Algorithmic Game Theory	**Problem Set 1**
CS 6840 Spring 2012	**Due Wednesday, February 15th**

There are 6 questions on this problem set of varying difficulty. For full credit you should solve 5 of the 6 problems. Solving all 6 results in extra credit. A full solution for each problem includes proving that your answer is correct. If you cannot solve a problem, write down how far you got, and why are you stuck.

You may work in pairs and hand in a shared homework with both of your names marked. You may discuss homework questions with other students, but closely collaborate only with your partner. Send me or the TAs email if you have trouble finding a partner, and we'll see if we can help. You may use any fact we proved in class without proving the proof or reference, and may read the relevant chapters of the book. However, you may **not use other published papers, or the Web to find your answer**.

Solutions can be submitted on CMS in pdf format (only), or handed in in class. If you have a partner, write both names on the solution, but only upload or submit it once. In any case, please type your solution to make it easier to read. If your solution is complex, say more than about half a page, please include a 3-line summary to help us understand the argument.

We will maintain a FAQ for the problem set on the course Web page.

(1) Consider the non-atomic multicommodity flow problem we have been discussing in class (defined on Friday, January 27). We defined Nash equilibria as a flow f with congestion x on the edges such that for all pairs of paths P and Q connecting the same pair of terminals, if $f_P > 0$ then $\sum_{e\in P} d_e(x_e) \le \sum_{e\in Q} d_e(x_e)$.

A maybe more intuitive definition would be as follows. A flow f is a Nash equilibrium, if the following holds. For any pairs of paths P and Q connecting the same pair of terminals, such that $f_P > 0$ and any $0 < \delta \le f_P$ if we define an alternate flow $\hat{f}$ by setting

$$\hat{f}_R = \begin{cases} f_P - \delta & \text{if } R = P \\ f_Q + \delta & \text{if } R = Q \\ f_R & \text{otherwise} \end{cases}$$

Let x and $\hat{x}$ denote the congestion of the flows f and $\hat{f}$ respectively. The alternate definition states that the flow f is at equilibrium if $\sum_{e\in P} d_e(x_e) \le \sum_{e\in Q} d_e(\hat{x}_e)$ for all choices of P, Q and δ as above.

This definition considers the flow f and a very small δ amount of flow on a path P, and wonders if this small amount of flow is happier on path P or should it switch to another path Q. Here we model the flow being non-atomic by allowing arbitrary small amounts of flow to switch, but we do not allow "zero" amount to switch, as it is less clear what that means.

Show that the two definitions are the same if the delay functions $d_e(x)$ are monotone increasing (nondecreasing) and continuous. Is this also true for functions $d(x)$ that are not necessarily monotone?

(2) We considered in class the notion of (λ, μ)-smooth delay functions. A class of functions is (λ, μ)-smooth if for all delays $d(x)$ in this class, and any two congestions x^* and x we have that

$$x^* d(x) \le \lambda x^* d(x^*) + \mu x d(x).$$

(a) Show that the class of nonnegative monotone increasing functions is $(1,1)$-smooth.

(b) We have seen in class (on Monday, January 30) that (λ, μ)-smooth for a $\mu < 1$ implies that the corresponding nonatomic congestion game has price of anarchy at most $\lambda/(1-\mu)$. Unfortunately the general bound in part (a) has $\mu = 1$. Show that the above class cannot be (λ, μ)-smooth for any constants λ and $\mu < 1$ by showing that the price of anarchy can be arbitrarily high. (Hint: enough to consider routing games on a network with two parallel edges, one with $d_e(x) = 1$. By setting the delay of the other edge appropriately, you can achieve arbitrarily high price of anarchy.

(c) Show that the following holds using (a). For any equilibrium flow f that with rate r_i for user type i and any flow g that satisfies rates $(1+\delta)r_i$ for each player type i, we can bound the total delay of f in terms of the delay in g. More formally, if x denotes the congestion of flow f and y denotes the congestion of g, prove a bound of the form

$$\sum_e x_e d_e(x_e) \leq F(\delta) \sum_e y_e d_e(y_e)$$

for some function $F(.)$ that is defined for all $\delta > 0$ (but can approach infinity as δ goes to 0).

(d) The following class of delay functions are often used to model capacities. Assume each edge e has two parameters a_e and u_e, and let $d_e(x) = \frac{a_e}{u_e - x}$. Note that this function has $d_e(0) = a_e$, and it models an edge with capacity u_e, as delay goes to infinity as the congestion approaches u_e. Now consider your bound from (c) for this class of functions. Instead of the flow g, we will now consider a flow $\hat{g}$ that has the same flow rate as f, but is running in a network with a scaled down capacity $\hat{u}_e = \frac{u_e}{1+\delta}$ for each edge e. Show a bound analogous to your bound in (c) comparing the cost of f and $\hat{g}$: an equilibrium flow, and a flow in a network with somewhat smaller capacities.

(e) Explain how the bound in (d) can be interpreted to suggest a choice between two options when delays are high in a network: (i) increase bandwidth or capacity of the edges or (ii) improve routing of the flow.

(3) Consider the load balancing game: there are n jobs, each controlled by a separate and selfish user. There are m servers S that can serve jobs, and each job j has an associated set $S_j \subseteq S$ of servers where it can possibly be served. Each server i has a load dependent response time: $r_i(x)$ is the response time of server i if its load is x. We assume that $r_i(x)$ is a monotone increasing function for all i. You may also assume that $r_i(x)$ is convex if that helps. Please note in your answer what assumption you are using and where (r monotone or convex?).

Assume first that each job has a equal load of 1, so the load on a machine is the number of jobs it serves. This is a congestion game, and hence Nash equilibria are local optima of the corresponding potential function Φ.

Hint: For this problem, you may use the fact that the minimum cost matching problem can be solved in polynomial time. This may be useful as a subroutine. **The minimum cost matching problem** is given by a bipartite graph G, costs on the edges and an integer k, and the problem is to find a matching in G of size k of minimum possible cost. You can see any book on algorithms or combinatorial optimization for algorithms for this problem.

(a) Give a polynomial time algorithm to find an equilibrium.

(b) Consider the assignment of jobs to servers that minimizes social welfare, the sum of all response times **over jobs** (or average response time), and give a polynomial time algorithm to find the best assignment for this objective function.

(c) Consider the special case when all jobs can be processed on any machine. Are all pure equilibria are socially optimal in this case?

(d) A mixed Nash equilibrium is a probability distribution for all players on strategies, so that no player can improve his or her expected cost by changing strategy. In this part consider the special case when all jobs can be processed on any machine and all machines have identical delay functions $r_i(x) = x$. An example of a mixed Nash equilibrium is when all players choose between the machines uniformly at random. In such a mixed Nash all players' expected cost is $1 + \frac{n-1}{m}$, and this example shows that the price of anarchy is at least $2 - 1/m$. (See example 17.4 in the book.) Show that this example is worst possible, and the price of anarchy is at most $2-1/m$, that is, the sum of the expected cost of the players in a mixed Nash equilibrium is at most $2 - 1/m$ times the minimum possible such cost.

(4) Consider a variant of the previous question where we assume each job j has a different load p_j, and the load on a server is the sum of the loads $\sum p_j$ over the jobs assigned to this particular server. This is no longer a potential game in the sense we defined in class, as congestion is no longer a function of just the number of jobs.

(a) Show that there is no potential Φ function with the property that for any pure strategy, and any single player changing strategy, the change in the potential function is same as the change in the players's cost who is changing strategy.

(b) Show that this game has a pure Nash equilibrium. Is the equilibrium necessarily unique?

(c) Is a sequence of best responses guaranteed to find a Nash equilibrium, or can a best response sequence cycle?

(5) Consider a variant of the question (3) where we assume each job j experiences congestion differently. Assume $r_i^j(x)$ is the cost experienced by player j on server i if its load is x, where x is the number of players that selected server i. This game models situations where congestion effects different players differently (depending on what type of jobs they have).

(a) Is a sequence of best responses guaranteed to find a Nash equilibrium, or can a best response sequence cycle?

(b) Prove that the game has a pure Nash equilibrium.

(c) Prove that the best response converges if there are only two machines.

(6) Consider the following grouping game. Players are nodes of a graph, and edges represent social relations. We assumed that edge (i, j) connecting nodes i and j has an associated value $u_{ij} \geq 0$. We will think of the value u_{ij} as the strength of their friendship, the benefit players i and j both incur if nodes i and j are together. In the game each player can choose between affiliating with one of two social groups A or B. The result is a partition of the nodes into two sets. The utility of a player $i \in A$ is $U_i(A, B) = \sum_{j \in A} u_{ij}$, and if player i is in B its utility is $U_i(A, B) = \sum_{k \in B} u_{ik}$, i.e., the some of the utilities for partners in the same group. Assume that the values are symmetric, namely that $u_{ij} = u_{ji}$ for all i and j.

(a) Show that this is a potential game with the social welfare $\frac{1}{2}\sum_i U_i(A, B)$ is a potential function.

(b) Show that the Price of Anarchy is at most 2, i.e., that any (pure strategy) Nash equilibrium has social welfare at least 1/2 of the maximum possible.

(c) It turns out that all potential games are congestion games with the right definition of congestion. Define a congestion game that is equivalent to this partition game.

(d) How does the game change if we do not assume symmetry? that is if u_{ij} may not equal u_{ji}? Is the resulting game still a potential game? Does it have a bound of 2 on the price of anarchy? is it still true that any sequence of best responses converges to a Nash equilibrium?