Selfish Routing

Algorithmic Game Theory Course

CoReLab (NTUA)

March 2016

Congestion Games

- set of selfish players
- finite set of resources
- congestion impairs the quality of the resources
- for every player a finite set of strategies

Each player minimizes individual cost!

routing games = congestion games + network infrastructure

- resources $\rightsquigarrow$ edges
- strategies $\rightsquigarrow$ paths

They can model:

- traffic networks
- telecommunication networks
- resource allocation settings
- habitat selection
- …

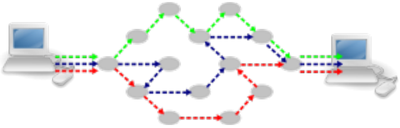

Categories of Congestion Games

- ***Atomic:*** $\rightsquigarrow$ finite set of players each one of them controlling a significant amount of traffic
- ***Non-atomic:*** $\rightsquigarrow$ infinite set of infinitesimal players

- additive/ non-additive
- bottleneck
- weighted
- congestion games with player-specific payoff functions
- ...

The Model

- single commodity directed network, $\mathcal{G} = (V, E)$ (parallel edges are allowed).
- an amount of traffic, r.
- for each edge, $e \in E$, a nonnegative, nondecreasing latency function, ℓ_e.

Flow

Vector $f = (f_p)_{p \in \mathcal{P}}$ splitting traffic among the paths of $\mathcal{G}$.

$$\textit{feasibility} : \begin{cases} f_p \geq 0, \ \ \forall p \in \mathcal{P} \\ \sum_{p \in \mathcal{P}} f_p = r \end{cases}$$

Edge Decomposition of f: $f_e = \sum_{p: e \in p} f_p$.

Example

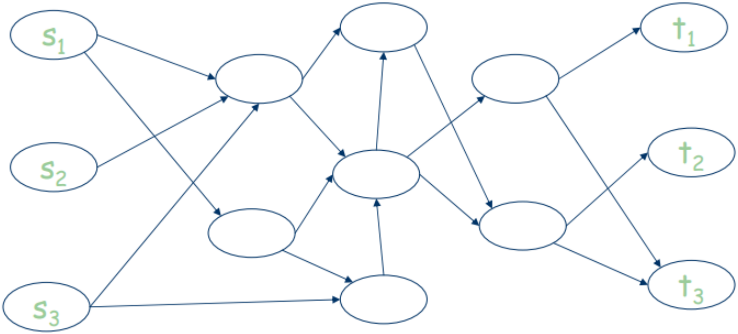

$r_1 = r_2 = r_3 = 1$ and $\ell_e(x) = x$ for all edges

Example

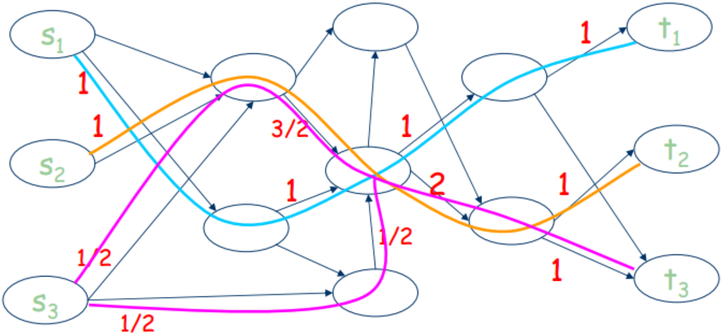

$r_1 = r_2 = r_3 = 1$ and $\ell_e(x) = x$ for all edges

The Model (cont.)

Individual Cost

- perceived cost of players on path p,
- $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$.

Social Cost

- measures quality of a flow, f,
- commonly used: the average of players cost,
- $C(f) = \sum_{p \in \mathcal{P}} f_p \ell_p(f) = \sum_{e \in E} f_e \ell_e(f_e)$.

Latency of a Flow

- $L(f) := max_{p: f_p > 0} \ell_p(f)$.

Example

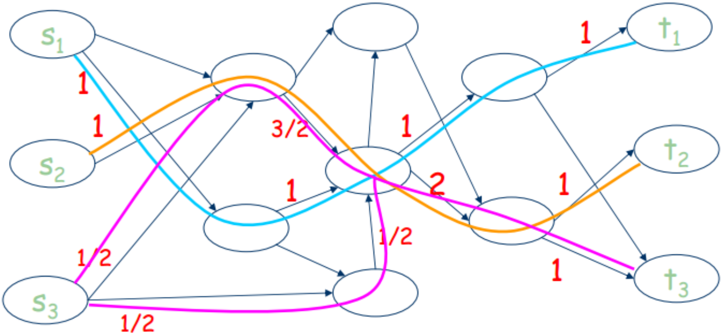

$r_1 = r_2 = r_3 = 1$ and $\ell_e(x) = x$ for all edges

Equilibrium

Definition

A feasible flow f is a *Wardrop equilibrium* if for every pair of paths $p, q \in \mathcal{P}$, with $f_p > 0$, it is $\ell_p(f) \leq \ell_q(f)$.

Intuitively, no player has incentive to deviate!

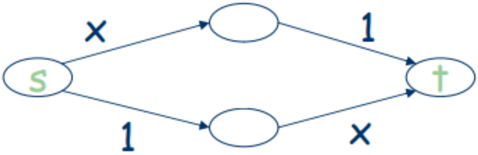

Equilibrium

Definition

A feasible flow f is a *Wardrop equilibrium* if for every pair of paths $p, q \in \mathcal{P}$, with $f_p > 0$, it is $\ell_p(f) \leq \ell_q(f)$.

Intuitively, no player has incentive to deviate!

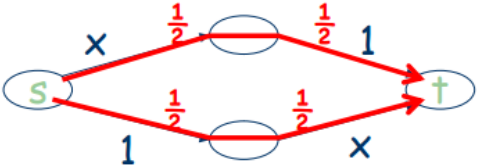

Equilibrium

Definition

A feasible flow f is a *Wardrop equilibrium* if for every pair of paths $p, q \in \mathcal{P}$, with $f_p > 0$, it is $\ell_p(f) \leq \ell_q(f)$.

Intuitively, no player has incentive to deviate!

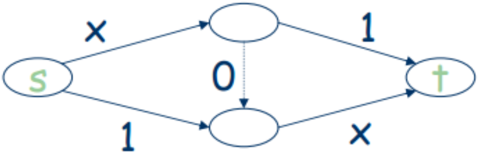

Equilibrium

Definition

A feasible flow f is a *Wardrop equilibrium* if for every pair of paths $p, q \in \mathcal{P}$, with $f_p > 0$, it is $\ell_p(f) \leq \ell_q(f)$.

Intuitively, no player has incentive to deviate!

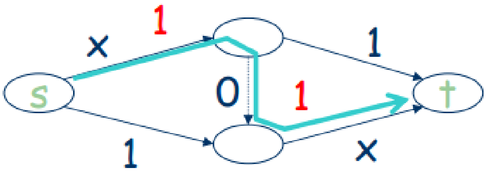

Equilibrium

Definition

A feasible flow f is a *Wardrop equilibrium* if for every pair of paths $p, q \in \mathcal{P}$, with $f_p > 0$, it is $\ell_p(f) \leq \ell_q(f)$.

Intuitively, no player has incentive to deviate!

Corollary 1

Every *used path* experiences the same latency at equilibrium (i.e. $\ell_p(f) = L(f)$, $\forall p \in \mathcal{P} : f_p > 0$).

Corollary 2

A flow is an equilibrium if and only if traffic travels *only* on shortest s-t paths.

Equilibrium

Theorem

a) Wardrop equilibrium always exists.

b) If f and f' are equilibrium flows then $\ell_e(f_e) = \ell_e(f'_e)$, $\forall e \in E$.

Proof of (a).

Method of Potential Function: define a function Φ on the outcomes of the game, so that the equilibria are exactly the outcomes that optimize Φ (local minima).

- Let $\Phi(f) := \sum_{e\in E} \int_0^{f_e} \ell_e(x)dx$
- Let f be a feasible flow, and f' be a flow that differs from f only in two paths, p, q: $f'_p = f_p - \delta$, $f'_q = f_q + \delta$, $(\delta \to 0)$.

$$\Phi(f') - \Phi(f) = \sum_{e\in E} \int_0^{f'_e} \ell_e(x)dx - \sum_{e\in E} \int_0^{f_e} \ell_e(x)dx$$

$$= \sum_{e\in q\setminus p} \int_{f_e}^{f_e+\delta} \ell_e(x)dx - \sum_{e\in p\setminus q} \int_{f_e-\delta}^{f_e} \ell_e(x)dx$$

$$\overset{\delta \to 0}{\approx} \sum_{e\in q\setminus p} \delta \ell_e(f'_e) - \sum_{e\in p\setminus q} \delta \ell_e(f_e)$$

$$= \delta(c_q(f') - c_p(f))$$

Equilibrium

Proof of (a).

Method of Potential Function: define a function Φ on the outcomes of the game, so that the equilibria are exactly the outcomes that optimize Φ (local minima).

- Let $\Phi(f) := \sum_{e \in E} \int_0^{f_e} \ell_e(x) dx$
- The set of feasible flows is compact (i.e. closed and bounded) and Φ is a continuous function on this set $\Rightarrow$ Φ achieves a minimum value.
- The first-order optimality conditions for Φ exactly match the definition of *WE*, i.e. f minimizes Φ iff f is a *WE*.

Equilibrium

Proof of (b).

Let f and f' be two equilibrium flows.

- any convex combination of them $\lambda f + (1-\lambda)f'$, $\forall \lambda \in [0,1]$ is also a feasible flow
- Φ is convex $\Rightarrow \Phi(\lambda f + (1-\lambda)f') \leq \lambda\Phi(f) + (1-\lambda)\Phi(f')$
- $\Phi(f)$ and $\Phi(f')$ are global minima
  $\Rightarrow \Phi(\lambda f + (1-\lambda)f')$ is also global minimum, $\forall \lambda \in [0,1]$
  $\Rightarrow \Phi(\lambda f + (1-\lambda)f') = \lambda\Phi(f) + (1-\lambda)\Phi(f')$
- every summand, $\int_0^t \ell_e(x)dx$, of Φ is convex
  $\Rightarrow$ *calculus* $\Rightarrow$ every summand, $\int_0^t \ell_e(x)dx$, must be linear between the values f_e and f'_e
  $\Rightarrow \ell_e(x)$ must be constant between the values f_e and f'_e

$$\Rightarrow \ell_e(f_e) = \ell_e(f'_e),\ \forall e \in E$$

□

Equilibrium

Corollary of (b)

Equilibrium is essentially unique (i.e. all *WE* have the same SC).

Proof.

- Theorem (b) $\Rightarrow \ell_p(f) = \ell_p(f'), \forall p \in \mathcal{P}$
- Let $\mathcal{P}^*(f) := \{p \in \mathcal{P} \mid \ell_p(f) \leq \ell_q(f), \forall q \in \mathcal{P}\}$, the set of shortest paths under a general flow f
- $\mathcal{P}^*(f) = \mathcal{P}^*(f') =: \mathcal{P}^* \ (\neq \emptyset)$
- Definition of *WE* $\Rightarrow$ all used paths have equal ***and*** minimum cost.
- $\Rightarrow$ both f and f' have support ***only*** in some subset of $\mathcal{P}^*$
- $\Rightarrow C(f) = C(f')$

□

Uniqueness

Obvious sufficient condition:

strictly increasing latency functions $\Rightarrow$
strictly convex potential function $\Rightarrow$
unique minimum $\Rightarrow$
unique equilibrium $\square$

Not-so-obvious necessary and sufficient condition:

$$\nexists p_1 \neq p_2 \in \mathcal{P} : \exists \epsilon > 0 :$$

$$\ell_{p_1}(f) = L(f), \forall x \in [f_{p_1} - \epsilon, f_{p_1}] \textit{ and } \ell_{p_2}(f) = L(f), \forall x \in [f_{p_2}, f_{p_2} + \epsilon]$$

(i.e. cannot modify equilibrium flow without changing path costs)

Characterizing Equilibrium

Variational Inequality

A flow f is a Wardrop equilibrium iff
$\sum_{e \in E} f_e \ell_e(f_e) \leq \sum_{e \in E} f_e^* \ell_e(f_e)$, for every feasible flow f^*.

Proof.

- $\sum_{p \in \mathcal{P}} f_p \ell_p(f) \leq \sum_{p \in \mathcal{P}} f_p^* \ell_p(f)$
  (equilibrium flow uses only minimum cost paths)
- Writing $\ell_p(f) = \sum_{e \in p} \ell_e(f_e)$ and reversing the order of summation on both sides proves the proposition.

□

Example: Pigou's network [Pigou, 1920]

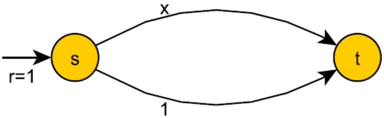

Unique equilibrium:

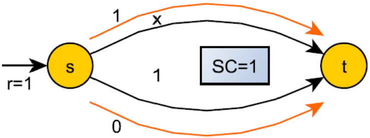

Social Optimum:

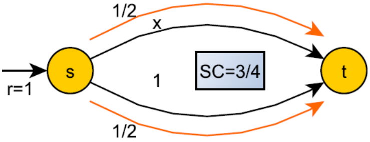

Equilibrium does not need to coincide with Social Optimum!

What About This Network?

(Non-linear Pigou's Network)

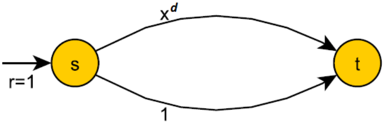

- *Equilibrium:* all traffic on the upper edge $\Rightarrow C(f) = 1$
- *Optimal:* routes $\epsilon-$fraction of the traffic on the lower edge

$$\Rightarrow C(o) \rightarrow 0 \text{ as } d \rightarrow \infty$$

Equilibrium can be arbitrarily inefficient!

Measuring Performance Degradation

Price of Anarchy (*PoA*) [Koutsoupias & Papadimitriou, '99]

Worst possible ratio between equilibrium and social optimum:

- for an instance: $PoA(\mathcal{I}) = sup\{\frac{C(f)}{C(o)} \mid f \text{ is equilibrium}\}$
- for a class of latency functions: $PoA(\mathcal{L}) = sup_{\mathcal{I} \in \mathcal{L}} PoA(\mathcal{I})$

For nonatomic routing games: *minimum performance degradation in order to achieve equilibrium!*

Bounding the Inefficiency of Equilibrium

If we do not restrict the class of allowable functions PoA grows unbounded (recall non-linear Pigou's network).

Approach #1: Focus on affine latency functions

Theorem [Roughgarden, Tardos, '00]

Let G be a network with affine latency functions. Then $PoA(G) \leq \frac{4}{3}$.

Proof.

$C(f) \leq \sum_{e \in E} o_e(\alpha_e f_e + b_e)$ *(variational inequality)*
$\leq \sum_{e \in E} \alpha_e f_e o_e + \sum_{e \in E} b_e o_e$
$\leq \sum_{e \in E} \alpha_e (o_e^2 + \frac{f_e^2}{4}) + \sum_{e \in E} b_e o_e$
$\leq C(o) + \frac{C(f)}{4} \Rightarrow PoA(G) \leq \frac{4}{3}$ □

Bounding the Inefficiency of Equilibrium

Theorem [Roughgarden, Tardos, '00]

Let G be a network with affine latency functions. Then $PoA(G) \leq \frac{4}{3}$.

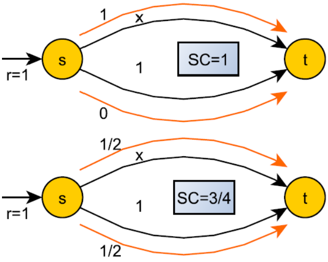

For affine latency functions Pigou's network is the worst-case instance.

Could this be the general case?

Bounding the Inefficiency of Equilibrium

[Roughgarden, '02]

- The Price of Anarchy is independent of the network topology!
- There is always a network with two parallel arcs achieving the maximum possible performance degradation.

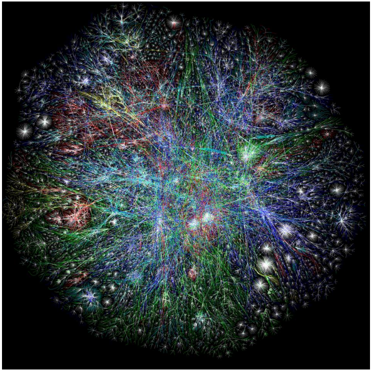

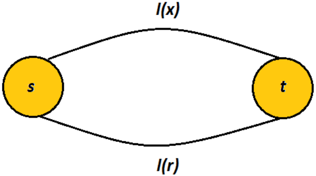

Bounding the Inefficiency of Equilibrium

Let $\mathcal{L}$ be a class of continuous and nondecreasing latency functions. Define $\beta(\mathcal{L}) = sup_{\ell \in \mathcal{L}} sup_{x \geq y \geq 0} \frac{y(\ell(x) - \ell(y))}{x\ell(x)}$.

Theorem [Correa, Schulz, Stier-Moses, '04]

The *PoA* of the instance with latency functions drawn from class $\mathcal{L}$ is bounded from above by $\rho(\mathcal{L}) := (1 - \beta(\mathcal{L}))^{-1}$ and the bound is tight.

Note that $0 \leq \beta(\mathcal{L}) < 1$, so $\rho(\mathcal{L})$ is well-defined!

Bounding the Inefficiency of Equilibrium

Proof.

Let $\mathcal{L}$ be a family of continuous, nondecreasing latency functions.

Consider the following Pigou-like network where $\ell \in \mathcal{L}$ and $r > 0$.

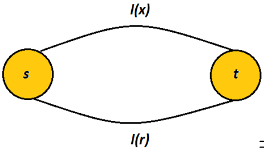

$$PoA(\mathcal{I}_{\ell,r}) = sup_{x \geq 0} \frac{r\ell(r)}{x\ell(x) + (r - x)\ell(r)}$$

$$= sup_{x \geq 0} \frac{1}{1 - \frac{x(\ell(r) - \ell(x))}{r\ell(r)}} = \frac{1}{1 - sup_{x \geq 0} \frac{x(\ell(r) - \ell(x))}{r\ell(r)}}$$

- Consider the worst traffic rate possible:
  $PoA(\mathcal{I}_{\ell}) = sup_{r \geq 0} PoA(\mathcal{I}_{\ell,r})$
- Consider the worst latency function possible:
  $PoA(\mathcal{I}_{\mathcal{L}}) = sup_{\ell \in \mathcal{L}} PoA(\mathcal{I}_{\ell})$

Bounding the Inefficiency of Equilibrium

Proof. (cont.)

Let $\mathcal{L}$ be a family of continuous, nondecreasing latency functions.

Define:

- $\beta(\ell, r) := sup_{x \geq 0} \frac{x(\ell(r) - \ell(x))}{r\ell(r)}$ $\rightsquigarrow$ *recover opt flow*
- $\beta(\ell) := sup_{r \geq 0} \beta(\ell, r)$ $\rightsquigarrow$ *worst traffic rate*
- $\beta(\mathcal{L}) := sup_{\ell \in \mathcal{L}} \beta(\ell)$ $\rightsquigarrow$ *worst instance in the class*
- $C^f(x) = \sum_{e \in E} x_e \ell_e(f_e)$

$$C(f) \leq C^f(x) \leq \sum_{e \in E} f_e \ell_e(f_e) \beta(\ell_e, f_e) + \sum_{e \in E} x_e \ell_e(x_e)$$

$$\leq \beta(\mathcal{L}) C(f) + C(x)$$

Setting $x = o$ completes the proof.

□

PoA Bounds

Description	Typical Representative	Price of Anarchy
Linear	$ax + b$	$\frac{4}{3} \approx 1.333$
Quadratic	$ax^2 + bx + c$	$\frac{3\sqrt{3}}{3\sqrt{3}-2} \approx 1.626$
Cubic	$ax^3 + bx^2 + cx + d$	$\frac{4\sqrt[3]{4}}{4\sqrt[3]{4}-3} \approx 1.896$
Polynomials of degree $\leq p$	$\sum_{i=0}^{p} a_i x^i$	$\frac{(p+1)\sqrt[p]{p+1}}{(p+1)\sqrt[p]{p+1}-p} = \Theta(\frac{p}{\ln p})$
M/M/1 Delay Functions	$(u-x)^{-1}$	$\frac{1}{2}\left(1 + \sqrt{\frac{u_{min}}{u_{min}-R_{max}}}\right)$

Alleviating Routing's Inefficiency

1. *taxing the edges of the network:* deliberately increase the perceived costs of some paths to prevent extensive usage.
2. *Stackelberg strategies:* small fraction of cooperative players influences the configuration of the rest of the users.
3. *eliminating Braess's paradox:* changing the network topology by making some edges unavailable.

Marginal Cost Tolls

- increase the latency of the edges to modify equilibrium
- compute social cost based on the initial latency functions (tolls only affect players perceived cost)

Theorem

The optimal flow for a network G with latency functions $\ell_e(x)$ is an equilibrium flow for the same network with latency functions $c_e(x) = (x\ell_e(x))'$.

Proof.

$\Phi_{G'}(f) = \sum_{e \in E} \int_0^{f_e} c_e(x)dx = \sum_{e \in E} \int_0^{f_e} (x\ell_e(x))'dx = \sum_{e \in E} f_e \ell_e(f_e) = C_G(f)$

$\Rightarrow$ the minimizer of C_G is a minimizer of $\Phi_{G'}$

$\Rightarrow$ o is an equilibrium flow for G' $\square$

What about restricted tolls?

Stackelberg Strategies

Two different sets of players:

- cooperative players, $s = \alpha r$, for $\alpha \in (0, 1)$
- selfish players, $t = (1 - \alpha) r$

Place the cooperative fraction of the flow arbitrarily in G. Then, the rest of the flow forms equilibrium based on the configuration of the cooperative players.

The flow is an equilibrium only for the selfish players!

Stackelberg Strategies

Theorem [Roughgarden, '01]

Computing optimal Stackelberg strategy is $NP-$hard even for affine latency functions and parallel-arc networks.

- Good performance guarantees for the following strategies:
  $\rightsquigarrow$ **Scale**: compute the optimal flow, o, then assign flow αo_p to every path p
  $\rightsquigarrow$ **LLF**: assign flow to the heavier paths
- Kumar & Marathe: FPTAS for Stackelberg strategies on parallel arcs.

Largest Latency First (LLF) Strategy

Focus on parallel-arc networks:

1. Compute the optimal flow, o, for G
2. Index the machines (edges) so that $\ell_1(o_1) \leq \ell_2(o_2) \leq \ell_m(o_m)$
3. Compute $k = argmin_{i \in [m]} \{\sum_{i=k+1}^{m} o_i \leq \alpha r\}$
4. Set $s_i = o_i, \forall i > k$, $s_k = \alpha r - \sum_{i=k+1}^{m} o_i$ and $s_i = 0, \forall i < k$

Largest Latency First (LLF) Strategy

Theorem [Roughgarden, '01]

For a parallel-link instance $\mathcal{I}$ with arbitrary latency functions, LLF strategy induces equilibrium with cost, at most $\frac{1}{\alpha}$ of that of the optimal flow.

Proof.

We will use induction on the number of edges, m.

W.l.o.g. assume $r = 1$. We examine two different cases:

1. $\exists e \in E : t_e = 0$, i.e. there is an edge that is not used by selfish players
2. $\forall e \in E : t_e > 0$. i.e. selfish players use all the edges

Largest Latency First (LLF) Strategy

Proof. (case 1)

- Partition the edges into two sets:
  $\rightsquigarrow E_1 = \{e \in E \mid t_e = 0\}$ *(not used by selfish players)*
  $\rightsquigarrow E_2 = \{e \in E \mid t_e > 0\}$ *(used by selfish players)*
- Let $\alpha_1 \rightsquigarrow$ amount of cooperative traffic in E_1 (α_2 in E_2 resp.)
- Let C_i be the social cost of the subinstance E_i (i = 1,2)

  $\Rightarrow C_2 = (1 - \alpha_1)L$ and $C_1 \geq \alpha_1 L$ *(L is the cost of edges in E_2)*

Largest Latency First (LLF) Strategy

Proof. (case 1, cont.)

- *Focus on M_2*:
  $\rightsquigarrow$ o is still an opt assignment for $(M_2, 1-\alpha_1)$
  $\rightsquigarrow$ s is an LLF strategy for $\mathcal{I}_2 = (M_2, 1-\alpha_1, \frac{\alpha_2}{1-\alpha_1})$

- Apply inductive hypothesis to $\mathcal{I}_2$:

$$C(o) \geq C_1 + \frac{\alpha_2}{1-\alpha_1} C_2$$

- It suffices to prove that $\alpha C_{LLF} = \alpha(C_1 + C_2) \leq C_1 + \frac{\alpha_2}{1-\alpha_1} C_2$

$\Rightarrow$ holds trivially when replacing C_1 with $\alpha_1 L$
and C_2 with $(1-\alpha_1)L$

$$\Rightarrow C(s+t) \leq \frac{1}{\alpha} C(o)$$

□

Largest Latency First (LLF) Strategy

Proof. (case 2)

- W.l.o.g. we assume that $\alpha < o_m$, i.e. we couldn't saturate the heavier edge
- $\ell_m(o_m) \geq L$, where L is the latency of equilibrium (otherwise $\ell_e(o_e) < L, \forall e \in E \Rightarrow ||o||_1 < r \rightsquigarrow$ *Contradiction!*)
- $C(o) \geq o_m \ell_m(o_m) \geq \alpha L = \alpha C(s+t)$

$$\Updownarrow$$

$$C(s+t) \leq \frac{1}{\alpha} C(o)$$

□

References

1. AGT Book: *Chapters 17, 18.*
2. J.R. Correa, A.S. Schulz and N.E. Stier-Moses: *Selfish Routing in Capacitated Networks.*, 2004
3. E. Koutsoupias and C. Papadimitriou: *Worst-case Equilibria*, STACS 1999
4. V.S.A. Kumar and M.V. Marathe: *Improved Results for Stackelberg Scheduling Strategies.*, ICALP 2002
5. T. Roughgarden: *Stackelberg Scheduling Strategies.*, STOC 2001
6. T. Roughgarden: *The Price of Anarchy is Independent of the Network Topology*, STOC 2002
7. T. Roughgarden: *Selfish Routing and the Price of Anarchy.*, Survey 2006
8. T. Roughgarden and E. Tardos: *How Bad is Selfish Routing?*, FOCS 2000