CS369N: Problem Set #3

Due in class on Tuesday, November 29, 2011

Instructions: Same as the first homework.

Problem 11

(20 points) Recall in Lectures #7-8 we discussed the Balcan-Blum-Gupta (BBG) k-median algorithm. We first discussed the simpler version of the problem where we assume that every cluster in the optimal solution has at least $2b + 2$ points, where $b = \epsilon n(1 + \frac{5}{\alpha})$. (This is in addition to the assumption that the instance is $(1 + \alpha, \epsilon)$-isolated.)

Our goal in lecture was to recover an ϵ-close clustering, rather than to optimize the k-median objective per se. Show that, nevertheless, the BBG algorithm gives an $O(1)$-approximation to the k-median objective in $(1 + \alpha, \epsilon)$-isolated instances that satisfy the large clusters assumption. (The constant can depend on α.)

[Hint: recall that after Step 3 of the BBG algorithm, all but the non-well-separated points are correctly classified. Use the Triangle Inequality to charge the cost of incorrectly classified points to the cost of the optimal solution.]

Problem 12

This problem considers γ-stable instances of metric Max Cut (recall Lectures #9 and #10). The input is a complete undirected graph $G = (V, E, w)$ with nonnegative edge weights w that satisfy the Triangle Inequality. Recall that such an instance is $(1 + \epsilon)$-stable if the maximum cut stays the same no matter how you multiply the edge weights by factors in $[1, 1 + \epsilon]$. (Even for such scalings that yield weights that violate the Triangle Inequality.) Assume that $\epsilon > 0$ is small but constant. Let (A, B) denote the maximum cut.

(a) (3 points) Prove that for every vertex v, the total weight of its incident edges that cross (A, B) is at least $(1 + \epsilon)$ times that of those that do not.

(b) (4 points) Without loss of generality, we can scale all the edge weights so that they sum to n^2. Define the weight of a vertex as the sum of the weights of its incident edges. Prove that every vertex has weight at least n.

(c) (5 points) Construct a (polynomial-size, non-metric) graph $G' = (V', E')$ as follows. For every edge $v \in V$ with weight w_v, add $\lfloor w_v \rfloor$ vertices to V' ("copies of v"). For each $u, v \in V$, add an edge to E' between each copy of u and of v, with weight $w_{uv}/\lfloor w_u \rfloor \lfloor w_v \rfloor$. Prove that the property in part (a) continues to hold for the graph G' (perhaps with a different constant ϵ'). Prove that a maximum cut of G can be recovered from one of G'.

(d) (6 points) Prove that in G', for every vertex u, the maximum weight of an edge incident to u is at most a constant factor times the average weight of an edge incident to u.

(e) (7 points) Give a polynomial-time approximation scheme (PTAS) for $(1 + \epsilon)$-stable metric Max Cut instances.

[Hint: use both random sampling and brute-force search.]

Problem 13

(15 points) This problem considers a planted model for graph coloring, to complement the ones we saw in lecture for the minimum bisection and maximum clique problems. Fix an integer k (which you should view as a constant), a number $p \in (0,1)$ (also constant), and an integer n (which you should think of as going to infinity). Consider generating a random k-colorable graph as follows:

1. Each vertex is independently given a label uniformly at random from $\{1, 2, \ldots, k\}$.
2. For each pair of vertices with endpoints with different labels, include the corresponding edge (independently) with probability p.

The ensuing graph is k-colorable with probability 1, with the color classes corresponding to the subsets of same-labeled vertices.

Design a polynomial-time algorithm that recovers the planted k-coloring in such a random graph with high probability (for large n).

Problem 14

Recall the setting of online decision-making discussed in Lecture #13.

(a) (8 points) In both algorithms we discussed in class, we assumed that the time horizon T was known a priori. For one of the two algorithms (your choice), show how to adapt it and its analysis to obtain similar regret bounds for the case where T is only revealed at the end of the sequence of cost vectors.

(b) (7 points) Consider a setting where the actions correspond to s-t paths in a fixed graph, and at each time step a cost vector on edges is revealed (with all edge costs between 0 and 1). Which of the two algorithms discussed in lecture, if any, lends itself to a polynomial-time (in the size of the graph) implementation for this setting? Discuss what the regret bounds look like in this case.

Problem 15

Recall the setting of revenue-maximizing auctions discussed in Lecture #14. Here we discuss two extensions.

(a) (10 points) Instead of unlimited supply, suppose you have k identical goods and $n > k$ bidders, each of whom wants one. It turns out that if every bidder's valuation is drawn i.i.d. from a distribution G (under mild assumptions which you shouldn't worry about), the truthful auction that maximizes the expected revenue is the Vickrey auction with a reserve price r (where $r \in \arg\max p \cdot (1 - F(p))$). This auction sells to all of the buyers i that have a valuation v_i above r and are also amongst the top k valuations overall. All winners pay either r or the $(k+1)$th highest valuation, whichever is larger. As usual, define $\mathcal{C}_D$ as the set of all such auctions (i.e., the Vickrey auction with all possible choices of the reserve r).

Assume that $k \geq 2$ and design a truthful auction such that, for every input $v_1, \ldots, v_n$, your (randomized) auction should have expected revenue at least a constant fraction of that of every auction in $\mathcal{C}_D$ that sells to at least 2 buyers. (You don't have to compete with auctions of $\mathcal{C}_D$ that sell to only one bidder on input v, just like in Lecture #14).

(b) (6 points) Let's return to the unlimited supply case. Suppose now that bidders' valuations are drawn independently from *different* distributions $F_1, \ldots, F_n$. What do optimal auctions (that maximize expected revenue) look like in this case? What is $\mathcal{C}_D$? What can you say about whether or not there are auctions that are α-instance optimal with respect to $\mathcal{C}_D$ (e.g., with the usual constraint of selling to at least two bidders)?

(c) (4 points) This part is also about the unlimited supply case. Consider the class $\mathcal{C}$ of auctions that attempt to sell to the bidders at prices $p_1, \ldots, p_n$ with $p_1 \geq p_2 \geq \cdots \geq p_n$. Call these *nonincreasing* auctions. Give examples of sets $\mathcal{D}$ of distributions for which the corresponding set $\mathcal{C}_D$ is precisely the set $\mathcal{C}$ of nonincreasing auctions.

(d) (Extra credit) For valuations $v_1, \ldots, v_n$, let $\mathcal{M}^{(2)}(v)$ denote the maximum revenue achieved by a nonincreasing auction that sells to at least two bidders. Is there an auction that has expected revenue at least a constant fraction of $\mathcal{M}^{(2)}(v)$ on every input? Say whatever you can, positive or negative.