CS880: Topics in Algorithmic Game Theory	Homework 3
Out: 3/21/2011	Due: 4/5/2011

Guidelines

Same as for HW1.

Problems

1. (4 points) Recall that in class we saw an example of Braess's paradox: adding an edge to the network in the selfish routing game can make the Nash equilibrium flow worse in terms of average latency. Give bounds on the extent to which adding edges can worsen the average latency of the equilibrium flow. You may assume that the network has a single source and a single sink, and the latency functions are continuous, differentiable, and convex. Your bound may depend on the types of latency functions in the network.

2. In class we saw a different way of bounding the cost of the Nash equilibrium flow in a selfish routing game: the cost is no more than the cost of sending twice as much flow in the network optimally. This kind of "bicriteria" bound is surprisingly versatile. In this question we will consider two variants.

   (a) (4 points) In our discussion in class we assumed that the equilibrium flow is extremely sensitive to delay: if there is a path with even slightly less delay than the current flow path, the flow tries to shift to that new path. A more reasonable way to model a stable state may be to use a notion of approximate equilibria: a flow f is called an approximate Nash flow if for any two source-sink paths P and Q with $f_P > 0$, $\ell_P(f) \leq (1+\epsilon)\ell_Q(f)$ for some fixed sensitivity parameter ϵ. Approximate Nash flows may be worse than Nash flows in terms of average cost. Show that a form of the bicriteria bound holds for all approximate Nash flows.

   (b) (6 points) Next consider a capacitated version of the selfish routing problem. Each edge is now endowed with two parameters, a capacity u_e, and a scaling factor a_e. The latency function of the edge is given by $\ell_e(x) = a_e/(u_e - x)$. Note that this is a convex function and approaches infinity as x approaches u_e (and therefore, capacities are enforced). Show that the average delay of a Nash flow in such a network is bounded above by the average delay of an optimal flow satisfying the same demand but in a network with only half the original capacity (that is where u_e for each edge e is replaced by $u_e/2$ without changing the a_es).

3. For this question once again consider the (nonatomic) selfish routing game with a single source and a single sink. Assume that latency functions on edges are nonnegative, linear, increasing functions, that is, functions of the form $\ell_e(x) = a_e x + b_e$. A different measure for the quality of routing in a network may be the maximum latency that a flow faces:

$$\text{maxcost}(f) = \max_{P: f_P > 0} \sum_{e \in P} \ell_e(f_e)$$

   The price of anarchy for this cost measure is then defined in the same way as before: the ratio of the maxcost of the worst Nash flow to the maxcost of the optimal flow.

Note that while our quality measure has changed, the utility function of any user is still the same – minimizing its own latency. Therefore, Nash flows under the original average cost measure are still Nash flows here, but optimal flows may be different. Also (convince yourself that), like Nash flows, there always exist optimal flows under the maxcost function that are *fair flows* – any two flow-carrying paths have equal end-to-end latency. (Note however that a fair flow is not necessarily a Nash flow because non-flow-carrying paths may have lower latency. Also note that the optimal flow may not be unique, and not all optimal flows are fair.)

(a) (3 points) Consider a single-source single-sink network with multiple parallel links between the source and the sink (as in Pigou's example). Prove that the price of anarchy with respect to maxcost in such a network is 1.

(b) (3 points) Prove that in general single-source single-sink networks, the price of anarchy can be larger than 1. (You only need to give an example where PoA is > 1.)

(c) (5 points) Prove that in general single-source single-sink networks, the price of anarchy cannot be larger than $4/3$. (Recall that all latency functions are linear and use the characterizations of Nash flows we came up with earlier.)

(d) (5 points) Let f^* be the optimal flow in the network with respect to the average cost measure. Note that f^* is not necessarily a fair flow. Prove that the "unfairness" of f^*, formally the ratio of the latency of the "longest" flow-carrying path in f^* to that of the "shortest" flow-carrying path, is at most 2. That is, the disparity between the latencies of any two flow-carrying paths in f^* is at most a factor of 2.