CPSC455b: Hour Exam

February 26, 2002

Answer four of the following five questions. Please remember to put your name and email address on the cover(s) of your blue book(s).

Question 1

Let $\{1, 2, \ldots, n\}$ be the set of agents in a mechanism-design problem. Let $T = T^1 \times T^2 \times \ldots \times T^n$ be the type space and $A = A^1 \times A^2 \times \ldots \times A^n$ be the strategy space. Individual type vectors and strategy profiles are denoted $t = (t^1, t^2, \ldots, t^n)$ and $(a^1, a^2, \ldots, a^n)$, respectively. O is the set of feasible outputs, and $u^i(o, t^i)$ is agent i's utility if his type is t^i and the output is o. If utilities are quasilinear, then u^i is of the form $v^i + p^i$, where $v^i(o, t^i)$ is agent i's valuation, and $p^i(a^1, a^2, \ldots, a^n)$ is the payment that the mechanism gives to agent i on input $(a^1, a^2, \ldots, a^n) \in A$. Identify each of the "solution concepts" defined in parts (a)–(c).

(a) (5 points) A strategy profile $(a^1, a^2, \ldots, a^n)$ satisfies this definition if, for every i and every $\overline{a}^i \in A^i$, $\overline{a}^i \neq a^i$,

$$u^i(o, t^i) \geq u^i(\bar{o}, t^i),$$

where o is the output of the mechanism on input $(a^1, \ldots, a^{i-1}, a^i, a^{i+1}, \ldots, a^n)$, and $\bar{o}$ is the output on input $(a^1, \ldots, a^{i-1}, \overline{a}^i, a^{i+1}, \ldots, a^n)$.

(b) (5 points) A strategy profile $(a^1, a^2, \ldots, a^n)$ satisfies the definition if, for all

i, all $s^{-i} \in A^1 \times \ldots \times A^{i-1} \times A^{i+1} \times A^n$, and all $\overline{a}^i \in A^i$, $\overline{a}^i \neq a^i$,

$$u^i(o, t^i) \geq u^i(\overline{o}, t^i),$$

where o is the output of the mechanism on input (a^i, s^{-i}), and $\overline{o}$ is the output on input $(\overline{a}^i, s^{-i})$.

(c) (5 points) Here we assume that there is a "common prior" on the distribution of agents' types, *i.e.*, a probability distribution on T that is known to all of the agents. Let $(a^1(\cdot), a^2(\cdot), \ldots, a^n(\cdot))$ denote a strategy profile, in order to emphasize that the strategy $a^i(t^i)$ played by i is a function of its type t^i. $(a^1(\cdot), a^2(\cdot), \ldots, a^n(\cdot))$ satisfies this definition if, for all i, all t^i, and all $\overline{a}^i \in A^i$, $\overline{a}^i \neq a^i$,

$$\tilde{u}^i(o, t^i) \geq \tilde{u}^i(\overline{o}, t^i)$$

where o and $\overline{o}$ are as in part (a) above, and $\tilde{u}^i$ is the *expected* utility of agent i. The expectation is computed over the common prior distribution on T.

(d) (5 points) Assume, as in part (c), that there is a common prior on agents' types and thus that each of the three notions given above is well defined for a particular mechanism-design problem. True or False: $(b) \succ (c) \succ (a)$. That is, a strategy profile that satisfies (b) also satisfies (c), and one that satisfies (c) also satisfies (a).

(e) (5 points) What is a truthful mechanism M called if it satisfies the following condition: For every type vector $t = (t^1, t^2, \ldots, t^n)$, if $o = M(t)$ and $o' \neq o$, then

$$\exists i \ u^i(o', t^i) > u^i(o, t^i) \Rightarrow \exists j \ u^j(o', t^j) < u^j(o, t^j).$$

Question 2

(a) (7 points) What is a *utilitarian* mechanism-design problem?

(b) (8 points) What is a *VCG mechanism*?

(c) (10 points) Recall the *task-allocation* mechanism-design problem, defined as follows:

- An instance is a set of tasks $Z = \{z_1, \ldots, z_k\}$.
- Agent i's type is $t^i = (t_1^i, \ldots, t_k^i)$, where t_j^i is the minimum time in which i can complete z_j.
- Feasible outputs of the mechanism consist of partitions $Z = Z^1 \sqcup Z^2 \sqcup \ldots \sqcup Z^n$. ($Z^i$ is the set of tasks assigned to agent i.)
- Agent i's valuation function is

$$v^i(Z, t^i) = - \sum_{z_j \in Z^i} t_j^i.$$

- The mechanism-design goal is to compute

$$\min_Z \max_i \sum_{z_j \in Z^i} t_j^i.$$

Nisan and Ronen propose the MinWork mechanism: Assign each task z_j to the agent i that declares the smallest completion time a_j^i, breaking ties arbitrarily. The payment to agent i is

$$\sum_{z_j \in Z^i} \min_{a_j^{i'} > a_j^i} a_j^{i'}.$$

Prove that truth telling is a (weakly) dominant strategy for the MinWork mechanism. (That is, you should prove that truth telling is a dominant strategy, but you need not prove that it is the only dominant strategy.)

Question 3

(a) (15 points) Recall that we studied two formulations of the lowest-cost path (LCP) mechanism-design problem, one used in the (first) Nisan-Ronen paper [NR1] and in the Hershberger-Suri paper [HS] and the other used in the Feigenbaum-Papadimitiou-Sami-Shenker paper [FPSS]. For five points each, give three reasons that the formulation of the problem in [FPSS] is more relevant to Internet routing than the formulation in [NR1, HS].

(b) (10 points) Recall that, in the mechanism given in [FPSS], agent k's type is c_k (its per-packet cost for carrying traffic) and its payment p^k is of the form

$$p^k = \sum_{i,j} T_{i,j} p^k_{ij},$$

where $\{T_{i,j}\}$ is the traffic matrix, and

$$p^k_{ij} = c_k I_k(c;i,j) + \left(\sum_{r \in N} I_r(c|^k_\infty; i,j) c_r - \sum_{r \in N} I_r(c;i,j) c_r \right).$$

Here, $I_k(c;i,j)$ is 1 if k is on the LCP from i to j and is 0 otherwise. If $c = (c_1, \ldots, c_n)$, then $c|^k_\infty = (c_1, \ldots, c_{k-1}, \infty, c_{k+1}, \ldots, c_n)$. For five points each, give two aspects of this (provably unique) payment function that are notable, in view of assumptions (or, more precisely, lack of assumptions) made in the [FPSS] formulation of the LCP mechanism-design problem.

Question 4

Recall that Goldberg, Hartline, Karlin, and Wright [GHKW] study one-round, sealed-bid auctions for *digital goods*, *i.e.*, goods in unlimited supply. They define the *m-optimal, single-price, ominiscient* auction $\mathcal{F}^{(m)}$ as follows:

Let b be a bid vector, and let v_i be the i^{th} largest bid in the vector b. Auction $\mathcal{F}^{(m)}$ on input b determines the value k such that $k \geq m$ and kv_k is maximized. All bidders with $b_i \geq v_k$ win at price v_k; all remaining bidders lose. The profit of $\mathcal{F}^{(m)}$ on input b is thus

$$\mathcal{F}^{(m)}(b) = \max_{m \leq k \leq n} kv_k$$

(a) (8 points) What does it mean to say that a truthful auction $\mathcal{A}$ for goods in unlimited supply is *β-competitive against* $\mathcal{F}^{(m)}$? (That is, what is the definition given in [GHKW]?)

(b) (5 points) What is the k-item Vickrey auction V_k?

(c) (12 points) Prove that, for any number n of bidders, any constant $\beta > 1$, and any $k = f(n)$, where $1 \leq f(n) < n$, the truthful auction $\mathcal{V}_k$ is not β-competitive against $\mathcal{F}^{(k)}$.

Question 5

(a) (5 points) What is the combinatorial-auction design problem (as defined in the Nisan-Ronen paper "Computationally feasible VCG-based mechanisms" [NR2])?

(b) (4 points) Why is there no known polynomial-time mechanism for this problem that always produces optimal allocations?

(c) (8 points) Let $k()$ be a mapping from agents' declarations, *i.e.*, vectors $(w^1(), \ldots, w^n())$ of valuation functions, to allocations. What does it mean to say that $m = (k(), p())$ is a "VCG mechanism based on $k()$"? That is, what is the definition given in [NR2]?

(d) (8 points) Prove that, although a VCG-based mechanism is not necessarily truthful, the only lies that an agent can benefit from telling are ones that improve the allocation computed by the mechanism.