CS 6840 – Algorithmic Game Theory (5 pages)	Spring 2012
Lecture 36 Scribe Notes: Fair sharing and networks	
Instructor: Eva Tardos	*Scribe: Jean Rougé (jer329)*

NOTE: This lecture is based on a paper by R. Johari and J.N. Tsitsiklis published in 2004 that you can find on the course's website.

1 Review of previous lectures

We've investigated fair sharing on a single link during the last two lectures. Let us recall the setting:

* n users compete for bandwidth on a single link of total capacity B
* each user i has his own utility function $U_i(x)$ for x amount of bandwidth
* for all i, we assume that U_i is concave, continously differentiable, monotone increasing, and positive
* user i pays an amount w_i to get an amount x_i of bandwidth, and receives net utility $U_i(x_i) - w_i$

We've seen two disctinct ways of allocating the resource :

1. **Price equilibirum:** we post a fixed price p per amount and we let every user choose his own amount x_i by solving the optmization problem

$$x_i = \arg\max_x (U_i(x) - px)$$

then we showed that this is an equilibrium if either $p = 0$ and $\sum_i x_i \leq B$, or $p > 0$ and $\sum_i x_i = B$; and we also showed that the solution for price equilibirum is socially optimum, i.e. that it maximizes $\sum_i U_i(x_i)$.

2. **Fair sharing as a game:** now each user offers an amount of money w_i, and as a result gets his *fair share*

$$x_i = \frac{x_i}{\sum_j w_j} B$$

We proved that any Nash equilibrium for this game satisfies that for every user i, either

$$U_i'(0) \leq p \text{ and } w_i = 0$$

or

$$U_i'(x_i)\left(1 - \frac{x_i}{B}\right) = p$$

where $p = \frac{\sum_i w_i}{B}$ is the "implicit" price at which the bandwidth gets sold.
Finally we've also seen that the price of anarchy in that game is bounded by $\frac{4}{3}$.

2 And now to networks

Now we want to extend these results to a network comprising a number of links.

Let us define this new setting:

* we consider a graph $G = (V, E)$, where each edge $e \in E$ has a bandwidth capacity $b_e \geq 0$
* user i wants to use links along a fixed path P_i in G
* each user i offers an amount of money $w_{i,e}$ for every edge $e \in P_i$ along his path
* and as a result, player i gets allocated an amount $x_{i,e}$ for each edge edge $e \in P_i$; and he actually enjoys bandwidth[1]
$$x_i = \min_{e \in P_i} x_{i,e}$$
* thus the net utility for user i is
$$U_i(x_i) - \sum_{e \in P_i} w_{i,e}$$

To be able to say something on the Nash equilibria in this game, we'll need to have a result comparable to the price equilibrium theorem we proved for the single-link setting, so that we'll be able to compare a Nash equilibrium to the price equilibrium.

2.1 Price equilibirum theorem

Like we did last time for a single link, let us define a price p_e for every edge $e \in E$; then each user i maximizes his utility by solving the optimization problem

$$x_i = \arg\max_x \left(U_i(x) - x \sum_{e \in P_i} p_e \right)$$

Definition. $(p_e)_{e \in E}$ defines a *price equilibrium* if for every edge $e \in E$, either

$$\sum_{i | e \in P_i} x_i = b_e$$

or

$$\sum_{i | e \in P_i} x_i \leq b_e \text{ and } p_e = 0$$

And it turns out that we get the same result as in the single-link case:

Theorem 1. A price equilibrium is socially optimal, i.e. it maximizes $\sum_i U_i(x_i)$.

[1] The paper by Johari and Tsitsiklis mentioned above is more general than this, in particular the following analysis can be extended to other definitions of x_i as a function of $(x_{i,e})_e$; what's important is that the utility depends on a global variable which in turn depends on the local allocations of the elemtary resources.

Proof. Let us consider a price equilibrium $(x_1, x_2, \ldots, x_n)$, and let us compare it to a socially optimum allocation $(x_1^*, x_2^*, \ldots, x_n^*)$. By definition of the x_i, for all user i,

$$U_i(x_i) - x_i \sum_{e \in P_i} p_e \geq U_i(x_i^*) - x_i^* \sum_{e \in P_i} p_e$$

Now by summing over all users

$$\sum_i U_i(x_i) \geq U_i(x_i^*) + \sum_i x_i \sum_{e \in P_i} p_e - \sum_i x_i^* \sum_{e \in P_i} p_e \tag{1}$$

Yet by reversing the order of summation

$$\sum_i x_i \sum_{e \in P_i} p_e = \sum_e p_e \sum_{i | e \in P_i} x_i = \sum_e p_e b_e$$

where the last equality comes from the definition of a price equilibirum : either $\sum_{i|e \in P_i} x_i = b_e$ or $p_e = 0$, hence in either case $p_e \sum_{i|e \in P_i} x_i = p_e b_e$. Besides, by reversing the order of summation again,

$$\sum_i x_i^* \sum_{e \in P_i} p_e = \sum_e p_e \sum_{i | e \in P_i} x_i^* \leq \sum_e p_e b_e$$

where the last inequality comes from the fact that we can't allow users to exceed an edge's capacity. Hence,

$$\sum_i x_i \sum_{e \in P_i} p_e - \sum_i x_i^* \sum_{e \in P_i} p_e \geq 0$$

and so (1) becomes

$$U_i(x_i) \geq U_i(x_i^*)$$

□

Note that this theorem does not state that a price equilibrium exists. Using convex optimization, one can prove that price equilibrium exists, when utilities are concave. However, we did not prove this here.

2.2 Network sharing as a game

Now let's get back to considering the game outlined at the beginning of this section : each user i offers an amount of money $w_{i,e}$ for every edge $e \in P_i$ along his path, and as a result gets allocated the amount of bandwidth $x_{i,e}$ according to fair-sharing:

$$x_{i,e} = \frac{w_{i,e}}{\sum_j w_{j,e}} b_e$$

Then the actual bandwitdh he actually enjoys is the minimum along his path, i.e.

$$x_i = \min_{e \in P_i} x_{i,e}$$

which results in the net utility

$$U_i(x_i) - \sum_{e \in P_i} w_{i,e}$$

Unfortunately, this natural definition for this game cannot have an equilibrium : indeed, if an user i is the only one competing for a given edge e, then any offer $w_{i,e} > 0$ will ensure that he'll have the whole link for himself alone; but if he offers $w_{i,e} = 0$ he won't get anything.

So we're going to consider a slightly modified version of the game to account for this:

1. each user i, for any edge along his path, either makes an offer $w_{i,e} > 0$ or asks for a free bandwidth $f_{i,e}$ over that edge
2. now for any edge $e \in E$:
   * if anyone offered money for e, we share e according to the fair-share rule
   * if no one offered money for e and if we can accomodate all the requests, i.e. $\sum_{i|e\in P_i} f_{i,e} \leq b_e$ then we give away the bandwidth for free, i.e. $x_{i,e} = f_{i,e}$
   * if no one offered money for e but we can't accomodate all the requests, i.e. $\sum_{i|e\in P_i} f_{i,e} > b_e$, then nobody gets anything, i.e. $x_{i,e} = 0$ (the idea being that this is an over-demanded resource, so we're not willing to give it away for free)

We are going to show for this game a similar result to the one we've seen for the single-link setting:

Theorem 2. The price of anarchy in this game is at most $\frac{4}{3}$. That is, if $(x_1, x_2, \ldots, x_n)$ and $(x_1^*, x_2^*, \ldots, x_n^*)$ respectively are the allocation at a Nash equilibrium and a socially optimal allocation, then

$$\sum_i U_i(x_i) \geq \frac{3}{4} \sum_i U_i(x_i^*)$$

The proof of this theorem wasn't completed during that lecture, by lack of time. We are only going to establish a caracterization of Nash equilibria here, and the rest of the proof will be derived in lecture 37 scribe notes.

This characterization of Nash equilibria we're looking for would be an analog of the result we've recalled at the beginning on Nash equilibria for the single-link setting. We had obtained this characterization as the result of a single-variate optimization problem in the player's offer w. Here, by contrast, each player makes a number of offers $w_{i,e}$, and we don't want to try and solve a multi-variate optimization problem. Let us see how we can translate this problem into a single-variate problem.

We only have to notice that at equilibrium, for every user i, and for all edge $e \in P_i$, $x_{i,e} = x_i$ (indeed, user i has no interest in having more bandwidth on one edge than on another, since he only enjoys the minimum of all them).

Then we can express all the $w_{i,e}$ variables as functions of x_i, since by definition

$$x_{i,e} = \frac{w_{i,e}}{\sum_j w_{j,e}} b_e$$

Re-arranging the terms, and using that at equilibrium $x_i = x_{i,e}$, we get that at equilibrium

$$w_{i,e} = x_i \frac{\sum_{j\neq i} w_{j,e}}{b_e - x_i}$$

And now we're down to a single-variate optimization problem: we're looking for

$$x_i = \arg\max_x \left(U_i\left(x\right) - \sum_{e \in P_i} \frac{\sum_{j \neq i} w_{j,e}}{b_e - x} \right)$$

Setting the derivative of the expression above to 0, we get the following characterization: for every user i, at equilibirum, either

$$U_i'\left(0\right) \leq \sum_{e \in P_i} p_e \text{ and } x_i = 0$$

or else

$$U_i'\left(x_i\right) = \sum_{e \in P_i} p_e \frac{1}{1 - \frac{x_i}{b_e}}$$

where p_e is the unit price at which edge e gets sold, namely $p_e = \frac{\sum_{i|e \in P_i} w_{i,e}}{b_e}$.

See lecture 37 scribe notes for the end of this proof.