CS364A: Problem Set #1

Due in class on Thursday, January 20, 2011

Instructions:

(1) Students taking the course for a letter grade should attempt all of the following 5 problems; those taking the course pass-fail should attempt the first 3.

(2) Some of these problems are difficult. I highly encourage you to start on them early and discuss them extensively with your fellow students. If you don't solve a problem to completion, write up what you've got: partial proofs, lemmas, high-level ideas, counterexamples, and so on. This is not an IQ test; we're just looking for evidence that you've thought long and hard about the material.

(3) You may refer to your course notes, and to the textbooks and research papers listed on the course Web page *only*. You cannot refer to textbooks, handouts, or research papers that are not listed on the course home page. Cite any sources that you use, and make sure that all your words are your own.

(4) Collaboration on this homework is *strongly encouraged*. However, your write-up must be your own, and you must list the names of your collaborators on the front page.

(5) No late assignments will be accepted.

Problem 1

(a) (3 points) [From Lecture #2.] Prove that for every false bid $b_i \neq v_i$ by a bidder in a Vickrey auction, there exist bids b_{-i} by the other bidders such that i's payoff when bidding b_i is strictly less than when bidding v_i.

(b) (4 points) [From Lecture #2.] Consider a Vickrey auction with n bidders and suppose a subset S of the bidders decide to collude, meaning that they submit false bids in a coordinated way to maximize the sum of their payoffs. Prove necessary and sufficient conditions on the set S (in terms of the private valuations of the bidders) such that the bidders of S can increase their collective payoff via non-truthful bidding.

(c) (4 points) [From Lecture #2.] We proved that the Vickrey auction is truthful under the assumption that every bidder's utility function is *quasi-linear* — of the form $u_i(v_i, p_i) = v_i \cdot x_i - p_i$. State some significantly weaker assumptions on the utility functions $u_i(v_i, p_i)$ under which truthful bidding is a dominant strategy for every bidder.

(d) (4 points) [From Lecture #3.] Prove that for every single-parameter problem, every implementable allocation rule is monotone.

Problem 2

Recall the sponsored search auction problem discussed in Lectures #2 and 3: there are k slots, the jth slot has a known click-through rate (CTR) of α_j (nonincreasing in j), and the payoff of bidder i in slot j is $\alpha_j(v_i - p_j)$, where v_i is the (private) value-per-click of the bidder and p_j is the price charged per-click in that slot. For historical reasons, modern search engines do not use the truthful auction discussed in class. Instead, they use auctions derived from the *Generalized Second-Price (GSP)* auction, defined as follows:

(1) Rank advertisers by bid; assume without loss that $b_1 \geq b_2 \geq \cdots \geq b_n$.

(2) For $i = 1, 2, \ldots, k$, assign the ith bidder to the i slot.

(3) For $i = 1, 2, \ldots, k$, charge the ith bidder a price of b_{i+1} per click.

(a) (4 points) Prove that for every $k \geq 2$ and sequence $\alpha_1 \geq \cdots \geq \alpha_k > 0$ of CTRs, there exist valuations for the bidders such that the GSP auction is not truthful.

(b) [Do not hand in.] Fix CTRs for slots and valuations-per-click for bidders. We can assume that $k = n$ by adding dummy slots with zero CTR (if $k < n$) or dummy bidders with zero valuation (if $k > n$). A bid vector b is an *equilibrium* of GSP if no bidder can increase its payoff by changing its bid. Verify that this translates to the following conditions, assuming that $b_1 \geq b_2 \geq \cdots \geq b_n$: for every i and higher slot $j < i$,

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_j);$$

and for every lower slot $j > i$,

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1}).$$

(Derive these by adopting i's perspective and "targeting" the slot j.)

(c) [Do not hand in.] A bid vector b with $b_1 \geq \cdots \geq b_n$ is *envy-free* if for every bidder i and higher slot $j < i$,

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1});$$

and for every lower slot $j > i$,

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1}).$$

Verify that an envy-free bid vector is necessarily an equilibrium. (The terminology "envy-free" stems from the following interpretation: write $p_j = b_{j+1}$, for the current price-per-click of slot j; then the above inequalities say: "each bidder i is as happy getting its current slot at its current price as it would be getting any other slot and that slot's current price".)

(d) (6 points) A bid vector is *locally envy-free* if the inequalities in (c) hold for adjacent slots (i.e., for every i and $j = i-1, i+1$). Prove that, as long as the CTRs are strictly decreasing, a locally envy-free bid vector must in fact be envy-free.

[Hint: you might want to first prove that the bidders must be sorted in nonincreasing order of valuations.]

(e) (5 points) Prove that, for every set of α_j's and v_i's, there is an equilibrium of the GSP auction for which the outcome (i.e., the assignment of bidders to slots) and the prices paid precisely match those of the truthful auction discussed in class. If you want, you can assume that the CTRs are strictly decreasing.

[Hint: Recall that you know a closed-form solution for the payments made by the truthful auction. What bids would yield these payments in a GSP auction? Part (d) might be useful for proving that they form an equilibrium.]

Problem 3

Recall our discussion of Bayesian-optimal and prior-free revenue-maximizing auctions (Lectures #3-5).

(a) (5 points) Consider a single-item auction, where the ith bidder's valuation v_i is drawn from a regular distribution F_i. Bidders' valuations are independent, but notice that the distributions are *not* identical. Is the Bayesian-optimal auction equivalent to the Vickrey auction with a suitably chosen reserve price? Explain.

(b) (7 points) Consider the following single-bidder prior-free pricing game. There is an unknown value $v \in [1, h]$. If you offer a price $p \leq v$ you get p dollars and otherwise you get nothing. Design a randomized pricing strategy (i.e., a probability distribution over prices) such that your expected revenue is at least v/α for every v, where α is as close to 1 as you can manage.

(c) (6 points) Prove the best lower bound that you can on what values of α are achievable in the game in (b).

[Hint: use the probabilistic method.]

(d) (5 points) Now consider a general setting, with n bidders each with private valuation v_i. Let $X \subseteq \{0,1\}^n$ be the set of feasible allocation vectors. Assume that X is *downward-closed*, meaning that if $x \in X$, and $y \in \{0,1\}^n$ with $y \leq x$ component-wise, then $y \in X$ as well.

Use Myerson's Lemma to prove that there is a truthful mechanism that always outputs an allocation that maximizes the social surplus $\sum_{i=1}^n v_i x_i$ over $x \in X$ (assuming truthful bidding). [Cf., the truthful sponsored search auction in Lecture #2.]

(e) (7 points) Continuing with the setting of part (d): suppose we know that every valuation v_i lies in $[1, h]$. Prove that there is a randomized truthful mechanism that, for every valuation profile $v \in [1, h]^n$, has expected revenue at least

$$\frac{1}{\alpha} \max_{x \in X} \sum_{i=1}^{n} v_i x_i,$$

where α is the same number as in part (b).

Problem 4

This problem explores *composition theorems*, which identify conditions under which truthful mechanisms can be safely combined.

(a) (7 points) Show by example that there is a downward-closed setting (as defined in Problem 3(d)) and monotone allocation rules x^1, x^2 such that the following allocation rule $x^{\max}$ is *not* monotone: given valuations v, output whichever allocation out of $x^1(v), x^2(v)$ has a higher surplus (i.e., a larger value of $\sum_{i=1}^n v_i x_i$). If possible, don't rely on weird tie-breaking rules in your argument.

(a) (8 points) Identify some additional conditions (beyond monotonicity) on the allocation rules x^1, x^2 such that the induced "better-of-two" allocation rule *is* guaranteed to be monotone. Make your conditions as weak as you can.

Problem 5

A general issue in theoretical computer science is to understand the power and limitations of adding randomness to a computational model. We are only beginning to understand this issue in mechanism design; this problem provides some positive and negative results in the simple setting of digital goods auctions (with n bidders and n identical goods).

(a) [Do not hand in.] We first develop a different randomized competitive auction based on the profit extraction subroutine that we covered in lecture. Consider a bid vector b with all bids in the range $[1, h]$, with the property that $\mathcal{F}^{(2)}(b) \geq 2h$. Write $OPT(b)$ for $\mathcal{F}^{(2)}(b)$ and $OPT_{-i}(b)$ for $\mathcal{F}^{(1)}(b_{-i})$ — the optimal fixed-price revenue from b_{-i}, where any number of winners is allowed. Observe that for every i, $OPT(b)/2 \leq OPT_{-i}(b) \leq OPT(b)$.

(b) (4 points) Let $r_1(x)$ and $r_2(x)$ denote the functions that round x to the nearest odd power of 2 and the nearest even power of 2, respectively. (E.g., $r_1(14) = 2^3 = 8$ and $r_2(14) = 2^4 = 16$, while $r_1(18) = 32$ and $r_2(18) = 16$.) Prove that for every bid vector b that satisfies the assumption in (a), there is always a choice of $j = 1, 2$ such that $r_j(OPT(b)) \leq OPT(b)$ and also $r_j(OPT_{-i}(b)) = r_j(OPT(b))$ for every i.

(c) (4 points) Recall the ProfitExtract subroutine from lecture. Suppose that running this subroutine on a bid vector b with revenue target R results in a price p being charged to the winning bidders S. Let $i \in S$ and set $b'_i = +\infty$. Prove that running ProfitExtract with the new bid vector (b'_i, b_{-i}) and the same revenue target R yields the same outcome as before (the same winning set S and price p).

(d) [Do not hand in.] Consider the following randomized digital goods auction: given a bid vector b, independently for each bidder i, perform three steps: (1) Choose $j = 1, 2$ uniformly at random and set $R_i = r_j(OPT_{-i}(b))$; (2) set $b'_i = +\infty$ and run the ProfitExtract subroutine on the bid vector (b'_i, b_{-i}) with revenue target R_i, terminating with a set S_i of winners at price p_i (with $p_i|S_i| = R_i$); (3) finally, offer bidder i a posted price of p_i. Convince yourself that this is a truthful auction.

(e) (7 points) Prove that for every bid vector b that satisfies the assumption of part (a), the expected revenue of the auction in part (d) is at least $OPT(b)/4$.

[Hint: Let S denote the winning bidders when ProfitExtract is called on the bid vector b with revenue target $r_j(OPT(b)) \leq OPT(b)$. Argue separately about each bidder of S.]

(f) (8 points) Derandomize the auction in part (d) while losing only an additive factor of h in the revenue guarantee. I.e., design a deterministic auction, closely related to the auction in (d), that on every bid vector b that satisfies the assumption in (a), obtains revenue at least $(OPT(b)/4) - h$.

[Hints: Argue that it suffices to ensure the following: for every ℓ, at least $\lfloor \ell/2 \rfloor$ of the top ℓ bidders choose $j = 1$ in step (1), and at least $\lfloor \ell/2 \rfloor$ of the top ℓ bidders choose $j = 2$. Do you see how to ensure that the bidders accomplish this, using a parity argument applied to the different b_{-i}'s?]

(g) (2 points) Prove that the auction in (f) obtains revenue at least $(OPT(b)/4) - h$ for every bid vector b (not only those satisfying the assumption in (a)).

(h) (6 extra-credit points) Unlike the 4-competitive RSPE auction covered in class, this auction suffers an additional additive loss term. Prove that this is necessary in the following sense: for every constant $c > 1$, no deterministic (asymmetric) auction obtains revenue at least $\mathcal{F}^{(2)}(b)/c$ for every bid vector b.