Controlling Congestion Games

Alfredo Garcia, Mingyi Hong
University of Virginia

June 2011

Motivation

- Congestion games have been used to model strategic interaction in communication and transportation networks in which selfish users (players) make autonomous decisions regarding routes and/or flows
- Can a congestion game be controlled ?
    - The controller lacks the ability to dictate players' choices *and* complete information on the players' payoffs yet...

Motivation

- Congestion games have been used to model strategic interaction in communication and transportation networks in which selfish users (players) make autonomous decisions regarding routes and/or flows
- Can a congestion game be controlled ?
    - The controller lacks the ability to dictate players' choices *and* complete information on the players' payoffs yet...
    - it is desirable to induce a stable state of the network with certain properties.

Motivation

- One solution to this "mechanism design" problem is the celebrated Vickrey-Clarke-Grove (VCG) mechanism in which the controller's goal is to induce a Nash equilibrium that maximizes aggregate surplus.
- In the VCG mechanism,

Motivation

- One solution to this “mechanism design” problem is the celebrated Vickrey-Clarke-Grove (VCG) mechanism in which the controller’s goal is to induce a Nash equilibrium that maximizes aggregate surplus.
- In the VCG mechanism,
  - the controller exchanges messages with each individual user regarding available routes and the associated congestion charges;

Motivation

- One solution to this “mechanism design” problem is the celebrated Vickrey-Clarke-Grove (VCG) mechanism in which the controller’s goal is to induce a Nash equilibrium that maximizes aggregate surplus.
- In the VCG mechanism,
  - the controller exchanges messages with each individual user regarding available routes and the associated congestion charges;
  - assuming route choices are verifiable, optimal network performance can be induced as a Nash equilibrium in dominant strategies (incentive compatibility).

Motivation

- Important drawbacks:
  - The VCG mechanism necessarily discriminates amongst heterogeneous users
- Mechanisms with lower computational burden

Motivation

- Important drawbacks:
  - The VCG mechanism necessarily discriminates amongst heterogeneous users
  - Significant computational overhead (e.g. message exchange, ability to verify users route choices).
- Mechanisms with lower computational burden

Motivation

- Important drawbacks:
  - The VCG mechanism necessarily discriminates amongst heterogeneous users
  - Significant computational overhead (e.g. message exchange, ability to verify users route choices).
- Mechanisms with lower computational burden
  - would be anonymous (can not price discriminate)

Motivation

- Important drawbacks:
  - The VCG mechanism necessarily discriminates amongst heterogeneous users
  - Significant computational overhead (e.g. message exchange, ability to verify users route choices).
- Mechanisms with lower computational burden
  - would be anonymous (can not price discriminate)
  - reach desired state over time (instead of in "one shot")

Motivation

- Important drawbacks:
  - The VCG mechanism necessarily discriminates amongst heterogeneous users
  - Significant computational overhead (e.g. message exchange, ability to verify users route choices).
- Mechanisms with lower computational burden
  - would be anonymous (can not price discriminate)
  - reach desired state over time (instead of in "one shot")
  - decentralized (instead of centralized)

Controlling a Congestion Game

- The controller determines bounds on aggregate utilization on a subset of critical links.
- The underlying congestion game is *controlled* when
  - congestion prices are identified online so that aggregate flows on the network stabilize (users reach equilibrium) and

Controlling a Congestion Game

- The controller determines bounds on aggregate utilization on a subset of critical links.
- The underlying congestion game is *controlled* when
  - congestion prices are identified online so that aggregate flows on the network stabilize (users reach equilibrium) and
  - aggregate utilization in critical links is within given bounds.

Setup

- A directed graph $(\mathcal{V}, \mathcal{E})$ which represents a network of interest, and a set $\mathcal{K}$ of network users.
- Each user $k \in \mathcal{K}$ is associated with an origin-destination pair $(o_k, d_k) \in \mathcal{V} \times \mathcal{V}$.
- A path $p \in \mathcal{P}_k$ is defined as an ordered set of links connecting o_k to d_k.

Setup

- We assume that flow of user $k \in \mathcal{K}$ can be continuously split and routed over a subset of available paths $p \in \mathcal{P}_k$.
- The delay associated with link $e \in \mathcal{E}$ is modeled by the function $c_e : \mathbb{R}_+ \mapsto \mathbb{R}_+$ which is assumed strictly increasing and twice-differentiable.
- Given joint path choice $\mathbf{p}$ and demand $\mathbf{r} = \{r_k : k \in \mathcal{K}\}$ the aggregate link utilization of link $e \in \mathcal{E}$ can be written as:

$$\lambda_e(\mathbf{p}, \mathbf{r}) = \sum_{k \in \mathcal{K}} r_k \mathbf{1}_{\{e \in p_k\}}$$

where $\mathbf{1}_{\{e \in p_k\}} = 1$ if link $e \in \mathcal{E}$ is in path p_k and $\mathbf{1}_{\{e \in p_k\}} = 0$, otherwise.

Setup

- Let $\pi_k(p)$ the proportion of user k's flow on path $p \in \mathcal{P}_k$. Aggregate link utilization can be written as:

$$\lambda_e(\boldsymbol{\pi}, \mathbf{r}) = \sum_{k \in \mathcal{K}} r_k \sum_{p \in \mathcal{P}_k} \pi_k(p) \mathbf{1}_{\{e \in p\}}.$$

- Congestion "prices" $\mu_e \in \mathcal{E}$ are introduced in the form of artificially added delay so that the total delay along path p_k given joint route choice $\boldsymbol{\pi}$ and demand $\mathbf{r}$ is

$$\sum_{e \in p_k} [c_e(\lambda_e(\boldsymbol{\pi}, \mathbf{r})) + \mu_e]$$

- User k's surplus is $U_k(r)$ where r is flow $(U' > 0, U'' < 0)$ and when minimum cost is v_k demand is

$$D_k(v_k) = \max_{r \geq 0} [U_k(r) - r v_k]$$

Nash Equilibrium

A routing policy combination $\boldsymbol{\pi}^*$ together with demand profile $\mathbf{r}^*$ is a Nash equilibrium iff for every $p \in \mathcal{P}_k$ with $\pi_k^*(p) > 0$:

$$p \in \arg\min_{p_k \in \mathcal{P}_k} \sum_{e \in p_k} [c_e(\lambda_e(\boldsymbol{\pi}^*, \mathbf{r}^*)) + \mu_e]$$

and

$$r_k^* = \arg\max_{r \geq 0} [U_k(r) - r v_k(\boldsymbol{\pi}^*, \mathbf{r}^*))]$$

where $v_k(\boldsymbol{\pi}, \mathbf{r}) \triangleq \min_{p_k \in \mathcal{P}_k} \sum_{e \in p_k} [c_e(\lambda_e(\boldsymbol{\pi}, \mathbf{r})) + \mu_e]$

Nash Equilibrium

- Nash equilibrium (NE) are minimizers of the potential function $\Phi(\boldsymbol{\pi}, \mathbf{r}; \boldsymbol{\mu})$ defined as

$$\Phi(\boldsymbol{\pi}, \mathbf{r}; \boldsymbol{\mu}) = \sum_{e \in \mathcal{E}} \int_0^{\lambda_e(\boldsymbol{\pi}, \mathbf{r})} [c_e(x) + \mu_e] dx - \sum_{k \in \mathcal{K}} \int_0^{r_k} D_k^{-1}(x) dx$$

- Convexity of Φ implies all Nash equilibria induce the same aggregate utilization on links, i.e. if $(\boldsymbol{\pi}^*, \mathbf{r}^*)$ and $(\boldsymbol{\pi}^{**}, \mathbf{r}^{**})$ are NE then $\lambda_e(\boldsymbol{\pi}^*, \mathbf{r}^*) = \lambda_e(\boldsymbol{\pi}^{**}, \mathbf{r}^{**}) = \lambda_e^*$

Controlling a Congestion Game

- Let $\{\bar{\lambda}_e > 0 : e \in \bar{\mathcal{E}}\}$ denote the set of desired bounds on aggregate utilization.
- With complete information, one can find congestion prices, say μ_e^*, $e \in \bar{\mathcal{E}}$ so that the unique aggregate demand for utilization in equilibrium "clears" the available "supply", i.e., $\mu_e^*(\bar{\lambda}_e - \lambda_e^*) = 0$ for all $e \in \bar{\mathcal{E}}$.
- Non-linear complementarity problem (Larsson and Patriksson (1999), Yang et al. (2010))

Controlling a Congestion Game

- Under incomplete information, the challenge is to identify $\mu_e^*, e \in \bar{\mathcal{E}}$ through online interaction with users
- Tools for control:
  - dynamically adjusting congestion prices and
  - rules governing the speed at which routes and flows can be adjusted

Controlling a Congestion Game

- Let λ_e^t and μ_e^t denote respectively, the utilization and congestion price for link $e \in \mathcal{E}$ at time $t > 0$
- Users are assumed to optimize path and flow choices:
  - the new desired path

$$p_k^{t+1} \in \arg \min_{p_k \in \mathcal{P}_k} \sum_{e \in p_k} [c_e(\lambda_e^t) + \mu_e^t],$$

Controlling a Congestion Game

- Let λ_e^t and μ_e^t denote respectively, the utilization and congestion price for link $e \in \mathcal{E}$ at time $t > 0$
- Users are assumed to optimize path and flow choices:
  - the new desired path

$$p_k^{t+1} \in \arg\min_{p_k \in \mathcal{P}_k} \sum_{e \in p_k} [c_e(\lambda_e^t) + \mu_e^t],$$

  - the new desired flow demand is

$$\tilde{r}_k^{t+1} = \arg\max_{r \geq 0} [U_k(r) - r v_k^t]$$

where v_k^t is the current mean (total) delay for user k, i.e.:

$$v_k^t = \sum_{p \in \mathcal{P}_k} \pi_k^t(p) \sum_{e \in p} [c_e(\lambda_e^t) + \mu_e^t]$$

Controlling a Congestion Game: Inertia

- The *actual* implemented routes and flows are adjusted according to:

$$\pi_k^{t+1}(p) = \frac{t}{t+1}\pi_k^t(p) + \frac{1}{t+1}\mathbf{1}_{\{p=p_k^{t+1}\}}$$

and

$$r_k^{t+1} = \frac{t}{t+1}r_k^t + \frac{1}{t+1}\tilde{r}_k^{t+1}$$

- Congestion prices are updated as follows:

$$\mu_e^{t+1} = \left[\mu_e^t + \rho\left(\lambda_e^t - \bar{\lambda}_e\right)\right]^+$$

where $\rho > 0$ and $e \in \bar{\mathcal{E}}$.

Assume users can adjust flows r_k but routes $\{\pi_k(p)\}$ are *fixed* $\forall k, p$.

Theorem

For all $e \in \bar{\mathcal{E}}$, $\mu_e^t \to \mu_e^*$.

Corollary

For all $e \in \bar{\mathcal{E}}$, $\mu_e^t(\lambda_e^t - \bar{\lambda}_e) \to 0$. *Moreover,*

$$\begin{aligned} r_k^{t+1} &\to r_k^* \\ \lambda_e^t &\to \lambda_e^* \leq \bar{\lambda}_e \end{aligned}$$

Controlling a Congestion Game: Flow

Assume users have *fixed* flows r_k but routes $\{\pi_k(p)\}$ can be changed $\forall k, p$.

Theorem

Assume that for every user k, there exists at least in path that consists solely of links $e \notin \bar{\mathcal{E}}$. Then $\mu_e^t \to \mu_e^*$ for all $e \in \bar{\mathcal{E}}$.

Corollary

For all $e \in \bar{\mathcal{E}}$, $\mu_e^t(\lambda_e^t - \bar{\lambda}_e) \to 0$. Moreover, $\lambda_e^t \to \lambda_e^* \leq \bar{\lambda}_e$.

Conclusions

- We have introduced a mechanism that enables the "control" of a two important classes of congestion games
- No guarantees for efficiency and/or speed of convergence
- Current work: examine robustness and efficiency with few users classes