CS 684 Algorithmic Game Theory — Scribe: Georgios Piliouras
Instructor: Éva Tardos — Monday, October 03, 2005

Bandwidth sharing

We will commence our analysis of bandwidth sharing games by introducing a very simple model, where we have only one link. In latter lectures, we will extend our results to more complex networks.

Game Description

We have a link of limited capacity B and k users who want a share of the bandwidth. Each user i has a specific utility function $U_i(x)$, which represents his satisfaction when he receives a bandwidth x. All utility functions $U_i(x)$ are assumed (for convenience of our analysis) to be strictly monotone increasing, continuously differentiable and non-negative.

Lastly, we require that all utility functions are strictly concave, which is going to be necessary for our analysis. This is inherently similar to the *marginal increasing utility* function. That is, given a specific amount of bandwidth, our gain due to a marginal increase to this bandwidth decreases, as the original amount of bandwidth increases.

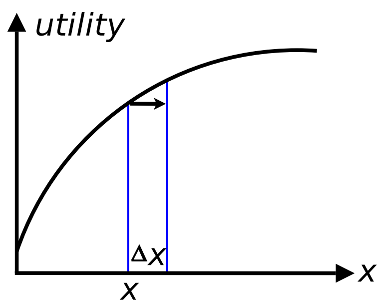

Figure 1: concave utility function

Now, in order to complete the description of the game, we must specify the policy that will determine the distribution of the given bandwidth B over the n players given their utility functions. We will begin by introducing a rather simplistic policy. By analyzing it, we will identify its weaknesses and confront them in order to implement a real-life applicable policy.

Ideal System

In an ideal system, where every agent is truthful and able to compute his utility function, we could achieve the optimal solution, simply by dividing the bandwidth among the players according to their stated utility functions. It is worthy of mentioning that we ignore all delays here. Specifically, our approach assumes that the link is good until we exceed its capacity B, where it crashes.

The optimal solution is

$$\max_x \sum_{i=1}^{k} U_i(x_i)$$

such that

$$x_i \geq 0 \quad \text{and} \quad \sum_{i=1}^{k} x_i \leq B$$

where x_i is the bandwidth given to user i.

However, this model encompasses a lot of implicit assumptions that do not apply to real-life scenarios.

- In our approach we assume that all players are truthful. However, this assumption is not well founded, since every player under this policy would have no reason to announce his true utility function. Clearly, he would announce a severely exaggerated utility function, so as to maximize his share of the bandwidth.
- Another problem of this approach is that every player is assumed to be able to identify his utility function U_i. However, the agent might face problems not only in realizing his true utility function "how should I respond to 1 Tera bps?" but also in describing it due to its large communication complexity.

It becomes immediately evident that we need a mechanism that forces the players to reveal their true utilities and at the same time overcomes the obstacle of having to explicitly declare complex utility functions. We will present such a mechanism below:

Simple mechanism

Our approach will be based on a pricing system, namely we will force every player to pay a specific price p for every unit of the available bandwidth. The price p is common for all players regardless of their utility functions.

Assume that utility is expressed in money (common currency). Given the price, the user i will want to maximize his gain

$$\max_{x \geq 0} (U_i(x) - xp)$$

As we have mentioned earlier we assume the utility function of each player to be strictly monotone, and strictly concave. Hence, we can easily optimize this function for every player, by applying simple calculus techniques. Indeed, for each player i we will achieve maximum at x_i where

$$U_i'(x_i) = p$$

Since U_i is strictly concave, U_i' will be strictly monotone decreasing, so we have a unique solution x_i. We should also add that if $U_i'(0) < p$, then we conclude that $x_i = 0$ (user doesn't gain anything).

Although our primary goal is to ensure that we can satisfy the demands of all players, that is

$$\sum_{i=1}^{k} x_i \leq B$$

at the optimal solution we'll actually have

$$\sum_{i=1}^{k} x_i = B$$

because we assumed that utilities are strictly optimal, and so the optimal solution cannot have leftover bandwidth.

Theorem 1 *If the price p results in utility requests $x_1, \ldots, x_k$ such that*

$$\sum_{i=1}^{k} x_i = B$$

then this is the optimal solution.

Proof: Consider a solution where user i gets y_i. If y_i solves the utility function, then we derive immediately that

$$y_i \geq 0 \quad \text{and} \quad \sum_{j=1}^{k} y_j = B$$

In addition, since U_i is concave we have $U_i(x) \leq U_i'(x_i)(x - x_i) + U_i(x_i)$. Hence, we imply that

$$\sum_{i=1}^{k} U_i(y_i) \quad \leq \quad \sum_{i=1}^{k} (U_i'(x_i)(y_i - x_i) + U_i(x_i)) \tag{1}$$

However,

$$\sum_{i=1}^{k} U_i'(x_i)(y_i - x_i) \leq p \sum_{i=1}^{k} (y_i - x_i) = p(B - B) = 0 \tag{2}$$

because $U_i'(x_i) \leq p$. Moreover, either $U_i'(x_i) = p$ or for all $x \geq 0$, $U_i'(x) < p$. In the latter case, we set $x_i = 0$, so $U_i'(x_i)(y_i - x_i) \leq p(y_i - x_i)$ holds since $U_i'(x_i) < p$ and $y_i - x_i \geq 0$. So, by combining the inequalities 1, 2 we derive that

$$\begin{aligned} \sum_{i=1}^{k} U_i(y_i) \quad &\leq \quad \sum_{i=1}^{k} (U_i'(x_i)(y_i - x_i) + U_i(x_i)) \\ &\leq \quad \sum_{i=1}^{k} U_i(x_i) \end{aligned}$$

Thus, the solution is optimal. □

One rather intuitively evident way to find this price p, which achieves the optimal solution, is to perform sampling over the possible values of p. Although this approach works for our simple one-link case, it is more difficulty to generalize to more complex networks. Hence, we will explore other approaches in finding this optimal price p.

Frank Kelly mechanism

Every user i determines the amount of money w_i he is willing to pay. Afterwards, the mechanism distributes the bandwidth proportionally to this amount. Namely, the price is set to

$$p = \frac{\sum_{i=1}^{k} w_i}{B}$$

and bandwidth received by the i^{th} player to

$$\frac{w_i}{p} = \frac{w_i}{\sum_{j=1}^{k} w_j} B$$

Now, the players use the given price p, to update their responses w_i. During this phase, Kelly's algorithm assumes that the users do not anticipate any change in prices. This model is reasonable for "small" users, that is if $w_i \ll \displaystyle\sum_{j=1}^{k} w_j$. Specifically, players optimize their expected gain

$$\max_{x \geq 0} U_i(x) - xp$$

Each user computes x_i, an optimal solution, and announces $w_i = x_i p$, the money he is willing to spend. Once the new w_i are announced, the algorithm updates the price p and the bandwidth shares according to them. These updates continue until the game reaches an equilibrium. That is, until the updated prices of each players match their previous prices.

Claim 2 *The prices at Nash equilibrium are optimal.*

The optimality of the bandwidth distribution follows immediately from theorem 1.

We conclude by clarifying that although Kelly's algorithm has the advantage that it does not assume knowledge of the individual utility functions, the game defined by it is not a proper game because of the explicit rule that users solve some optimization problem.