perturbation model assumed that each item weight in the Knapsack instance was drawn independently from a distribution with density function bounded above by $1/\sigma$.

(iii) In Lecture #15, our perturbation model assumed that each value in the instance of a binary optimization problem was drawn independently from a distribution with density function bounded above by $1/\sigma$.

Suppose we relax these independence assumptions to the following: for every point/item weight/value, conditioned arbitrarily on the other points/item weights/values, the conditional distribution has density function bounded above by $1/\sigma$.

Which of our three smoothed analyses from Lecture #13-15 continue to hold, with the same proofs, under this relaxed assumption? Justify your answer.

Lecture 16 Exercises

Exercise 51

Let D_1, D_2 denote two probability distributions on a finite set Ω. Define the *statistical distance* between D_1 and D_2 as

$$\max_{S \subseteq \Omega} \left| \mathbf{Pr}_{D_1}[S] - \mathbf{Pr}_{D_2}[S] \right| .$$

Define the ℓ_1 *distance* as

$$\sum_{\omega \in \Omega} \left| \mathbf{Pr}_{D_1}[\omega] - \mathbf{Pr}_{D_2}[\omega] \right| .$$

Prove that the statistical distance is precisely half the ℓ_1 distance.

Problems

Problem 24

Recall the scheduling problem mentioned in lecture, which provides another example of a binary optimization problem that can be solved in pseudopolynomial (and hence, by Lecture #15, smoothed polynomial) time. The input consists of n jobs, each with a known a positive processing time p_j, deadline d_j, and cost c_j. You should assume that all costs are integral. The feasible solutions are orderings of these jobs on a single machine. The *finishing time* F_j of a job j in an ordering is the sum of p_j and the processing times of all the jobs scheduled prior to j. A job is *late* in an ordering if its finishing time is strictly larger than its deadline. The goal of the problem is to compute the ordering of the jobs that minimizes the total cost of the late jobs.

(a) (3 points) Prove that if a subset S of jobs can all be scheduled to finish by their deadlines, then scheduling them in order of increasing deadline accomplishes this.

(b) (7 points) Give a dynamic programming algorithm that solves the scheduling problem in time polynomial in n and C, where $C = \max_{j=1}^{n} C_j$.

Problem 25

(15 points) The point of this problem is show that the pseudopolynomial-time (and hence, by Lecture #15, polynomial smoothed complexity) algorithm for the Knapsack problem extends to the case where there are k knapsacks, where k is a constant.

Precisely, there are again n items. For $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, k$, there is a positive value v_{ij} of placing item i in the jth knapsack and a positive weight w_{ij} of item i in the jth knapsack. For $j = 1, 2, \ldots, k$, the jth knapsack has a capacity W_j. You should assume that all the w_{ij}'s and W_j's are integers. The goal is to place items in knapsacks, respecting all knapsack capacities, to maximize the total value in all of the knapsacks.

Give an algorithm for the multiple knapsack problem that runs in time polynomial in n and $\max_{j=1}^{k} W_j$. The running time can have arbitrary dependence on the number of knapsacks k.

Problem 26

(15 points) The point of this problem is to investigate analogs of the Isolation Lemma that accommodate random constraints rather than a random objective function value. Consider a binary optimization problem with a maximization objective. Let the objective function $\max \sum_{i=1}^n v_i x_i$ be fixed, with $v_i > 0$ for every i. Also fixed is a preliminary feasible set $F \subseteq \{0,1\}^n$. You can assume that for every $i = 1, 2, \ldots, n$, there are members $\mathbf{x}$ of F with $x_i = 0$ and with $x_i = 1$. In addition, we consider a random linear constraint of the form $\sum_{i=1}^n w_i x_i \le W$. Assume that W is fixed and at least t, where $t \ge 0$ is the minimum number of 1s in a member of F. Assume that each w_i is drawn independently from $[0,1]$ according to one of our usual smoothed distributions, with density function $f_i : [0,1] \to [0, \frac{1}{\sigma}]$ for a parameter σ. The final (random) feasible set is defined as the set of $\mathbf{x} \in F$ with $\sum_{i=1}^n w_i x_i \le W$. Note that under our assumptions, the feasible set is non-empty with probability 1.

Define the *loser gap* L as follows: let V^* denote the maximum value of a feasible solution, let $\bar{\mathbf{x}}$ minimize $\sum_{i=1}^n w_i x_i$ over all $\mathbf{x} \in F$ with $\sum_{i=1}^n v_i x_i > V^*$ and $\sum_{i=1}^n w_i x_i > W$, and set $L = \sum_{i=1}^n w_i \bar{x}_i - W$. Note that L is a random variable. Prove that the probability (over the w_i's) that L is less than ϵ is at most $n\epsilon/\sigma$.

[Hint: follow the same proof template as for analyzing the winner gap. Analyze the probability that a variable x_i is "ϵ-bad," in the sense that if $\bar{\mathbf{x}}^{(i)}$ denotes the maximum-value feasible solution with $x_i = 0$, then there is a solution $\mathbf{x}^{(i)} \in F$ with $x_i = 1$, $\sum_{j=1}^n v_j x_j^{(i)} > \sum_{j=1}^n v_j \bar{x}_j^{(i)}$, and $\sum_{j=1}^n w_j x_j^{(i)} \in (W, W + \epsilon)$.]

Problem 27

(15 points) Recall from Lecture #16 that we proved the following (the Leftover Hash Lemma). Suppose X is a random variable with collision probability $cp(X)$ at most $1/K$. Suppose $\mathcal{H}$ is a (2-)universal family of hash functions (from the range of X to the range $\{0, 1, 2, \ldots, M-1\}$), and h is chosen uniformly at random from $\mathcal{H}$. Then the statistical distance between the joint distribution of $(h, h(X))$ and of the uniform distribution (on $\mathcal{H} \times \{0, 1, 2, \ldots, M-1\}$) is at most $\frac{1}{2}\sqrt{M/K}$.

For this problem, assume that you have a sequence $X_1, \ldots, X_T$ of random variables, with the property that for every i and fixed values of $X_1, \ldots, X_{i-1}$, the (conditional) collision probability of X_i is at most $1/K$ (i.e., a "block source"). Prove that the statistical distance between the joint distribution of $(h, h(X_1), \ldots, h(X_T))$ and of the uniform distribution is at most $\frac{T}{2}\sqrt{M/K}$.

[Hint: One high-level approach is to prove, by downward induction on i, a bound of $\frac{(T-i)}{2}\sqrt{M/K}$ on the statistical distance between $(h, h(X_{i+1}), \ldots, h(X_T))$ and the uniform distribution for every fixed value of $X_1, \ldots, X_i$. The increase in statistical distance in the inductive step should come from the Triangle Inequality.]