601.436/636 Algorithmic Game Theory **Spring 2022**
Homework #6 **Due:** April 28, 2022, 4:30pm

Remember: you may work in groups of up to three people, but must write up your solution entirely on your own. Collaboration is limited to discussing the problems – you may not look at, compare, reuse, etc. any text from anyone else in the class. Please include your list of collaborators on the first page of your submission. You may use the internet to look up formulas, definitions, etc., but may not simply look up the answers online.

Please include proofs with all of your answers, unless stated otherwise. Your solution must by typeset (*not* handwritten), and must be submitted by gradescope.

---

1 Bulow-Klemperer (50 points)

This problem considers a variation on the Bulow-Klemperer theorem. Consider selling $k \geq 1$ identical items (with at most one given to each bidder) to bidders with valuations drawn i.i.d. from F. Prove that for every $n \geq k$, the expected revenue of the Vickrey auction (with no reserve) with $n + k$ bidders is at least that of the revenue-optimal auction for F with n bidders. Recall that the Vickrey auction in this context would give the k items to the k highest bidders, and charge each one of them the $(k + 1)$st highest bid.

Solution: We mimic the proof of the original Bulow-Klemperer theorem. Let $\mathcal{A}^*$ denote the optimal auction for F, i.e., the IC auction which maximizes expected revenue (over all IC auctions). Let $\mathcal{V}$ denote the Vickrey auction. We will be a little vague with notation and let $Rev(\mathcal{A}(b_1, b_2, \dots, b_m))$ denote the revenue of mechanism $\mathcal{A}$ (for any mechanism $\mathcal{A}$) when given bids $b_1, b_2, \dots, b_m$. Since both $\mathcal{A}^*$ and $\mathcal{V}$ are IC, we may assume that bids are the true valuations, so we are trying to prove that

$$\mathbf{E}_{v_1,\dots,v_{n+k}\sim F}[Rev(\mathcal{V}(v_1, \dots, v_{n+k}))] \geq \mathbf{E}_{v_1,\dots,v_n\sim F}[Rev(\mathcal{A}^*(v_1, \dots, v_n))]. \quad (1)$$

Let $\mathcal{A}$ denote the "fake auction" on $n + k$ players which works as follows: we simulate $\mathcal{A}^*$ for the first n players, and if item i is not sold to one of the first n players then we give it to player $n + i$ for free.

Since $\mathcal{A}$ exactly runs $\mathcal{A}^*$ on the first n players, and then does not obtain any revenue from the final k players, we know that

$$\mathbf{E}_{v_1,\dots,v_{n+k}\sim F}[Rev(\mathcal{A}(v_1, \dots, v_{n+k}))] = \mathbf{E}_{v_1,\dots,v_n\sim F}[Rev(\mathcal{A}^*(v_1, \dots, v_n))] \quad (2)$$

Now we will show that

$$\mathbf{E}_{v_1,\dots,v_{n+k}\sim F}[Rev(\mathcal{V}(v_1, \dots, v_{n+k}))] \geq \mathbf{E}_{v_1,\dots,v_{n+k}\sim F}[Rev(\mathcal{A}(v_1, \dots, v_{n+n}))], \quad (3)$$

which together with (2) will imply (1) and thus the theorem.

But this is just like the Bulow-Klemperer argument from class. It is easy to see that $\mathcal{A}$ is IC, since for any of the first n players $\mathcal{A}$ is just like $\mathcal{A}^*$ (which by definition is IC), and for the final k players their bid has no effect on their utility. Hence no player has incentive to lie about their valuation. Since both $\mathcal{V}$ and $\mathcal{A}$ are IC, we know that their expected revenue is equal to their expected virtual welfare, so we can switch to analyzing expected virtual welfare. The Vickrey auction gives items to the k largest valuations, and since F is regular the k largest valuations are also the k largest virtual valuations. Thus the Vickrey auction maximizes virtual welfare among all IC mechanisms which always sell all the items. $\mathcal{A}$ is an IC mechanism which always sells all the items, and thus Vickrey has expected virtual welfare at least as large as the expected virtual welfare of $\mathcal{A}$, which implies (3).

2 Unit-Demand Valuations (50 points)

Consider a combinatorial auction with n players and item set M with $m = |M|$. A player i has a *unit-demand* valuation if there exist parameters $v_i^1, v_i^2, \ldots, v_i^m \in \mathbb{R}_{\geq 0}$ (one parameter per item) such that $v_i(S) = \max_{j \in S} v_i^j$ for all $S \subseteq M$ (and $v_i(\emptyset) = 0$). In other words, the value of a bundle for player i is determined by the single most valuable element in that bundle (from the perspective of player i).

Give a mechanism for combinatorial auctions in which all players have unit-demand valuations which satisfies the following properties:

(a) It is incentive-compatible,

(b) It maximizes social welfare (i.e., maximizes $\sum_{i=1}^n v_i(S_i)$ where S_i is the bundle given to player i by the mechanism), and

(c) It runs in time polynomial in n and m.

Hint: Show that VCG can be implemented in polynomial time (in this setting) by reducing it to a well-known graph algorithm problem.

Solution: We will show that the VCG mechanism can be implemented in polynomial time. Since we proved that VCG is IC and maximizes social welfare, this will imply the theorem. To show that VCG can be done in polynomial time, we need to prove that the allocation and the prices can be computed in polynomial time. Recall that the VCG allocation is the social-welfare maximizing allocation. To prove that we can compute this efficiently, we introduce a bipartite *auction graph* $G = (U, M, E)$ where $U = [n]$ and $E = U \times M$ is all possible edges between U and M. We put weight $w(i,j) = v_i^j$ on every edge $(i,j) \in U \times M$.

Consider some allocation which assigns player i the bundle S_i, and let $x(i) = \arg\max_{j \in S_i} v_i^j$. Then clearly the set of edges $E' = \{(i, x(i)\}_{i \in [n] \wedge S_i \neq \emptyset}$ is a matching in G, and this matching has total weight

$$\sum_{(i,x(i)) \in E'} w(i, x(i)) = \sum_{(i,x(i)) \in E'} \max_{j \in S_i} v_i^j = \sum_{i \in [n]} v_i(S_i).$$

Thus the weight of this matching equals the value of the allocation. Similarly, suppose that there is some matching $\hat{E}$ in G. Then by assigning every $i \in [n]$ the item that it is matched with in $\hat{E}$ (or the empty set if it does not participate in the matching), we get an allocation with value equal to the weight of $\hat{E}$. Combining these two directions implies that the maximum weight matching in G is a maximum value allocation, i.e., the allocation which maximizes social welfare. Since there are well-known algorithm for computing maximum weight matchings in polynomial time, we can compute a social-welfare maximizing allocation in polynomial time.

Now we need to show that we can compute VCG prices efficiently. Recall that the price for player i is their externality: the social welfare of all other players if i had not participated, minus the social welfare of all other players when i does participate. Both are easy to calculate: for the latter, we just take our social-welfare maximizing matching and add up the edges which do not include i, and for the former we simply remove i from the graph and run our favorite maximum-weight bipartite matching algorithm, adding up the total weight. Because of our proof of equivalence between allocations and matchings, this will exactly be the VCG price.

Thus we can implement VCG in polynomial time, so it will have all three of the desired properties.