

Lecture 4: 10.01.03

Lecturer: Christos Papadimitriou

Scribe: Rodrigo Fonseca

Disclaimer: *These notes have not been subjected to the usual scrutiny reserved for formal publications. They may be distributed outside this class only with the permission of the Instructor.*

4.1 Congestion Games: Allocating Bandwidth

We continue the discussion from last time, on Shenker's paper. Recall the setting of the bandwidth allocation, in which we have a switch that is shared among n users. Each user presents to the switch a rate r_i , and is serviced, according to the service discipline \bar{A} , reflected in an expected queue length $q_i = A_i(\bar{r})$. This is represented in the diagram of Figure 4.1.

We observe that the game created in this setting, by setting the payoffs to be equal to the utilities of each user,

$$p_i(\bar{r}) = U_i(r_i, A_i(\bar{r})),$$

actually depends on the A_i chosen.

This leads to the following question, which is central to Mechanism Design: can we create a game, by cleverly choosing the service discipline (and therefore the allocation function), such that the Nash equilibria are of 'high quality'? The following theorem is an impossibility result in this direction, answering this question negatively if we consider all possible sets of utility functions.

Theorem 4.1 *There is no A such that for all U_i , all Nash equilibria are Pareto optimal.*

Sketch of proof:

Any (\bar{r}, \bar{q}) point can be a Nash equilibrium given appropriate U_i 's.

It is always possible to find utility functions that give high enough utility to these specific queue lengths. They must be concave, and there are some subtleties involved to make things work out. This is detailed in Lemma 5, Appendix A, of [?].

For an allocation to be a Nash equilibrium, a necessary (but not sufficient) condition is that for all i , $\frac{dU_i}{dr_i} = 0$,

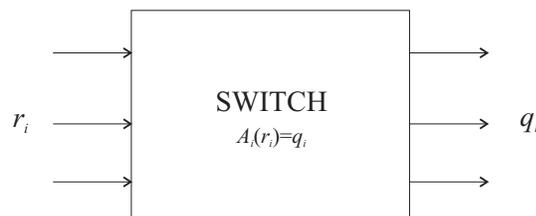


Figure 4.1: The setting for the bandwidth allocation.

or

$$\frac{\partial U_i}{\partial r_i} + \frac{\partial U_i}{\partial q_i} \frac{\partial A_i}{\partial r_i} = 0.$$

On the other hand, if an allocation is Pareto Optimal, then

$$\frac{\partial U_i}{\partial r_i} + \frac{\partial U_i}{\partial q_i} \frac{\partial (\sum q_i)}{\partial r_i} = 0.$$

For a Nash equilibrium to be Pareto optimal, we then must have

$$\frac{\partial A_i}{\partial r_i} = \frac{\partial \left(\frac{\sum r_i}{1 - \sum r_i} \right)}{\partial r_i} = \frac{1}{(1 - \sum r_i)^2}$$

for all r , which is a contradiction, since it states that the way my behavior affects my delay is independent of what I do.

Another important result is the following.

Theorem 4.2 *If for some U_i 's a Nash equilibrium is Pareto optimal, then $r_i = r_j \forall i, j$, and Fair Share could have achieved this.*

These two theorems state that while it is impossible to guarantee that Nash equilibria are Pareto optimal, there are times when this happens. The Fair Share allocation mechanism achieves all such points. The proportional allocation (the allocation that FIFO implements), and any other allocation function that has $\frac{\partial A_i}{\partial r_j} > 0$ (meaning that users are not insulated from each other), never have Pareto optimal Nash equilibria.

4.2 Equilibrium Theory

We now move on, and touch on another subject, that of Market Equilibria and the Price mechanism. Consider the model:

We have

A set of goods, $1..k$

Agents, $1..n$

Agents have endowments $x_i \in \mathcal{R}_+^k$, quantities of goods,
and utilities $U_i : \mathcal{R}^k \rightarrow \mathcal{R}$, strictly concave, that map combinations of goods to real numbers.

In this setting, as agents meet, they will barter items with one another in pairs, to improve their utility. For example, A can give 2 liters of Olive oil in exchange for 1 kg of wheat from B. This process constitutes a potentially exponential search to reach the optimal allocation of goods. However, the Price Mechanism comes to rescue, as the the following theorem states.

Theorem 4.3 *There is a price vector $\bar{p} \in \mathcal{R}_+^k$ such that if $y_i = \operatorname{argmax}(U_i(y) : \bar{p}y \leq \bar{p}x_i)$, then*

$$\sum_i y_i = \sum_i x_i.$$

The proof uses Brouwer's fixpoint theorem, or Kakutani's fixpoint theorem.

This vector attributes a *price* to each good, and the allocation has the property that each agent ends up with the best thing she can afford, and that there is no excess demand and no excess merchandise in the market.

Another theorem states that the final allocation is Pareto optimal. Also, allocation that dominates the initial endowment can be achieved by prices. Using the price mechanism, trade can be realized without resorting to the exponential search process.

A variation to the model is to factor in another component, namely *factories*. Factories are consumers without an initial endowment who have a choice (determined by the presence of raw materials and technology).

Factories have $z_j \in T_j \subseteq \mathcal{R}$, strictly convex.

Then the theorem can be altered to include the factories:

Theorem 4.4 (Factories) *There is a price vector $\bar{p} \in \mathcal{R}_+^k$ such that if $y_i = \operatorname{argmax}\{U_i(y) : \bar{p}y \leq \bar{p}x_i\}$ and $z_j = \operatorname{argmax}\{p z : z \in T_j\}$, then*

$$\sum_i y_i = \sum_i x_i + \sum_j z_j.$$

There are other extensions possible, which we didn't cover in class, such as consumer having shares of the factories.

4.2.1 Price equilibrium in the bandwidth sharing problem

Let us look at the bandwidth allocation problem discussed above as a market equilibrium situation. All users have an utility $U_i(r_i, q_i)$, as represented in the left of figure 4.2. In the Figure, we represent U_i as a function of r_i and $c_i = q_i/r_i$ (the delay).

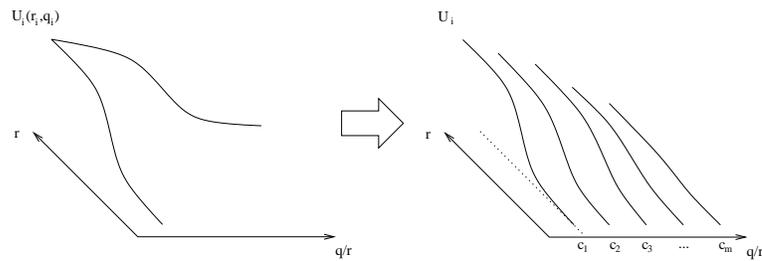


Figure 4.2: Discretization of the utility function

Let us assume that we can discretize U , as in the right hand side of the figure.

$$U_i(r_i^1, r_i^2, \dots, r_i^m) \rightarrow \mathcal{R}_+,$$

where r_i^j is user i 's rate under the queue length $c_j r_i^j$, and these are the goods.

The endowments are

$x_i = (0, 0, \dots, 0, 1)$, i.e., there are no rates, and 1 dollar.

The switch is a factory,

$$r^1 \dots r_m : \sum r_i c_i \geq.$$

By prices, one can then schedule, and achieve Pareto optimality. Is there a price equilibrium in this 'market'? This is still a hard problem, though, and all of the algorithms known have complexity that is exponential in the worst case.