CS364A: Problem Set #1

Due to the TAs by noon on Friday, October 11, 2013

Instructions:

(1) Form a group of at most 3 students and solve as many of the following problems as you can. You should turn in only one write-up for your entire group.

(2) Turn in your solutions directly to one of the TAs (Kostas or Okke). You can give them a hard copy or send a soft copy by email to `cs364a-aut1314-submissions@cs.stanford.edu`. Please type your solutions if possible and feel free to use the LaTeX template provided on the course home page.

(3) 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.

(4) Except where otherwise noted, 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. If you do use any approved sources, make you sure you cite them appropriately, and make sure that all your words are your own.

(5) You can discuss the problems verbally at a high level with other groups. And of course, you are encouraged to contact the course staff (via Piazza or office hours) for additional help.

(6) No late assignments will be accepted.

Problem 1

(10 points) Identify a real-world system in which the goals of some of the participants and the designer are fundamentally misaligned, leading to non-trivial strategic behavior by the participants. You should also give proposals for how the system could be improved to mitigate the incentive problems. You should include:

- A description of the system, detailed enough that you can express clearly the incentive problems and your solutions for them.
- Evidence that participants are gaming the system in undesirable ways. This could be anecdotal evidence, something that has or can be measured, or something that you demonstrate yourself. (Feel free to include screen shots or pictures.)
- A convincing argument why your proposed changes would reduce or eliminate the strategic behavior that you identified.

We emphasize that a "system" could be any number of things — a Web site, a competition, a political process, etc.

Problem 2

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

(b) (4 points) 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) (3 points) We proved that the Vickrey auction is truthful under the assumption that every bidder's utility function is *quasi-linear* — of the form $u_i(v_i, p_i) = v_i \cdot x_i - p_i$ (where x_i is 1 if i wins and 0 otherwise). 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.

Problem 3

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 of generality 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 4

Recall the Knapsack problem studied in lecture, and the DSIC 2-approximation mechanism that we gave for it. 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 $\sum_{i \in S} c_i \leq C$. We assume that $c_i \leq C$ 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 $\text{poly}(n) \cdot C$ when the costs and budget are integral, and in time $\text{poly}(n) \cdot \max_i v_i$ when the valuations are integral). In lecture we discussed a well-known 2-approximation algorithm. See any algorithms textbook (or the instructor's Coursera lectures on algorithms) for proofs of these facts.

(a) [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.

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

(c) (4 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.

(d) (7 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 $1/\epsilon$.

[Hint: Under what conditions does taking the better of two monotone allocation rules yield another monotone allocation rule?]

(e) (5 points) Suppose, for this part of the problem only, that we have two knapsacks with known capacities. An obvious approach is to use the 2-approximation algorithm discussed in lecture 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 an explicit counterexample.

(f) (6 points) Suppose, for this part of the problem only, that each bidder i has a private valuation v_i and also a *private* size c_i. The first step of a sealed-bid mechanism thus accepts a reported valuation b_i and a reported capacity a_i from each bidder i. The second step of a mechanism decides how much capacity y_i to award to each bidder i, subject to the constraint that $\sum_{i=1}^{n} y_i \leq C$. The third step of a mechanism charges a price p_i to each bidder i. The utility of a bidder i in an outcome is defined as $v_i x_i - p_i$, where x_i is 1 if the bidder gets its required capacity ($y_i \geq c_i$) and 0 otherwise (if $y_i < c_i$).

The mechanism from class extends naturally to this setting of private sizes. Specifically, the mechanism uses the 2-approximation algorithm (with the reported data $\mathbf{b}$ and $\mathbf{a}$) to select a subset S of winning bidders whose reported sizes fit in the knapsack (i.e., $\sum_{i \in S} a_i \leq C$). Second, the mechanism awards capacity a_i to each winner $i \in S$ and capacity 0 to each losing bidder. Third, the mechanism charges payments as if the reported sizes were known a priori: that is, the payment of a winning bidder $i \in S$ is the minimum bid b'_i at which it would continue to win (holding its reported size a_i fixed).

Is this extended mechanism DSIC? Prove it or give an explicit counterexample.

Problem 5

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. If bidder i receives the goods $A_i \subseteq M$ at a total price of p, then its utility is $v_i x_i - p$, where x_i is 1 if $A_i \supseteq T_i$ and 0 otherwise.

(a) (5 points) A subset S of bidders is feasible if every good of M is sought by at most one bidder of S — that is, if $T_i \cap T_j = \emptyset$ for each distinct $i, j \in S$. Consider the problem of computing the feasible subset of bidders that maximizes the social surplus $\sum_{i=1}^{n} v_i x_i$. Prove that this problem is NP-hard. Specifically, prove that computing a maximum-weight Independent Set of a graph with maximum degree d reduces to the problem of maximizing the social surplus of bidders who each desire a set T_i with cardinality at most d.[1]

(b) (5 points) Here is a natural greedy algorithm for the above social surplus maximization problem, given a reported bid from each player:

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

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

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

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

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

(c) (5 points) Prove that if all bidders report truthfully and have sets T_i of cardinality at most d, then the outcome of the mechanism in (b) has social surplus at least $\frac{1}{d}$ times that of the maximum possible.

Problem 6

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).

[1]The Independent Set problem is discussed in every textbook that covers NP-completeness.

(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.]

Extra Credit

(Up to 20 points, redeemable by December 13th, 2013) Can you produce a better physical demonstration of Braess's Paradox than those currently on YouTube? Possible dimensions for improvement include the magnitude of the weight's rise, production values, and dramatic content.