CO759: Algorithmic Game Theory — Fall 2010

Instructor: Chaitanya Swamy

Assignment 2

Due: By Dec 7, 2010

You may use anything proved in class directly. I will maintain a FAQ about the assignment on the course webpage. *Acknowledge all collaborators and any external sources of help or reference.* All questions carry equal weightage.

Q1: Recall the load balancing game from class: there are m machines and n jobs. Each job j is a player: it has a certain processing time p_j and its strategy is to choose a machine to get assigned to. Given an assignment of jobs to machines, the load L_i of machine i is the total processing time of the jobs assigned to it, and the cost incurred by a job j is the load of the machine to which it is assigned. In this question, we address the issue of how quickly the natural heuristic of letting players make improving-moves converges to a Nash equilibrium. (You might want to read the proof of Theorem 20.6 in the book before attempting this question.)

Given an assignment $\{i(j)\}$, we say that a job j is *unsatisfied* if it can move to some machine (other than $i(j)$) and reduce its cost. A *best-response move* of job j is a move that minimizes its cost (given the current assignment).

(a) Consider the round-robin policy, where jobs are considered in some arbitrary order and if a job is unsatisfied it makes a best-response move. Prove that this policy leads to a (pure) Nash equilibrium in at most n^2 steps. (Each iteration counts as a step regardless of whether the job considered is unsatisfied or not.)

(b) Now consider the following randomized policy: in each step, pick one of the n jobs at random and if it is unsatisfied make a best-response move. Prove that with probability at least $1 - \frac{1}{n}$, this policy leads to a Nash equilibrium in $O(n^2 \log n)$ steps.

Q2: Consider a nonatomic routing game instance $\mathcal{I} = \big(G, \{\ell_e(.)\}, \{(s_i, t_i, r_i)\}_{i=1}^k\big)$ where G is the underlying directed graph, the $\ell_e(.)$'s are the continuous nondecreasing latency functions on the edges, and r_i units of flow have to be routed from the source s_i to sink t_i for each $i = 1, \ldots, k$. Let $\mathcal{P}_i$ denote the set of all s_i-t_i paths. Recall that $C(f) = \sum_e f_e \ell_e(f_e) = \sum_i \sum_{P \in \mathcal{P}_i} f_P \ell_P(f)$ denotes the total cost of a feasible flow f. In class, we proved that $C(f^{NE}) \leq C(f^*)$, where f^{NE} is a Nash flow for $\mathcal{I}$ and f^* is an optimal flow for the instance $2\mathcal{I} = \big(G, \{\ell_e(.)\}, \{(s_i, t_i, 2r_i)\}\big)$.

Given a class $\mathcal{L}$ of latency functions, recall that $\alpha(\mathcal{L}) = \sup_{\ell \in \mathcal{L}} \alpha(\ell)$ and $\alpha(\ell) = \sup_{a,b \geq 0} \frac{b \cdot \ell(b)}{a \cdot \ell(a) + (b-a) \cdot \ell(b)}$. Define $\rho = \rho(\mathcal{L}) = 2 - \frac{1}{\alpha(\mathcal{L})} < 2$. Show that if all the latency functions belong to class $\mathcal{L}$, then the above bound can be improved to $C(f^{NE}) \leq C(\tilde{f})$, where $\tilde{f}$ is an optimal flow for the instance $\rho\mathcal{I} = \big(G, \{\ell_e(.)\}, \{(s_i, t_i, \rho r_i)\}\big)$.

Q3: In class, we restricted our attention to pure Nash equilibria (NE) and defined the price of anarchy (PoA) as the maximum over all pure NE of the ratio of the cost of the equilibrium to that of the optimum. But one can define PoA (and price of stability) for any equilibrium concept. In

particular, define the *mixed PoA* of a game as the maximum, *over all mixed NE* of the game, of the ratio of the expected cost of the mixed equilibrium and the cost of the optimum. In this question, we investigate the mixed PoA of atomic routing games.

(a) Prove that for linear latency functions, the mixed PoA is at most 5/2.

(b) Prove that the mixed PoA with latency functions that are polynomials of degree p with non-negative coefficients is at most $p^{O(p)}$.

Q4: In this question, we consider the PoA (of pure NE) of atomic and nonatomic routing games with respect to the *maximum player-cost objective* that we looked into in the load balancing game. We are given a directed graph G with nondecreasing latency functions $\{\ell_e(.)\}$ on the edges. We will restrict our attention to instances where all traffic has to be routed from a common source s to a common sink t. Given an s-t flow $f = (f_P)$ that sends flow $f_P \geq 0$ on the s-t path P, let $L(f) = \max_{P:f_P>0} \ell_P(f)$ (where $\ell_P(f) = \sum_{e\in P} \ell_e(f_e)$) denote the cost of the flow according to the maximum-cost objective.

(a) Consider the atomic routing game where there are k players, each having source s and sink t. Let $\{P_i\}_{i=1}^k$ be a (arbitrary) Nash equilibrium with associated flow-vector f, so that $L(f) = \max_i \ell_{P_i}(f)$. Show that the PoA for the maximum-cost objective $L(.)$ with linear latency functions is at most $\frac{5}{2}$.

(b) Now consider the nonatomic game where k units of flow have to be sent from s to t and any fractional amount of flow can be routed along any s-t path. Let α be the PoA for the *total-cost objective* $C(f) = \sum_P f_P \ell_P(f) = \sum_e f_e \ell_e(f_e)$ with the given latency functions. Show that the PoA for the maximum-cost objective $L(.)$ is at most α.

Q5: Recall that for the atomic routing game, we used the (exact) potential function, $\Phi(f) = \sum_e \sum_{x=1}^{f_e} \ell_e(x)$ to prove the existence of a pure NE. Here we show how the potential function can also be useful in finding an *approximate* Nash equilibrium quickly.

Consider the atomic routing game on a directed graph G with nondecreasing latency functions $\{\ell_e(.)\}$, and k players, all having a common source s and common sink t. Consider the total-cost objective $C(f) = \sum_e f_e \ell_e(f_e) = \sum_i \ell_{P_i}(f)$, so Φ satisfies $\Phi(f) \leq C(f)$ for all f. Also, let $\alpha \geq 1$ be such that $\ell_e(x+1) \leq \alpha \cdot \ell_e(x)$ for all e and all $x \in \{1, \ldots, k-1\}$. Let $\mathcal{P}$ denote the set of all s-t paths. Given a flow-vector f corresponding to a path-selection $\{P_i\}_{i=1}^k$, where each P_i is an s-t path, define the *neighborhood* of f to be

$$N(f) \;:=\; \big\{g : \exists i \in \{1, \ldots, k\} \text{ and } Q \in \mathcal{P} \text{ s.t. } g \text{ is the flow-vector associated with } (Q, P_{-i})\big\}.$$

Say that $\big(\{P_i\}, f\big)$ is an *ϵ-approximate local optimum* of Φ, if for all $g \in N(f)$, we have $\Phi(f) - \Phi(g) \leq \epsilon \cdot \Phi(f)$. Define $\big(\{P_i\}, f\big)$ to be an *ϵ-Nash equilibrium* if for all i, for all $Q \in \mathcal{P}$, we have $\ell_{P_i}(f) - \ell_Q(g) \leq \epsilon \cdot \ell_{P_i}(f)$, where $g \in N(f)$ is the flow-vector that results when i deviates from P_i to Q.

(a) Show that if $\big(\{P_i\}, f\big)$ is an ϵ-approximate local optimum of Φ, then $\big(\{P_i\}, f\big)$ is a δ-NE, where $\delta = \frac{\epsilon k \alpha}{1 - k\epsilon}$.

(**Hint:** Try to lower bound $\ell_{P_i}(f)$ in terms of $\Phi(f)$ for every $i = 1, \ldots, k$.)

(b) Prove the following "converse" of part (b). Given $\big(\{P_i\}, f\big)$, let player i^* and $Q^* \in \mathcal{P}$ be such that among all the possible moves where a single player switches to an alternative strategy, the move where i^* switches over to Q^* yields the maximum reduction in the deviating player's cost. In other words, letting g^* denote the flow-vector in $N(f)$ obtained when i^* switches over to Q^*, we have $\ell_{P_{i^*}}(f) - \ell_{Q^*}(g^*) = \Phi(f) - \Phi(g^*) = \max_{g \in N(f)} \big(\Phi(f) - \Phi(g)\big)$. Show that if $\ell_{P_{i^*}}(f) - \ell_{Q^*}(g^*) \geq \epsilon \cdot \ell_{P_{i^*}}(f)$ then $\Phi(f) - \Phi(g^*) \geq \frac{\epsilon}{k(\alpha+1)} \cdot \Phi(f)$.

Part (a) shows that *any* algorithm for computing an approximate local-optimum of Φ can be used to compute an approximate Nash equilibrium. Using part (b), one can argue that an approximate local optimum may be computed by choosing improving moves suitably. Together, parts (a) and (b) yield a simple, efficient, "improving-moves" algorithm for computing a δ-NE. Suppose that $\Phi(f) \geq 1$ for all flow-vectors f resulting from a valid path-selection. Let C be some trivial upper bound on the maximum cost incurred by a player in a NE. For example, one can take $C = \min_{P \in \mathcal{P}} \sum_{e \in P} \ell_e(k)$. Assume that $\delta \leq 1$ without loss of generality. We set $\epsilon = \frac{\delta}{2k\alpha}$, so $\frac{\epsilon k \alpha}{1 - k\epsilon} \leq \delta$, and $\varepsilon = \frac{\epsilon}{k(\alpha+1)}$. We start from any state with $\Phi(.)$ value at most kC. At each step, given the current flow-vector f, letting i^*, Q^*, g^* be as defined in part (b), we make the improving move where i^* switches to path Q^* if the reduction $\ell_{P_{i^*}}(f) - \ell_{Q^*}(g^*)$ is at least $\epsilon \cdot \ell_{P_{i^*}}(f)$. We terminate when this no longer holds. Part (b) shows that each such improving move decreases the potential by at least a $(1 - \varepsilon)$-factor. So since the final potential is at least 1, the number of steps to termination is at most $\ln(kC)/\ln\big((1-\varepsilon)^{-1}\big) \leq \frac{8k^2\alpha^2}{\delta} \cdot \ln(kC)$. At termination, we have

$$\max_{g \in N(f)} \big(\Phi(f) - \Phi(g)\big) = \ell_{P_{i^*}}(f) - \ell_{Q^*}(g^*) \leq \epsilon \cdot \ell_{P_{i^*}}(f) \leq \epsilon \cdot \Phi(f).$$

So f is an ϵ-approximate local optimum of Φ, and part (a) shows that hence, f is a δ-NE.