CS364A: Problem Set #2

Due in class on Thursday, February 3, 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 6

This problem shows that, for Bayesian-optimal mechanism design, "sufficient competition" can obviate the need for a reserve price.

(a) (3 points) Consider a distribution F that is regular in the sense of Lecture #4, and let φ denote the corresponding virtual valuation function. Prove that the expected virtual value $\varphi(v_i)$ of a valuation v_i drawn from F is zero.

(b) (6 points) Consider selling $k \geq 1$ identical items 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 Bayesian-optimal auction for F with n bidders.

[Thus, modest additional competition is at least as valuable as knowing the distribution F and employing a corresponding optimal reserve price.]

[Hints: To develop intuition explore the case of $k = n = 1$ and $F(x) = x$ on $[0, 1]$. In general, use Myerson's characterization of the expected revenue of a truthful auction. Condition on the values of the first n bidders and then argue about the expected impact of the final k bidders on the revenue of the no-reserve Vickrey auction, using part (a).]

(c) (6 points) In the same setup as (b), assume that $n \geq 2$ and consider the following alternative mechanism: pick one of the n bidders uniformly at random, say with bid r, and run the Vickrey auction with reserve r on the other $n - 1$ bidders. Prove that the expected revenue of this auction is at least $(n - 1)/2n$ times that of the Bayesian-optimal auction (with n bidders).

Problem 7

(a) (5 points) Consider a general mechanism design problem, with a set Ω of outcomes, and n players, where player i has a private real-valued valuation $v_i(\omega)$ for each outcome $\omega \in \Omega$. Suppose the function $f : \Omega \to \mathcal{R}$ has the form

$$f(\omega) = c(\omega) + \sum_{i=1}^{n} w_i v_i(\omega),$$

where c is a publicly known function of the outcome, and where each w_i is a nonnegative, public, player-specific weight. Such a function is called an *affine maximizer.*

Show that for every affine maximizer objective function f and every subset $\Omega' \subseteq \Omega$ of the outcomes, there is a truthful mechanism that optimizes f over Ω'.

[Hint: modify the VCG mechanism. Don't worry about individual rationality.]

(b) (3 points) For the rest of this problem, consider a combinatorial auction with a set S of m goods and n bidders. Assume that the valuation $v_i(\cdot)$ of bidder i depends only on its bundle T_i of goods; that it is nondecreasing (so $T_1 \subseteq T_2$ implies that $v_i(T_1) \le v_i(T_2)$); that $v_i(\emptyset) = 0$; and that v_i is *subadditive*, meaning that $v_i(T_1) + v_i(T_2) \ge v_i(T_1 \cup T_2)$ for every pair T_1, T_2 of disjoint subsets of goods.

In this and the next two parts, we consider only the winner determination problem (i.e., we don't worry about payments or truthfulness, just polynomial-time surplus maximization). Given S and $v_1, \ldots, v_n$, call the winner determination problem *lopsided* if there is an optimal allocation of goods in which at least half of the total surplus of the allocation is due to players that were allocated a bundle with at least $\sqrt{m}$ goods. (I.e., if $2\sum_{i \in A} v_i(T_i^*) \ge \sum_{i=1}^{n} v_i(T_i^*)$, where $\{T_i^*\}$ is the optimal allocation and A is the subset of bidders i with $|T_i^*| \ge \sqrt{m}$.)

Show that in a lopsided problem, there is an allocation that gives all of the goods to a single player and achieves an $\Omega(1/\sqrt{m})$ fraction of the maximum-possible surplus.

(c) (4 points) Show that in a problem that is not lopsided, there is an allocation that gives at most one good to each player and achieves an $\Omega(1/\sqrt{m})$ fraction of the maximum-possible surplus.

[Hint: use subadditivity.]

(d) (4 points) Give a polynomial-time $O(\sqrt{m})$-approximate winner determination algorithm for subadditive valuations.

[Hint: make use of a graph matching algorithm.]

(e) (4 points) Give a polynomial-time, $O(\sqrt{m})$-approximate, truthful combinatorial auction for subadditive valuations.

[Hint: use part (a).]

Problem 8

Consider the following single-parameter mechanism design problem based on the Knapsack problem. There are n bidders, each with a private valuation v_i and a publicly known size c_i. There is a publicly known budget (or "knapsack capacity") C. The feasible allocations correspond to subsets S of bidders for which $\sum_{i \in S} c_i \le C$. We assume (without loss) that $c_i \le C$ for every i. The problem of computing the surplus-maximizing feasible allocation is precisely the Knapsack problem. Recall that this problem is NP-hard but can be solved in pseudo-polynomial time using dynamic programming (in time poly$(n) \cdot C$ when the costs and budget are integral, and in time poly$(n) \cdot \max_i v_i$ when the valuations are integral). See any algorithms textbook for proofs of these facts.

(a) (4 points) Consider the following polynomial-time approximation algorithm for the Knapsack problem. First, sort all the bidders in decreasing order of valuation, greedily pack the knapsack (i.e., add the next bidder to the knapsack if and only if it fits), and call the resulting allocation S_1. Second, sort all the bidders in decreasing order of the ratio v_i/c_i, greedily pack the knapsack, and call the resulting allocation S_2. The algorithm returns either S_1 or S_2, whichever is better. Convince yourself (or look up the fact) that this is a $\frac{1}{2}$-approximation algorithm.

Prove that this approximation algorithm defines a monotone allocation rule, and hence can be extended to a truthful approximation mechanism.

[Note: You might have proved something useful for this back in Problem 4.]

(b) (5 points) Suppose, for this part of the problem only, that we have two knapsacks with known capacities. An obvious approximation algorithm is to use the algorithm in part (a) to pack one knapsack, and then to use it again to pack the other knapsack using the remaining bidders. Does this algorithm define a monotone allocation rule? Either prove that it does or give a counterexample.

(c) [Do not hand in.] We next review a fully polynomial-time approximation scheme (FPTAS) for the Knapsack problem. Given a parameter $\epsilon > 0$ consider the following algorithm $\mathcal{A}_\epsilon$:

- Round each v_i up to the nearest multiple of $V\epsilon/n$, call it v'_i.
- Multiply through by $n/V\epsilon$ to get integers $v''_1, \ldots, v''_n$.
- Solve the Knapsack problem for the v''_i's exactly using a pseudo-polynomial-time algorithm.

Recall (or look up the fact) that, if we set the parameter V to be $\max_i v_i$, then the algorithm $\mathcal{A}_\epsilon$ runs in polynomial time and gives a $(1-\epsilon)$-approximation for the Knapsack problem.

(d) (3 points) Prove that if V is fixed up front, then the algorithm $\mathcal{A}_\epsilon$ defines a monotone allocation rule (for any ϵ).

(e) (5 points) Suppose we try to first set $V = \max_i v_i$ and then run the algorithm $\mathcal{A}_\epsilon$. Prove that this combined algorithm does *not* always define a monotone allocation rule.

(f) (8 points) Give a truthful fully polynomial-time approximation scheme for the Knapsack problem with private valuations. The running time of your algorithm should be polynomial in both n and $1/\epsilon$.

[Note: Again, results along the lines of Problem 4 could be useful.]

Problem 9

This problems considers auctions that provide revenue guarantees of various forms.

(a) (3 points) Consider a digital goods auction (n bidders, n identical goods) with a twist: the auctioneer incurs a fixed production cost of 1 if there is at least one winner; if no goods are sold, then no such cost is incurred. Call an auction for this problem *budget-balanced* if, whenever there is at least one winner, the prices charged to the winners sum to exactly the cost incurred (namely, 1). Define the *surplus* of an outcome with winners S to be 0 if $S = \emptyset$ and $-1 + \sum_{i \in S} v_i$ otherwise.

Note that the surplus can be truthfully maximized in this problem using the extension of the VCG mechanism described in Problem 7(c). Prove that with the standard VCG payments (in which losers pay 0), the VCG mechanism is not budget-balanced — in fact, it can generate 0 revenue even when the auctioneer incurs cost 1.

(b) (2 points) Explain how to instantiate the ProfitExtract subroutine from Lecture #5 to obtain a non-trivial truthful, budget-balanced mechanism for the above problem (in which losers pay 0 and winners pay at most their bids).

(c) (8 points) The mechanism in (b) does not generally maximize the surplus. Precisely, show that the largest-possible difference (over all possible valuation profiles $\mathbf{v}$) between the maximum surplus and the surplus achieved by this mechanism is exactly $-1 + \sum_{i=1}^{n} \frac{1}{i}$, which is roughly $\ln n - 1$.

(d) (5 points) We can generalize the result in (b) as follows. Consider a digital goods auction with players N, in which the auctioneer incurs a (publicly known) cost of $C(S)$ when the set of winners is $S \subseteq N$. Assume that $C(\emptyset) = 0$, that C is nondecreasing (meaning $C(S) \leq C(T)$ whenever $S \subseteq T$), and that C is *submodular*, meaning that

$$C(T \cup \{i\}) - C(T) \leq C(S \cup \{i\}) - C(S)$$

whenever $S \subseteq T$ and $i \notin T$. (This is a set-theoretic type of "marginal returns". For example, when $C(S)$ depends only on $|S|$, submodularity becomes discrete concavity.)

The *Shapley value of i in S*, denoted $\chi_{Sh}(i, S)$, is defined as follows. For an ordering π of the players of S, let T_π denote those preceding i in π. Then $\chi_{Sh}(i, S) := \mathbf{E}_\pi[C(T_\pi \cup \{i\}) - C(T_\pi)]$, where π is chosen uniformly at random. In other words, assuming that the players of S are added to the empty set 1-by-1 in a random order, $\chi_{Sh}(i, S)$ is the expected jump in cost caused by i's arrival.

Prove that, under the assumptions on C above, $\chi_{Sh}(i, S) \geq \chi_{Sh}(i, T)$ whenever $S \subseteq T$.

(e) (7 points) By using Shapley values as prices, generalize the budget-balanced truthful mechanism in (b) to a digital goods auction with an arbitrary nondecreasing, submodular cost function. Be sure to prove that your mechanism is truthful.

Problem 10

Recall the multi-parameter pricing problem from Lecture #6, where there are m goods and a single bidder, who has a private valuation $v_j \sim F_j$ for each good j (where the F_j's can be different but are independent and regular). The bidder wants at most one of the goods. In class, we showed that there is a simple posted price (specifically, a common "virtual price") that gets at least half the revenue of the optimal posted price.

(a) (5 points) Prove (by example) that the optimal posted price need not be a constant virtual price. Show the largest gap (between the expected revenue of the optimal posted price and of the best common virtual price) that you can.

(b) (5 points) Prove (by example) that a randomized mechanism can get higher expected revenue than the best (deterministic) posted price.

[Hint: Consider supplementing a posted price by also offering a "lottery" that is a coin flip (to be resolved after the sale) between two different goods.]

(c) (15 points) Consider the optimal randomized mechanism, which you can assume is simply a list of lotteries over goods (as in the hint in (b)), with a price for each lottery. (The bidder receives the lottery that maximizes its expected utility.) Prove that there is a (deterministic) posted price that gets a constant fraction of the expected revenue of this optimal randomized mechanism.

Extra Credit

(Up to 25 extra credit points, redeemable by March 17th, 2011) Consider the following variant on the prior-free profit-maximization problem for digital goods. There are n bidders, named $1, 2, \ldots, n$, with private valuations $v_1, \ldots, v_n$. In this problem, we will *not* assume that the v_i's are in sorted order; as we will see shortly, such an assumption would not be without loss of generality. There are n identical goods and each bidder wants only one of them. Obtain $v^{(2)}$ from v by replacing the largest valuation with a copy of the second-largest valuation. (We do this to avoid the usual technical issues with inputs that have only one very high-value bidder.)

A vector $p_1, \ldots, p_n$ of take-it-or-leave-it offers is *monotone* if $p_1 \geq p_2 \geq \cdots \geq p_n$. For a valuation profile v, let $\mathcal{M}(v^{(2)})$ denote the largest revenue obtainable from the profile $v^{(2)}$ using a monotone take-it-or-leave-it-offer. For example, if $v^{(2)}$ happens to satisfy $v_1^{(2)} \geq \cdots \geq v_n^{(2)}$, then $\mathcal{M}(v^{(2)})$ is simply $\sum_{i=1}^n v_i^{(2)}$. If $v^{(2)}$ happens to be sorted in the opposite order, then $\mathcal{M}(v^{(2)})$ coincides with $\mathcal{F}^{(1)}(v^{(2)})$ (as you should verify).

Your goal is to design a (randomized) truthful auction that, for every (not necessarily sorted) valuation profile v, has (expected) revenue at least $\mathcal{M}(v^{(2)})/\beta$, where β is a constant. Make β as small as you can.