Lecture 11

Introduction to Pricing Mechanisms

VCG Mechanism

- Suppose we have a network with a set of traffic sources S and a set of links L.
- Each link $l \in L$ has a finite fixed capacity c_l.
- Each source in S is associated with a route $r \subset L$ along which it transmits at some rate x_r.
- Note that we can use the index r to indicate both a route and the source that sends traffic along that route and we will follow this notation.
- The utility that the source obtains from transmitting data on route r at rate x_r is denoted by $U_r(x_r)$.
- We assume that the utility function is continuously differentiable, non-decreasing and strictly concave.

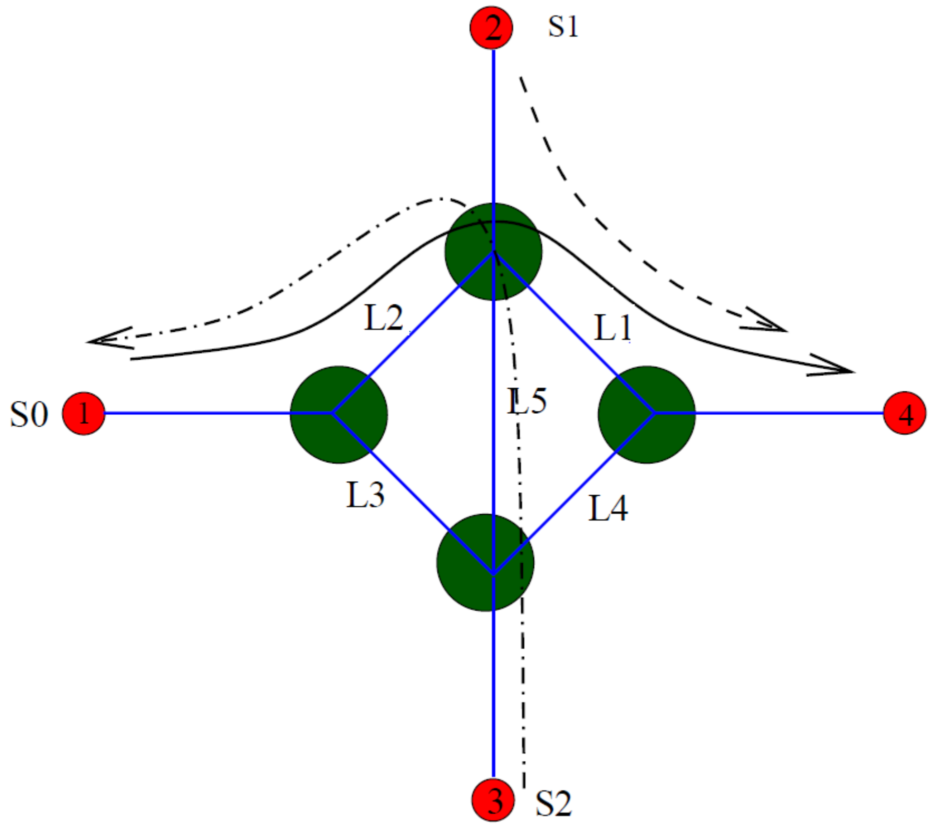

- It is straightforward to write down the problem as an optimization problem of the form:

$$\max_{x_r} \sum_{r \in \mathcal{S}} U_r(x_r)$$

subject to the constraints

$$\sum_{r:l \in r} x_r \leq c_l, \ \forall l \in \mathcal{L},$$

$$x_r \geq 0, \ \forall r \in \mathcal{S}.$$

- The above inequalities state that the capacity constraints of the links cannot be violated and that each source must be allocated a nonnegative rate of transmission.
- It is well known that **a strictly concave function has a unique maximum over a closed and bounded set**.
- In the above problem, the utility function is strictly concave, and the constraint set is closed (since we can have aggregate rate on a link equal to its capacity) and bounded (since the capacity of every link is finite).
- In addition, the constraint set for the utility maximization problem is convex which allows us to use the ***method of Lagrange multipliers*** and the **Karush-Kuhn-Tucker (KKT)** theorem.

VCG Mechanism

Suppose that the network planner asks each user to reveal their utilities and user r reveals its utility function as $\tilde{U}_r(x_r)$, which may or may not be the same as $U_r(x_r)$.

Users may choose to lie about their utility function to get a higher rate than they would get by revealing their true utility function.

the network solves the maximization problem

$$\max_{x \geq 0} \sum_r \tilde{U}_r(x_r)$$

subject to

$$\sum_{r:l \in r} x_r \leq c_l, \forall l$$

and allocates the resulting optimal solution $\tilde{x}_r$ to user r.

What is a proper incentive mechanism?

In return for allocating this rate to user r, the network charges a certain price p_r.

- The price is calculated as follows:
- The network planner calculates the reduction in the sum of the utilities obtained by other users in the network due to the presence of user r, and **collects this amount as the price from** user r.

Specifically, the network planner first obtains the optimal solution $\{\bar{x}_s\}$ to the following problem:

$$\max_{x \geq 0} \sum_{s \neq r} \tilde{U}_s(x_s)$$

subject to

$$\sum_{s \neq r : l \in s} x_s \leq c_l, \forall l.$$

In other words, the network planner first solves the utility maximization problem without including user r.

- The price p_r is then computed as

$$p_r = \sum_{s \neq r} \tilde{U}(\bar{x}_s) - \sum_{s \neq r} \tilde{U}(\tilde{x}_s),$$

which is the difference of sum utilities of all other users without $(\{\bar{x}\})$ and with $(\{\tilde{x}\})$ the presence of user r.

- The network planner announces this mechanism to the users of the network, i.e., it states that

  once the users reveal their utilities, it will allocate resources by solving the utility maximization problem and will charge a price p_r to user r.

- Now the question for the users is the following:

  what utility function should user r announce to maximize its *payoff*?

- The payoff is the utility minus the price:

$$U_r(\tilde{x}_r) - p_r.$$

- We will now see that an optimal strategy for each user is to truthfully reveal its utility function.

- We will show this by proving that announcing a false utility function can only result in reduction in the payoff for user r.

- Suppose user r reveals its utility function truthfully, while the other users may or may not.

- In this case, the payoff for user r is given by

$$\mathcal{U}^t = U_r(\tilde{x}_r^t) - \left(\sum_{s \neq r} \tilde{U}_s(\bar{x}_s^t) - \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^t) \right),$$

- the payoff for user r is given by

$$\mathcal{U}^t = U_r(\tilde{x}_r^t) - \left(\sum_{s \neq r} \tilde{U}_s(\bar{x}_s^t) - \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^t) \right),$$

$\{\tilde{x}_s^t\}$ is the allocation given to the users by the network planner

$\{\bar{x}_s^t\}$ is the solution of the network utility maximization problem when user r is excluded from the network.

- The superscript t indicates that user r has revealed its utility function truthfully.

- Next, suppose that user r lies about its utility function and denote the network planner's allocation by $\tilde{x}^l$.
- The superscript l indicates that user r has lied.
- Now, the payoff for user r is given by

$$\mathcal{U}^l = U_r(\tilde{x}_r^l) - \left(\sum_{s \neq r} \tilde{U}_s(\bar{x}_s^t) - \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^l) \right).$$

- the payoff for user r is given by

$$\mathcal{U}^t = U_r(\tilde{x}_r^t) - \left(\sum_{s \neq r} \tilde{U}_s(\bar{x}_s^t) - \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^t) \right),$$

$$\mathcal{U}^l = U_r(\tilde{x}_r^l) - \left(\sum_{s \neq r} \tilde{U}_s(\bar{x}_s^t) - \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^l) \right).$$

- If truth-telling were not optimal, $\mathcal{U}^l > \mathcal{U}^t$.

  If this were true, by comparing the two expressions for $\mathcal{U}^t$ and $\mathcal{U}^l$, we get

$$U_r(\tilde{x}_r^l) + \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^l) > U_r(\tilde{x}_r^t) + \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^t),$$

$$U_r(\tilde{x}_r^l) + \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^l) \overset{?}{>} U_r(\tilde{x}_r^t) + \sum_{s \neq r} \tilde{U}_s(\tilde{x}_s^t),$$

- contradicts the fact that $\tilde{x}^t$ is the optimal solution to

$$\max_{x \geq 0} U_r(x_r) + \sum_{s \neq r} \tilde{U}_s(x_s)$$

subject to the capacity constraints.

- Thus, truth-telling is optimal under the VCG mechanism.
- Note that truth-telling is optimal for user r independent of the strategies of the other users.
- A strategy which is optimal for a user independent of the strategies of other users is called a *dominant strategy* in game theory.
- Thus, truth-telling is a dominant strategy under the VCG mechanism.

Network Externalities and Distributed Pricing

What is Externality?

Definition (Externality)

An externality is any side effect (benefit or cost) that is imposed by the actions of a player on a third-party not directly involved.

A Case Study: Water Pollution

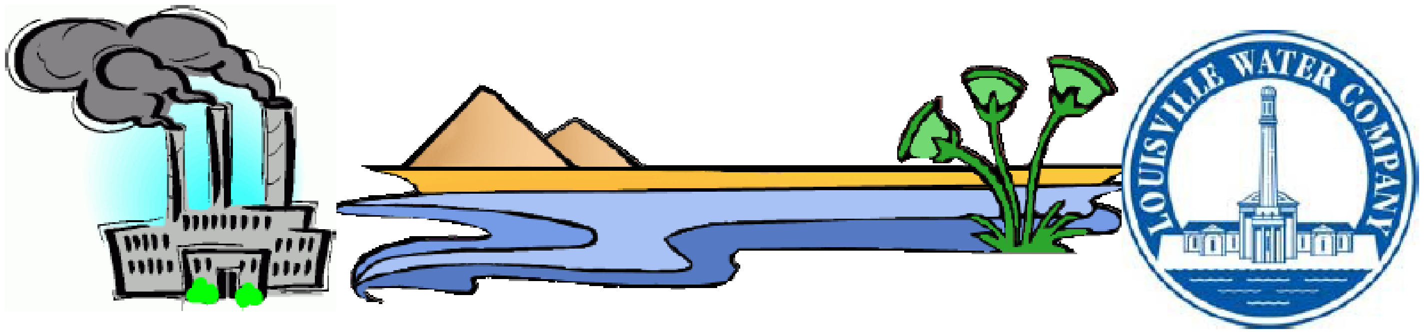

- The chemical company produces chemical products and discharges wastewater into the river.
- The water company produces bottle water by drawing water from the river.
- Water pollution increases the production cost of the water company.

Wireless Power Control

- Single-hop transmissions.
- A user = a transmitter/receiver pair.
- Transmit over one or multiple parallel channels.
- Distributed power control in wireless ad hoc networks
- Elastic applications with no SINR targets
- Want to maximize the total network performance
- Interferences in the same channel (negative externality).

Single Channel Communications

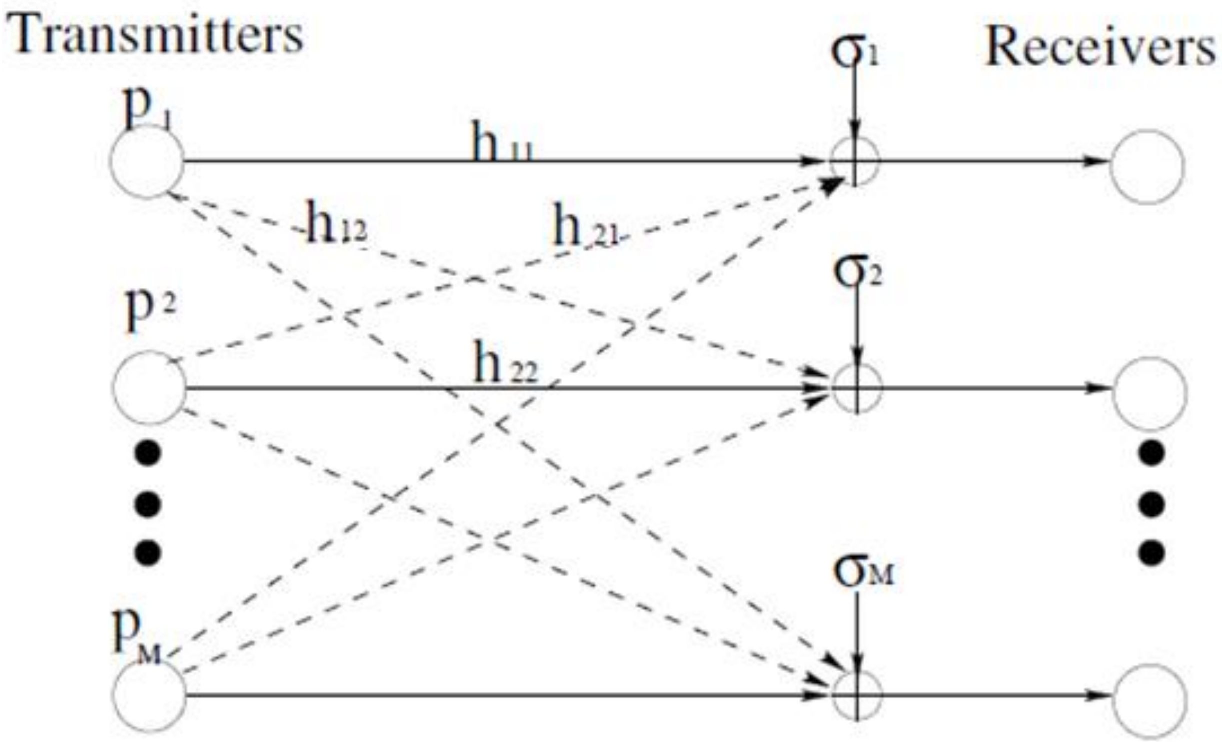

- We focus on a single channel.
- For each user $i \in \{1, ..., M\}$:
  - Power constraint: $p_i \in [P_i^{min}, P_i^{max}]$.
  - Received SINR (signal-to-interference plus noise ratio):

$$\gamma_i = \frac{p_i h_{i,i}}{\sigma_i + \sum_{j \neq i} p_j h_{ji}}.$$

  - Utility function $U_i(\gamma_i)$: increasing, differentiable, strictly concave.

Network Utility Maximization (NUM) Problem

NUM Problem

$$\max_{\{P_i^{\min} \leq p_i \leq P_i^{\max}, \forall n\}} \sum_i U_i(\gamma_i).$$

An example utility function is a *logarithmic utility function* $u_i(\gamma_i) = \theta_i \log(\gamma_i)$, where θ_i is a user-dependent priority parameter.

- Technical Challenges:
  - Coupled across users due to interferences.
  - Could be non-convex in power
- We want: efficient and distributed algorithm, with limited information exchange and fast convergence.

Any local optimum, $\boldsymbol{p}^* = (p_1^*, \ldots, p_M^*)$, of this problem must satisfy the KKT conditions where the stationarity condition is replaced with a relaxed zero gradient local optimality condition:

For any local maximum $\mathbf{p}^*$, *there exist unique Lagrange multipliers*

$$\lambda_{1,u}^*, \ldots, \lambda_{M,u}^* \text{ and } \lambda_{1,l}^*, \ldots, \lambda_{M,l}^*$$

such that for all $i \in \mathcal{M}$,

$$\frac{\partial u_i(\gamma_i(\boldsymbol{p}^*))}{\partial p_i} + \sum_{j \neq i} \frac{\partial u_j(\gamma_j(\boldsymbol{p}^*))}{\partial p_i} = \lambda_{i,u}^* - \lambda_{i,l}^*,$$

$$\lambda_{i,u}^*(p_i^* - P_i^{\max}) = 0,\ \lambda_{i,l}^*(P_i^{\min} - p_i^*) = 0,\ \lambda_{i,u}^*, \lambda_{i,l}^* \geq 0.$$

Let

$$\pi_j(p_j, p_{-j}) = -\frac{\partial u_j(\gamma_j(p_j, p_{-j}))}{\partial I_j(p_{-j})},$$

where $I_j(p_{-j}) = \sum_{k \neq j} p_k h_{kj}$ is the total interference received by user j.

Here, $\pi_j(p_j, p_{-j})$ is always nonnegative and represents user j's marginal increase in utility per unit decrease in total interference.

Now, we have:

$$\frac{\partial u_j(\gamma_j(p_j, p_{-j}))}{\partial p_i} = \frac{\partial u_j(\gamma_j(p_j, p_{-j}))}{\partial I_j(p_{-j})} \times \frac{\partial I_j(p_{-j})}{\partial p_i}$$

h_{ij}

$$= -\pi_j(p_j, p_{-j})\, h_{ij}$$

$$\frac{\partial u_i\left(\gamma_i\left(\boldsymbol{p}^*\right)\right)}{\partial p_i}+\sum_{j \neq i} \frac{\partial u_j\left(\gamma_j\left(\boldsymbol{p}^*\right)\right)}{\partial p_i}=\lambda_{i,u}^*-\lambda_{i,l}^*,$$

$$\frac{\partial u_i\left(\gamma_i\left(\boldsymbol{p}^*\right)\right)}{\partial p_i}-\sum_{j \neq i} \pi_j\left(p_j^*, p_{-j}^*\right) h_{ij}=\lambda_{i,u}^*-\lambda_{i,l}^*. \qquad (☆)$$

Viewing $\pi_j \left(= \pi_j\left(p_j, p_{-j}\right)\right)$ as a *price* charged to other users for generating interference to user j, condition (☆) is a necessary and sufficient optimality condition for the problem in which each user i specifies a power level to

$$\begin{aligned} \text{maximize} \quad & s_i\left(p_i; p_{-i}, \pi_{-i}\right)=u_i\left(\gamma_i\left(p_i, p_{-i}\right)\right)-p_i \sum_{j \neq i} \pi_j h_{ij}, \\ \text{s.t.} \quad & p_i \in \mathcal{P}_i \end{aligned}$$

assuming fixed p_{-i} and π_{-i}.

User i therefore maximizes the difference between its utility minus its payment to the other users in the network due to the interference it generates.

The payment is its transmit power times a weighted sum of other users' prices, with weights equal to the channel gains between user i's transmitter and the other users' receivers.

This pricing interpretation of the KKT conditions motivates the following asynchronous distributed pricing (ADP) algorithm. ⇨

Algorithm The ADP Algorithm

(1) INITIALIZATION: For each user $i \in \mathcal{M}$ choose some power $p_i(0) \in \mathcal{P}_i$ and price $\pi_i(0) \geq 0$.

(2) POWER UPDATE: At each $t \in T_{i,p}$, user i updates its power according to

$$p_i(t) = \mathcal{W}_i\left(p_{-i}(t^-), \pi_{-i}(t^-)\right).$$

(3) PRICE UPDATE: At each $t \in T_{i,\pi}$, user i updates its price according to

$$\pi_i(t) = \mathcal{C}_i\left(\mathbf{p}(t^-)\right).$$

- In the ADP algorithm,

  each user announces a single price and

  all users set their transmission powers based on the received prices.

- Prices and powers are asynchronously updated.

  For $i \in \mathcal{M}$, let

  $T_{i,p}$ time instances at which user i updates its power

  $T_{i,\pi}$ time instances at which user i updates its price

- User i updates its power according to

$$\mathcal{W}_i(p_{-i}, \pi_{-i}) = \arg\max_{\hat{p}_i \in \mathcal{P}_i} s_i\left(\hat{p}_i; p_{-i}, \pi_{-i}\right),$$

- Each user updates its price according to $\mathcal{C}_i(\mathbf{p}) = -\dfrac{\partial u_i(\gamma_i(\mathbf{p}))}{\partial I_i(p_{-i})}$,

  Note that in addition to being asynchronous across users,

  each user also need not update its power and price at the same time.