CS364A: Problem Set #1

Due in class on Thursday, October 9, 2008

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) (2 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 #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 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) (4 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 a locally envy-free bid vector must in fact be envy-free.

(e) (7 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.

[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) For prior-free multi-item auctions, prove that the "limited supply" case reduces to that of "unlimited supply", in the following sense. Let k and n denote the number of identical items and of bidders, respectively. Suppose that, for some $c \geq 1$, there is a (possibly randomized) truthful auction for the $k = n$ case with expected revenue at least a $1/c$ fraction of the fixed-price benchmark

$$\mathcal{F}^{(2)}(b) := \max_{2 \leq i \leq n} i \cdot b_i,$$

for every bid vector b (we are assuming without loss that $b_1 \geq b_2 \geq \cdots \geq b_n$).

Using this assumption, prove that, for every $k \in \{2, 3, \ldots, n\}$, there is a (possibly randomized) truthful auction for the case with only k identical goods that has expected revenue at least a $1/c$ fraction of

$$\mathcal{F}^{(2,k)}(b) := \max_{2 \leq i \leq k} i \cdot b_i,$$

the optimal fixed-price revenue subject to the supply constraint.

(b) (5 points) Use Myerson's Lemma to prove that every *deterministic* truthful auction with identical goods is equivalent to an auction of the following form: given bid vector b, offer each bidder i a posted price (a "take-it-or-leave-it" offer, recall Lecture #2) of $t_i(b_{-i})$, where t_i is an arbitrary function of the other bids, with range $[0, +\infty]$, and ties (when $b_i = t_i(b_{-i})$) broken arbitrarily.

(c) (5 points) An auction of the form in (b) is *symmetric* if all of the functions $t_1(\cdot), \ldots, t_n(\cdot)$ are a common function $t(\cdot)$, which is itself symmetric (i.e., invariant under permutations of its arguments). Prove that for every constant $c > 1$, no deterministic symmetric auction for digital goods (i.e., with $k = n$) is c-competitive with respect to the fixed-price benchmark $\mathcal{F}^{(2)}(b)$.

[Hint: Consider bid vectors with only "high" and "low" bids.]

(d) [Do not hand in.] For the rest of this problem, consider a single-good auction ($k = 1$). Recall Myerson's Bayesian-optimal auction (Lecture #4) — i.e., the auction that maximizes expected revenue (with respect to the prior F) over all truthful auctions. Suppose the prior distribution F (with density f) satisfies the *monotone hazard rate (MHR)* condition, meaning that $f(x)/(1 - F(x))$ is nondecreasing over the support of f. Verify that such a distribution is regular in the sense of Lecture #4.

(e) (5 points) Prove that if F satisfies the MHR condition, then the probability (over the draw of bidder valuations) that Myerson's optimal auction for F successfully awards the good to a bidder is at least $1/e$.

[Hint: first (and for partial credit) prove the result for exponential distributions. What is the hazard rate of such a distribution?]

(f) (5 points) Instead of i.i.d. draws from a distribution F, suppose we know that the ith valuation v_i is drawn from the distribution F_i with positive density f_i on $[a_i, b_i]$. Assume that each F_i is regular in the sense of Lecture #4. Identify the Bayesian-optimal auction in this case, and prove its optimality.

[Hint: this shouldn't be that hard if you've absorbed the lemmas we proved for the i.i.d. case. Of course, whenever such a lemma applies, you can use it directly in your proof.]

Problem 4

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 $b_1(\cdot), \cdots, b_n(\cdot)$ 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) (7 points) Suppose each valuation is an independent draw from the uniform distribution on $[0, 1]$. Prove that one equilibrium is given by setting $b_i(v_i) = v_i(n - 1)/n$ for every i and v_i.

(b) (8 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 auction (where in both cases the expectation is over the valuation draws).

(c) (8 extra-credit points) Extend the conclusion in (b) to the case of an arbitrary distribution F with positive and differentiable density f on support $[0, 1]$.

[Hint: You can prove this directly, but Myerson's Lemma will shorten the argument somewhat.]

Problem 5

A general issue in theoretical computer science is to understand the power and limitations of adding randomness to a computational model. This issue is currently poorly understood 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., $r_1(12) = 2^3 = 8$ while $r_2(12) = 2^4 = 16$.) 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.

Bonus Problem

(10 extra credit points) [From Ken Steiglitz's *Snipers, Shills, & Sharks*.] eBay is the dominant online auction in most of the world, but Japan are China are important exceptions. Yahoo dominates eBay in Japan, and Taobao is fighting it out with eBay in China. Try to explain the reasons for the success of Yahoo and Taobao in penetrating these Internet auction markets by studying their business histories and practices, rules, rate structures, feedback reputation systems, and interfaces. To what extent can their relative success be attributed to cultural differences?