Georgia Tech S'22	CS 6550/8803: Advanced Algorithms & Uncertainty	
	Homework 4	
Out: *March 16*		Due: *April 5*

Instructions:

- Each problem is worth twenty-five points. **Attempt any 4 of the 5 problems.** If you submit all 5 then your total score will be the **sum of top 4 problem scores**.
- Upload your solutions to the problems as a single PDF on Gradescope. Please anonymize all your submissions (i.e., do not list your name in the PDF), but if you forget, it's OK.
- You may collaborate with any classmates, textbooks, Internet, etc. Please upload a brief "collaboration statement" listing any collaborators as the last page of your non-extra-credit solutions PDF on Gradescope. But after the collaboration, always **write your solutions individually**.
- If you choose to do extra credit, upload your solution to the extra credits as a single separate PDF file to Gradescope. Please again anonymize your submission.
- For each problem, you should aim to keep your writeup below one page. For some problems, this may be infeasible, and for some problems you may write significantly less than a page. This is not a hard constraint, but part of the assignment is figuring out how to easily convince the grader of correctness, and to do so concisely. "One page" is just a guideline: if your solution is longer because you chose to use figures (or large margins, display math, etc.) that's fine.

Problems:

§1 (VCG with Approximation Algorithms) Prove that the VCG reduction doesn't work for the following greedy approximation algorithm (i.e., if we find the assignment and the prices using this greedy approximation algorithm):

- Initialize $S_i = \emptyset$ for all bidders i.
- For j = 1 to m, let bidder $i_j := \text{argmax}_i\{v_i(S_i \cup \{j\}) - v_i(S_i)\}$ be the bidder who gets the most marginal benefit from item j. Award item j to bidder i (update $S_i := S_i \cup \{j\}$) and continue.

(Recall that in HW-1 we proved that for submodular bidders this greedy algorithm gives a 2-approximation to the optimal welfare. But there we assumed that the $v_i(\cdot)$ functions were given to us, instead of being reported by bidders.)

Specifically, prove that VCG with this algorithm is not Dominant Strategy Truthful: provide two valuations $v_1(\cdot), v_2(\cdot)$ over two items such that bidder 1 is strictly better off by lying to the VCG mechanism that uses the above approximation algorithm when their value is $v_1(\cdot)$ and bidder 2 reports $v_2(\cdot)$.

§2 (Fixed-Price-Auctions) In this problem we want to use Fixed-Price-Auctions (FPAs) from Lecture 19 to design a truthful approximation algorithm for Combinatorial Auctions with XOS bidders (without any Bayesian assumptions).

(a) (Robustness) Consider n bidders with XOS valuations. Suppose O is an allocation with supporting prices $\mathbf{q} \in \mathbb{R}^m_{\geq 0}$. Let $A \subseteq [m]$ be a subset of items j with $\delta q_j \leq p_j \leq \frac{1}{2} q_j$. Prove that the Fixed-Price-Auction with prices $\mathbf{p}$ gives welfare at least $\delta \cdot \mathbf{q}(A)$. (Thus, FPAs are robust to knowing supporting prices approximately, and that too for only a subset of the items.)

(Hint: Modify the proof from Lecture 19 where we proved this statement for $A = [m]$ and $\delta = 1/2$.)

(b) (Randomized FPA) Suppose we are given the value of the optimal total welfare OPT (just the total value, not the allocations or the supporting item prices). Consider the following set of $O(\log m)$ candidate prices $C := \{\mathrm{OPT}, \mathrm{OPT}/2, \mathrm{OPT}/2^2, \ldots, \mathrm{OPT}/(2^{2\log m})\}$. Consider running a Fixed-Price-Auction where each item j independently chooses its price uniformly at random in the set C. Prove that the expected welfare of this Fixed-Price-Auction gives an $O(\log m)$-approximation.

§3 (XOS contains Submodular) Prove that any monotone submodular function f over m items with $f(\emptyset) = 0$ can written as an XOS function (max over non-negative linear functions).

(Hint: Design a non-negative linear function $a_S(\cdot) : 2^{[m]} \to \mathbb{R}_{\geq 0}$ for every subset $S \subseteq [m]$ of items such that $a_S(\mathbf{1}_S) = f(S)$ and $a_S(\mathbf{1}_T) \leq f(T)$ for every $T \subseteq [m]$. Now prove that the XOS function $\max_S\{a_S(\cdot)\}$ equals submodular function $f(\cdot)$.)

§4 (Two Player Minimax Theorem) We say that a two player game with each player having m pure strategies is Zero-Sum if the sum of the payoffs is zero, i.e., for each $(s_1, s_2) \in [m] \times [m]$ we have $p_1(s_1, s_2) = -p_2(s_1, s_2) = -A(s_1, s_2)$ where $A \in \mathbb{R}^{n \times n}$ is a given payoff matrix. Let $\Delta := \{x \in [0,1]^m$ with $\|x\|_1 = 1\}$ denote the m-dimensional simplex. In this problem we will prove the *minimax theorem* where the x-player is trying to minimize and the y-player is trying to maximize to show

$$\min_{x \in \Delta} \max_{y \in \Delta} \sum_{ij} x_i y_j A_{ij} = \max_{y \in \Delta} \min_{x \in \Delta} \sum_{ij} x_i y_j A_{ij}. \tag{1}$$

(a) Prove that the RHS $\leq$ LHS in (1) by arguing that for the x-player playing second only helps in minimizing the value. (Alternatively, this can be thought of as for the y-player playing second only helps in maximizing the value.)

(b) Next, we use Online Learning to prove the other direction of (1), which would show that in a 2-player zero-sum games it doesn't matter whether you go first or second against a "smart" opponent.

Assume that matrix $A \in [-1,1]^{n \times n}$ after rescaling. Consider an Experts problem setup (Lecture 10) where in round $t \in \{1, \ldots, T\}$ an Online Learning algorithm decides to play $y^{(t)} \in \Delta$ and receives the best-response reward from player-x, i.e., the reward is $\min_{x^{(t)} \in \Delta} \sum_{ij} x_i^{(t)} A_{ij} y_j^{(t)}$.

(i) Prove that $\sum_t \sum_{ij} x_i^{(t)} A_{ij} y_j^{(t)}$ is at most T times the RHS of (1).

(ii) If R denotes the average regret for the Experts problem when the per-step costs are in $[-1, 1]$, prove that

$$\sum_t \sum_{ij} x_i^{(t)} A_{ij} y_j^{(t)} \geq \max_{y \in \Delta} \sum_t \sum_{ij} x_i^{(t)} A_{ij} y_j - TR.$$

(iii) Deduce from the above two parts that in (1) we have RHS $\geq$ LHS $-R$. Since the average regret R tends to 0 with the number of rounds T, we have equality in (1).

§5 (PoA for Max-Cost) Consider the non-atomic flow problem from Lecture 21. Imagine there is a single source s and a single destination t, and there is 1 unit of traffic flowing between them. Suppose the edge costs are affine (i.e., $c_e(x_e) = a_e x_e + b_e$ for some $a_e, b_e \in \mathbb{R}_{\geq 0}$) and each participant takes their greedy shortest path. In this problem we will prove the Price of Anarchy (PoA) for the max-cost objective (i.e., the ratio of the max-cost for equilibrium flow to that of the max-cost for the optimal flow), where max-cost means the maximum total cost of any path taken by some participant.

(a) Show a graph where the PoA becomes 4/3 for the max-time objective.
(Hint: We have seen this graph in class.)

(b) Prove that the PoA for max-time is at most 4/3.

Extra Credit:

§6 (extra credit) In Problem 2(b), remove the assumption that the value of OPT is known. In other words, if we have n XOS-bidders with unknown valuations over m items (we don't even know any distributions on their valuations), design an $O(\log m)$-approximation Truthful mechanism.

(Hint: Randomly discard half the bidders: use these discarded bidders to get an approximate value of OPT without assigning them any items. Note that you can ask discarded bidders their valuations since they don't have any incentive to misreport. Now run the Fixed-Price-Auction on the remaining bidders. You need to handle the corner case separately where one bidder contributes nearly the entire value of OPT.)