601.436/636 Algorithmic Game Theory **Spring 2022**
Homework #5 **Due:** April 14, 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 Auctions with Distinct Goods (50 points)

Consider an auction setting with a set M of distinct goods. Each bidder i has a publicly known subset $T_i \subseteq M$ of goods that it wants, and a private valuation v_i of getting them. An allocation is a partition $(A_1, A_2, \ldots, A_n)$ of the goods (or, equivalently, a function from the goods to the bidders). The *social surplus* of an allocation is the sum over all bidders who receive their desired items of their valuations: if $(A_1, A_2, \ldots, A_n)$ is an allocation then the social surplus is $\sum_{i:T_i \subseteq A_i} v_i$.

(a) (16 points) Prove that this is a single-parameter environment.

Solution: Let $X = \{x \in \{0,1\}^n : T_i \cap T_j \neq \emptyset \implies x_i x_j = 0\}$.

Consider some allocation $(A_1, A_2, \ldots, A_n)$. Let x be defined as $x_i = 1$ if $T_i \subseteq A_i$ and $x_i = 0$ otherwise. Since the allocation is a partition of the goods, if $T_i \cap T_j \neq \emptyset$ then either $T_i \not\subseteq A_i$ or $T_j \not\subseteq A_j$, and hence either x_i or x_j is 0 and so $x_i x_j = 0$. Thus $x \in X$. The social surplus of the allocation is equal to the social surplus of x, since $\sum_{i:T_i \subseteq A_i} v_i = \sum_{i:x_i=1} v_i = \sum_{i=1}^n v_i x_i$.

Now consider some $x \in X$. Define an allocation $(A_1, A_2, \ldots, A_n)$ by setting $A_i = T_i$ if $x_i = 1$. For every item which is not assigned this way (it is not in A_i for any i with $x_i = 1$), add the item to A_1 (the goods assigned to bidder 1). To show that this is a partition, note that by the definition of x we know that if $T_i \cap T_j \neq \emptyset$ then $x_i x_j = 0$, and thus any item is assigned to at most one bidder. And every item is assigned to at least one bidder because we took all unassigned items and assigned them to bidder 1. Hence this is a partition. Moreover, the social surplus of the allocation is equal to the social surplus of x, since $\sum_{i=1}^n v_i x_i = \sum_{i:x_i=1} v_i = \sum_{i:T_i \subseteq A_i} v_i$.

Thus for every $x \in X$ there is an allocation with the same social surplus, and for every allocation there is an $x \in X$ with the same social surplus. Hence this is a single-parameter environment.

(b) (17 points) Here is a natural greedy allocation rule, given a reported bid b_i from each player i:

(a) Initialize $S = \emptyset$, $X = M$.

(b) Sort and re-index the bidders so that $b_1 \geq b_2 \geq \cdots \geq b_n$.

(c) For $i = 1, 2, 3, \ldots, n$:

- If $T_i \subseteq X$, then:
  - Delete T_i from X.
  - Add i to S.

(d) Return S (give the bidders in S their desired items)

Does this algorithm define a monotone allocation rule? Prove it or give an explicit counterexample.

Solution: Yes, this allocation rule is monotone. Consider some bidder i and other bids $\boldsymbol{b}_{-i}$. Let $y > z \in \mathbb{R}_{\geq 0}$. We want to show that $x_i(y, \boldsymbol{b}_{-i}) \geq x_i(z, \boldsymbol{b}_{-i})$. Since the allocation is either 0 or 1 (bidder i either gets all items in T_i or does not), if $x_i(z, \boldsymbol{b}_{-i}) = 0$ then this is trivially true. So suppose that $x_i(z, \boldsymbol{b}_{-i}) = 1$. Let $S(z)$ be the bidders who bid more than z and which get their desired items from the algorithm when $b_i = z$. Then these bidders do not desire any items that are in T_i, or else $x_i(z, \boldsymbol{b}_{-i})$ would be 0. Let $S(y)$ be the bidders who bid more than y and which get their desired items from the algorithm when $b_i = y$. Since the algorithm will run identically in the two cases until it considers bidder i at bid y, we know that $S(y) \subseteq S(z)$. Then since $\sup_{j \in S(z)} T_j \cap T_i = \emptyset$, we get that $\sum_{j \in S(y)} T_j \cap T_i = \emptyset$. Thus the algorithm will give bidder i all of T_i if they bid y, and hence $x_i(y, \boldsymbol{b}_{-i}) \geq x_i(z, \boldsymbol{b}_{-i})$. So the allocation rule is monotone.

(c) (17 points) Prove that if all bidders report truthfully and have sets T_i of cardinality at most d, then the outcome of the allocation rule in (b) has social surplus at least $1/d$ times the social surplus of the optimal (surplus-maximizing) allocation.

Solution: Let $S^* \subseteq [n]$ be the optimal solution, so $T_u \cap T_v = \emptyset$ for all $u, v \in S^*$ and $\sum_{u \in S^*} v_u = OPT$. Let $S \subseteq [n]$ be the solution output by the algorithm (note again that $T_u \cap T_v = \emptyset$ for all $u, v \in S$). We will create a "charging function" $f : S^* \setminus S \to S \setminus S^*$. Let $i \in S^* \setminus S$. Because $i \notin S$, we know that the algorithm did not choose it. The only reason the algorithm would not choose it is because not all of its desired items are still available, and hence there is some $j \in S$ with $v_j \geq v_i$ such that $T_j \cap T_i \neq \emptyset$. Note that $j \notin S^*$, since otherwise S^* would not be feasible. Let $f(i) = j$ (if there are multiple such j, choose one arbitrarily).

Note that if $f(i) = f(j) = k$ for $i, j \in S^*$ with $i \neq j$, then since $T_i \cap T_j = \emptyset$ we know that $(T_i \cap T_k) \cap (T_j \cap T_k) = \emptyset$. This implies that $|\{i \in S^* \setminus S : f(i) = j\}| \leq |T_j| \leq d$ for all $j \in S$. Thus

$$\begin{aligned} OPT = \sum_{i \in S^*} v_i &= \sum_{i \in S^* \cap S} v_i + \sum_{i \in S^* \setminus S} v_i \leq \sum_{i \in S^* \cap S} v_i + \sum_{i \in S^* \setminus S} v_{f(i)} \\ &= \sum_{i \in S^* \cap S} v_i + \sum_{i \in S \setminus S^*} v_i \cdot |\{j \in S^* \setminus S : f(j) = i\}| \\ &\leq \sum_{i \in S^* \cap S} v_i + d \sum_{i \in S \setminus S^*} v_i \leq d \sum_{i \in S} v_i. \end{aligned}$$

Thus the surplus of the solution given by the algorithm is at least $1/d$ times the optimal surplus.

2 First-Price Auctions (50 points)

In this problem we compare the revenue achieved by first- and second-price auctions for a single good. Analyzing what happens in a first-price auction is not trivial; the easiest way to proceed is to assume that each valuation v_i is drawn i.i.d. from a known prior distribution F. A strategy of a bidder i in a first price auction is then a predetermined formula for (under)bidding: formally, a function $b_i(\cdot)$ that maps its valuation v_i to a bid $b_i(v_i)$. You should conceptually think of this strategy (i.e., this function) as being announced to all of the other bidders in advance; but of course, the other bidders do not know the actual value of v_i (and hence do not know the corresponding bid $b_i(v_i)$). We will call such a family $b_1(\cdot), \ldots, b_n(\cdot)$ of bidding functions a *Bayes-Nash equilibrium* if for every bidder i and every valuation v_i, the bid $b_i(v_i)$ maximizes i's expected payoff, where the expectation is with respect to the random draws of the other bidders' valuations (which, via their bidding functions, induce a distribution over their bids).

(a) (25 points) Suppose each valuation is an independent draw from the uniform distribution on $[0, 1]$. Prove that one Bayes-Nash equilibrium is given by setting $b_i(v_i) = v_i(n-1)/n$ for every i and v_i.

Solution: Consider some player i with arbitrary valuation v_i. Suppose that player i bids b. If $b > b_j(v_j)$ for all other players j then i wins the item and gets utility $v_i - b$, and otherwise player i gets utility 0. For any other player j, the probability that $b_j(v_j) < b$ is the probability that $v_j < \frac{n}{n-1}b$ by the definition of b_j. Since v_j is drawn uniformly from $[0, 1]$, this probability is precisely $\frac{n}{n-1}b$. Since each players valuation is independent of the others, we get that the expected utility of player i bidding b is

$$(v_i - b)\left(\frac{n}{n-1}b\right)^{n-1} = v_i\left(\frac{n}{n-1}\right)^{n-1} b^{n-1} - \left(\frac{n}{n-1}\right)^{n-1} b^n.$$

We can find the maximum of this by taking the derivative with respect to b, which is

$$v_i\left(\frac{n}{n-1}\right)^{n-1} (n-1)b^{n-2} - \left(\frac{n}{n-1}\right)^{n-1} nb^{n-1}.$$

When we set this to 0 and solve for b, we get that $b = \frac{n-1}{n} v_i$. Since the expected utility is 0 when $b = 0$ and is nonnegative when $b \leq v_i$, this must be a maximum (rather than a minimum). Thus player i should bid $\frac{n-1}{n} v_i$ no matter what v_i is, so setting $b_i = \frac{n-1}{n} v_i$ for all i and v_i is a Bayes-Nash equilibrium.

(b) (25 points) Prove that the expected revenue of the seller at this equilibrium of the first-price auction is exactly the expected revenue of the seller with truthful bidding in a Vickrey (second-price) auction (where in both cases the expectation is over the valuation draws).

Solution: First consider the equilibrium of the above first-price auction. The revenue of the seller will be $\mathbf{E}[\max_{i \in [n]} b_i(v_i)] = \mathbf{E}[\max_{i \in [n]} \frac{n-1}{n} v_i] = \frac{n-1}{n} \mathbf{E}[\max_{i \in [n]} v_i]$. So we just need to figure out the expectation of the maximum valuation. Clearly the maximum valuation is at most x if all valuations are at most x, which occurs with probability x^n. So the probability density function of the max is nx^{n-1}, and thus the expectation is $\int_0^1 nx^n dx = \frac{n}{n+1}[x^{n+1}]_0^1 = \frac{n}{n+1}$. Thus the expected revenue is $\frac{n-1}{n} \cdot \frac{n}{n+1} = \frac{n-1}{n+1}$.

Now consider the Vickrey auction. Since it is incentive-compatible we may assume that every player bids their true valuation, and thus the expected revenue is the expected value of the second-highest valuation. There are n possibilities for the highest bidder and $n-1$ possibilities for the second highest, so the pdf of the second highest value is $n(n-1)x^{n-2}(1-x)$. Thus the expected value is

$$\begin{aligned} \int_0^1 \left(n(n-1)x^{n-2}(1-x)x\right) dx &= n(n-1) \int_0^1 (x^{n-1} - x^n) dx \\ &= n(n-1) \left[\frac{1}{n} x^n - \frac{1}{n+1} x^{n+1}\right]_0^1 = n(n-1) \frac{1}{n(n+1)} = \frac{n-1}{n+1}. \end{aligned}$$