The Inefficiency of Equilibria in Congestion Games

Nicolás Stier-Moses
Columbia Business School
`ns2224@columbia.edu`

Joint work with
J. Correa (UAI, Chile), and A. Schulz (MIT)

DIMACS workshop on large scale games

April 19, 2005

©2004 Erina.L

How Efficient are Equilibria?

Some computational studies show that equilibria are typically efficient

Examples:

→ Vehicular Traffic Networks Jahn, Möhring, Schulz & S.-M. OR'05

For seven "real world" instances, the total travel time of equilibria was 2%, 2%, 4%, 7%, 10%, 14% and 15% larger than a system optimum (average $\approx 7.5\%$)

→ Telecommunication Networks Qiu, Yang, Zhang & Shenker SIGCOMM'03

Can that be formalized with theoretical bounds?

Outline

- Model, Nash Equilibrium & System Optimum
- Price of Anarchy
  - Affine Delays
  - General Delays
- Improved Models
- Extensions
- Concluding Remarks

The Traffic Model

- Network $G = (N, A)$ with OD pairs of rate r_k, $k \in K$
- Nondecreasing latency functions $\ell_a : \mathbb{R}_+ \to \mathbb{R}_+$ for $a \in A$
  - belong to a given set $\mathcal{L}$ (e.g. polynomials)
- Example: Braess' Instance

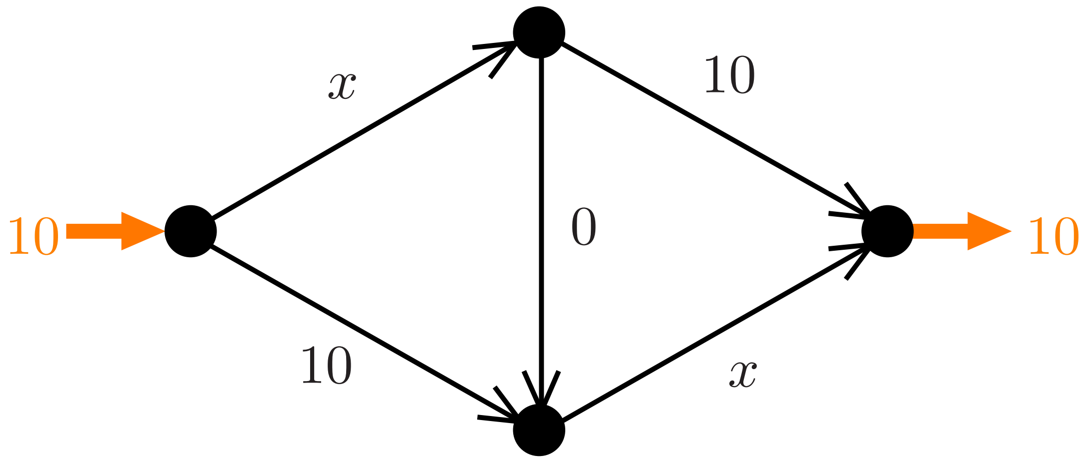

System Optimum

- Total travel time $C(x) \; := \; \sum_{a \in A} \ell_a(x_a) x_a$

- A **SO** flow x^{SO} is a feasible flow that minimizes $C(x)$

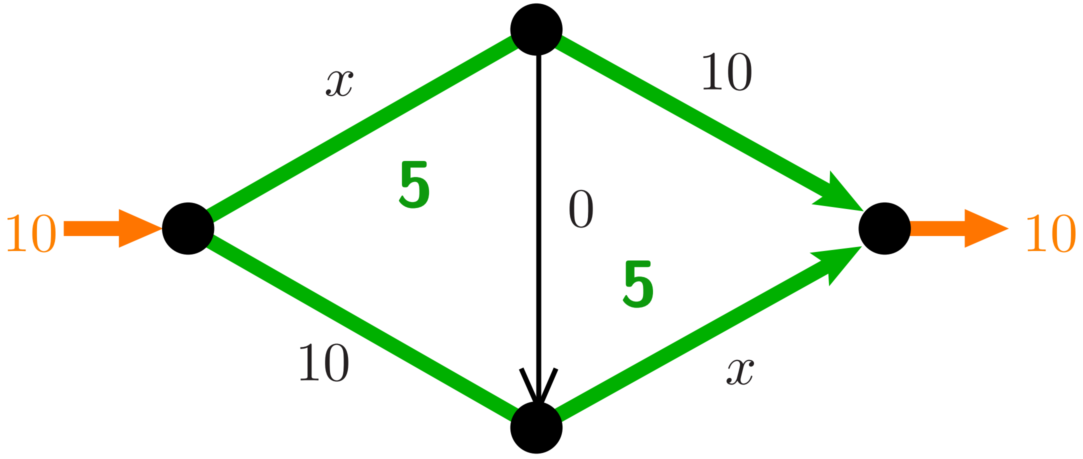

Nash Equilibrium of Nonatomic Network Game

Definition: A flow x^{NE} is a **NE** of the network game if nobody can switch to a path with smaller travel time Wardrop'52

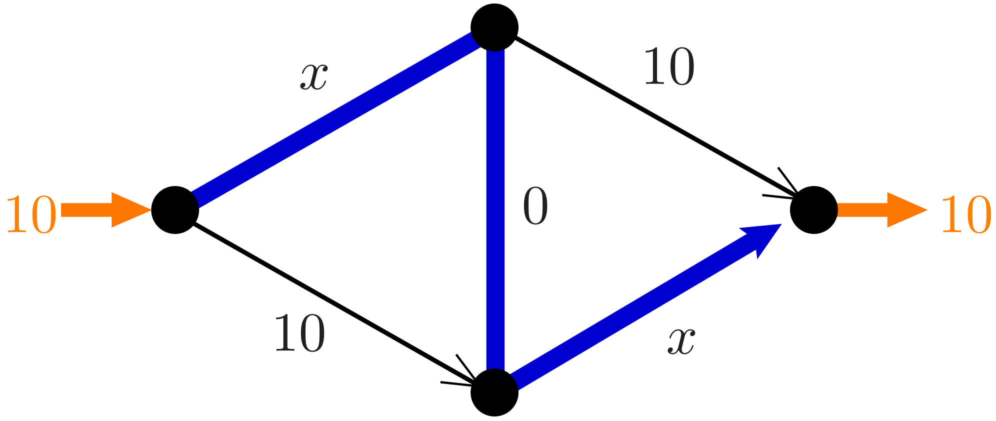

- **NE** characterized by a Variational Inequality: Smith'79

  $$\sum_{a\in A} \ell_a(x_a^{\text{NE}})x_a^{\text{NE}} \leq \sum_{a\in A} \ell_a(x_a^{\text{NE}})x_a \text{ for all } x$$ Dafermos'80

- **NE** exists & essentially unique Beckmann, McGuire & Winsten'56

Price of Anarchy

Price of Anarchy measures impact of lack of central coordination

Papadimitriou STOC'01

$$\boxed{\text{Price of Anarchy } \rho := \max_{\text{instances}} \frac{C(\textbf{NE})}{C(\textbf{SO})}}$$

- For unrestricted latency functions, ρ is unbounded
- We will assume a fixed set of latency functions $\mathcal{L}$

Price of Anarchy — Affine Latencies

Theorem.
In networks with *affine* latencies,

Roughgarden & Tardos JACM'02
Correa, Schulz & S.-M. MOR'04
Correa, Schulz & S.-M. IPCO'05

$$C(\textbf{NE}) \leq \frac{4}{3} C(\textbf{SO})$$

Selfishness drives the system close to optimality

Corollary. Braess' Instance is worst possible

Proof of 4/3 Result

Correa, Schulz & S.-M. IPCO'05

$$C(x^{\mathsf{NE}}) = \sum \ell_a(x_a^{\mathsf{NE}})x_a^{\mathsf{NE}} \leq \sum \ell_a(x_a^{\mathsf{NE}})x_a \qquad (\leftarrow \textbf{VI})$$

$$= \sum \ell_a(x_a)x_a + \sum (\ell_a(x_a^{\mathsf{NE}}) - \ell_a(x_a))x_a$$

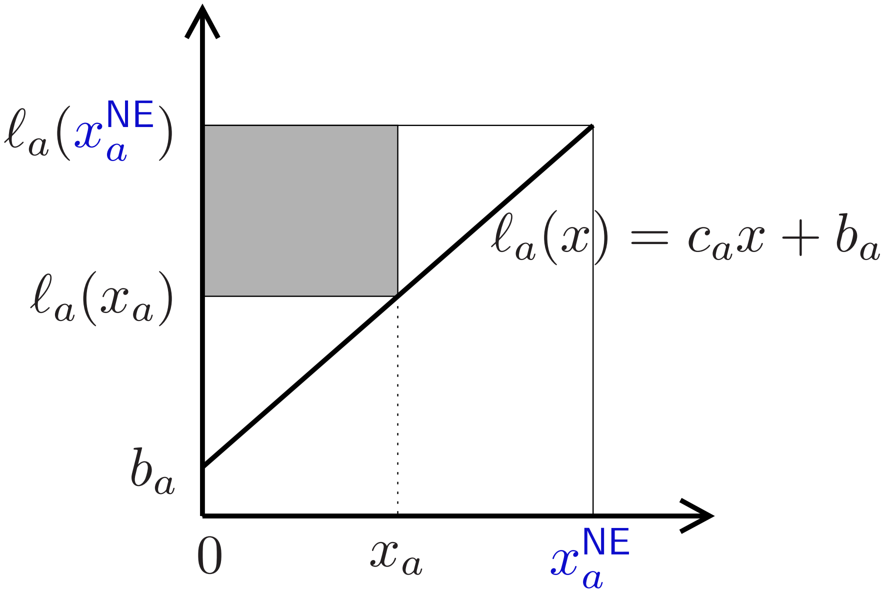

$$\leq C(x) + \tfrac{1}{4} \sum \ell_a(x_a^{\mathsf{NE}})x_a^{\mathsf{NE}} = C(x) + \tfrac{1}{4} C(x^{\mathsf{NE}})$$

Price of Anarchy — General Latencies

Roughgarden JCSS'03

Correa, Schulz & S.-M. MOR'04

Correa, Schulz & S.-M. IPCO'05

Let $\beta(\mathcal{L}) = \max_{\ell \in \mathcal{L}} \left\{ \frac{\text{shaded area}}{\text{big rectangle}} \right\}$

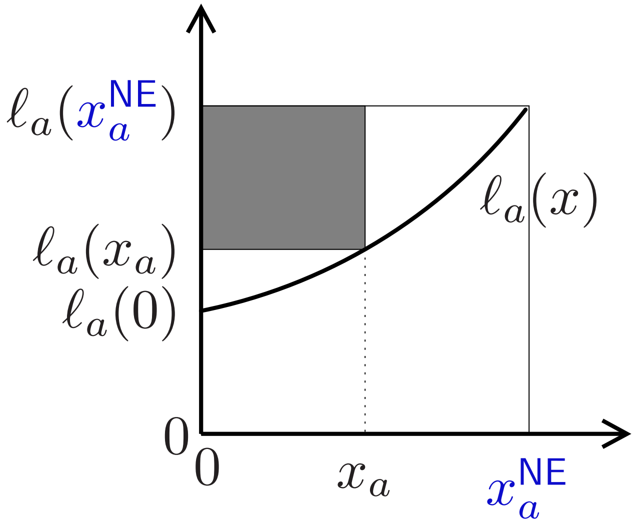

Theorem. If latencies are drawn from a family of *continuous* latencies $\mathcal{L}$,

$$\boxed{C(\textbf{NE}) \le (1-\beta(\mathcal{L}))^{-1}\ C(\textbf{SO})}$$

Example: Computing β(polynomials of degree n)

- Assume polynomials have positive coefficients

- Then, for $c \in [0,1]$: $\quad \ell(cx) \geq c^n \ell(x)$

- $$\begin{aligned} \beta(v,\ell) &= \max_{0\le x\le v} \left\{ \frac{x(\ell(v)-\ell(x))}{v\ell(v)} \right\} \\ &= \max_{0\le x\le v} \left\{ \frac{x}{v}\left(1-\frac{\ell(x)}{\ell(v)}\right) \right\} \qquad \text{(rewriting } x \text{ as } v\frac{x}{v}\text{)} \\ &\le \sup_{0\le x\le v} \left\{ \frac{x}{v}\left(1-\left(\frac{x}{v}\right)^n\right) \right\} \\ &= \sup_{0\le x\le 1} \left\{ x\left(1-x^n\right) \right\} = \frac{n}{(n+1)^{1+1/n}} \end{aligned}$$

Example: Computing POA

For a degree-n polynomial, last eqn. gives:

$$\text{POA} = \left(1 - \frac{n}{(n+1)^{1+1/n}}\right)^{-1}$$

polynomials of degree	2	3	4	$\ldots$	p
POA	1.626	1.896	2.151	$\ldots$	$\Omega(p/\ln p)$

These are tight: there are instances with same $C($**NE**$)/C($**SO**$)$

Pseudo-Approximation Results (aka Bicriteria)

Roughgarden & Tardos JACM'02

Chakrabarty 2004

Correa, Schulz & S.-M. IPCO'05

To be fair to the **NE**: Make the coordinator's life more difficult

Definition. Let **SO**$^{\alpha}$ be a system optimum of an instance with demands equal to $\alpha\, r_k$, for $k \in K$

Theorem.

$$\boxed{C(\textbf{NE}) \leq C(\textbf{SO}^{1+\beta(\mathcal{L})})}$$

Example. Affine: 1.25, Quadratic: 1.4, . . . Arbitrary: 2

Proof of Pseudo-Approximation Result

$$
\begin{aligned}
C(x^{\mathsf{NE}}) &= (1+\beta(\mathcal{L}))\sum \ell_a(x_a^{\mathsf{NE}})x_a^{\mathsf{NE}} - \beta(\mathcal{L})C(x^{\mathsf{NE}}) \\
&\le \sum \ell_a(x_a^{\mathsf{NE}})x_a - \beta(\mathcal{L})C(x^{\mathsf{NE}}) && (\leftarrow \textbf{VI}) \\
&\le C(x) && (\leftarrow \beta(\mathcal{L})\text{'s defi})
\end{aligned}
$$

With $x = \textbf{SO}^{1+\beta(\mathcal{L})}$, we get the result

Outline

- Model, Nash Equilibrium & System Optimum
- Price of Anarchy
  - Affine Delays
  - General Delays
- Improved Models
- Extensions
- Concluding Remarks

Are these Worst-Case Bounds Accurate?

polynomials of degree	2	3	4	...	p
POA	1.626	1.896	2.151	...	$\Omega(p/\ln p)$

↑

115%

In practice, we have: Jahn, Möhring, Schulz & S.-M. OR'05

In our runs the total travel time of equilibria

was not more than **15%** larger than a system optimum

Is the worst case too pessimistic?

Improved Models

1. Limited Congestion (low traffic)

2. No Fixed Delays (heavy traffic)

We give improved bounds in both cases

Limited Congestion

When system is not working under heavy traffic assumptions, it is reasonable to assume that for some $\eta \in [0,1]$:

$$\ell_a(0) \geq \eta\, \ell_a(x_a^{\mathsf{NE}})$$

[travel time at equilibrium is not much larger than travel time ‘at night’]

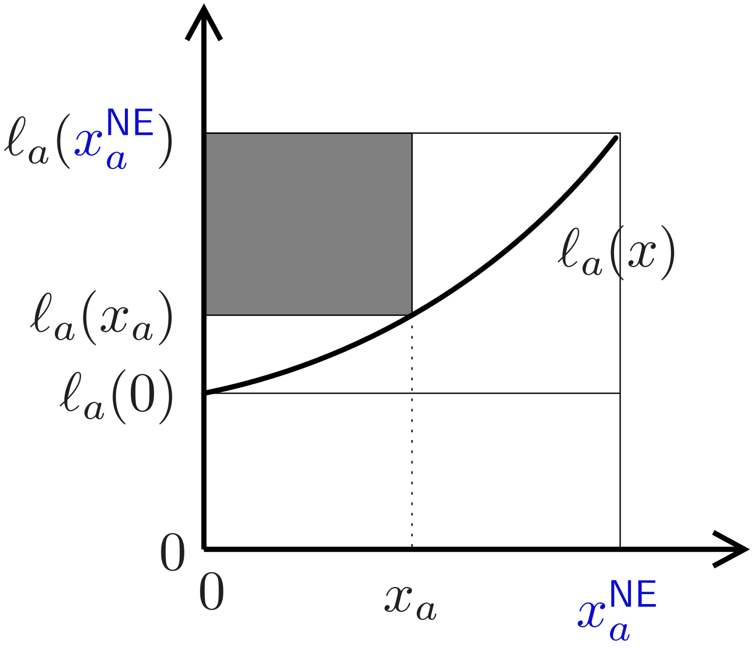

Price of Anarchy — Limited Congestion

Under those assumptions, **POA** is:

$$\boxed{(1-(1-\eta)\beta(\mathcal{L}))^{-1}}$$

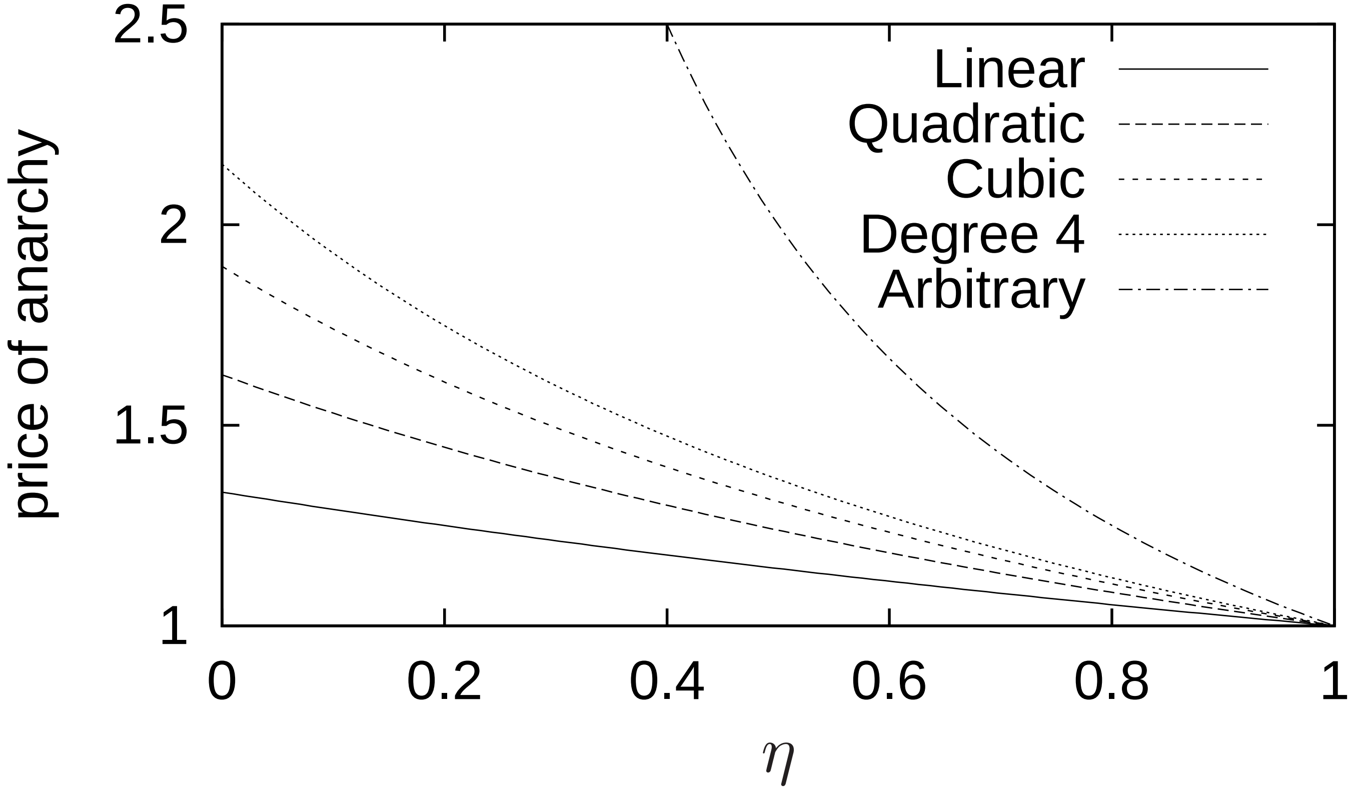

Pseudo-Approximation — Limited Congestion

The min ξ for which $C(\textbf{NE}) \leq C(\textbf{SO}^{\xi})$ is $\boxed{1 + (1-\eta)\beta(\mathcal{L})}$

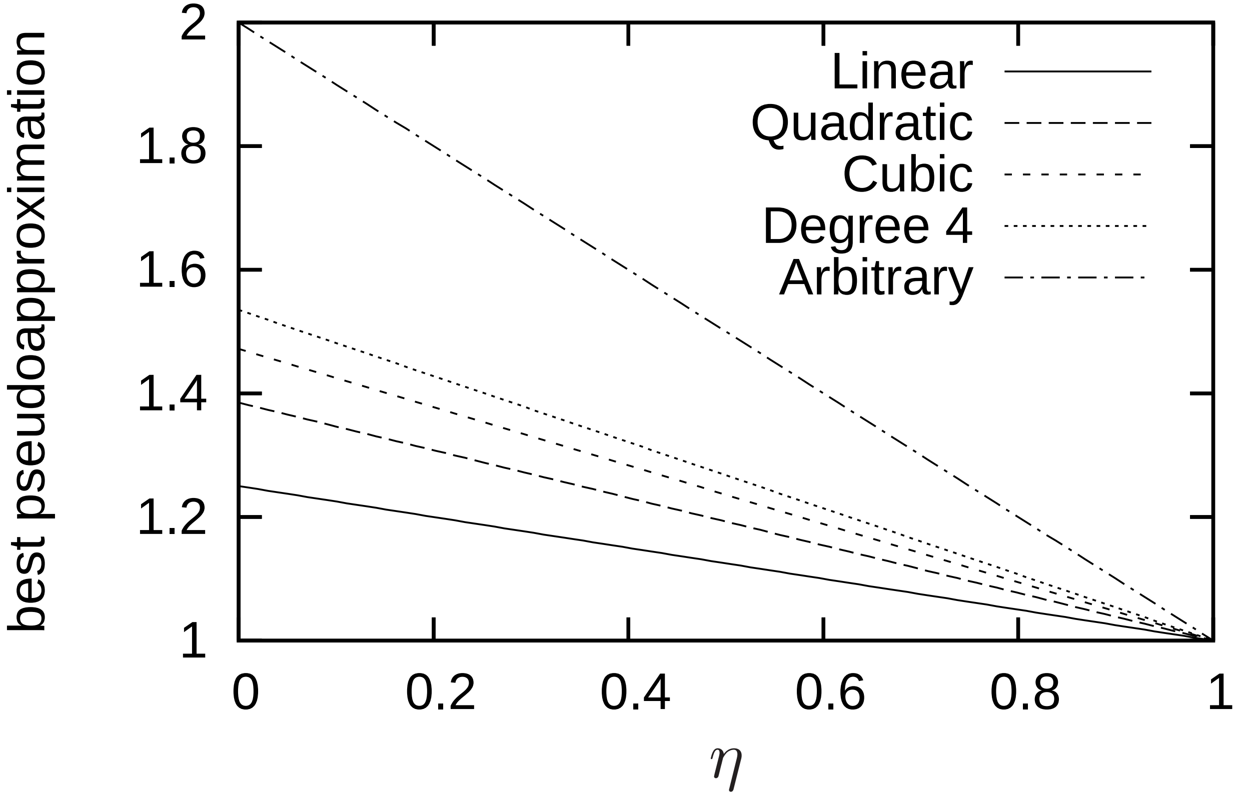

No Fixed Delays

Under heavy traffic, fixed costs have marginal importance:

$$\boxed{\text{Assume } \ell_a(0) = 0}$$

Set $\mathcal{L}$ of allowable latency functions	Price of Anarchy		
	$\ell_a(0) = 0$		$\ell_a(0)$ arbitrary
	LB	UB	
linear functions	1	1	4/3
quadratic functions	1.035	1.185	1.626
cubic functions	1.098	1.25	1.896
degree 4 polynomials	1.167	1.999	2.151

Example

Proposition. If $\ell_a(x_a) = c_{a,3}x_a^3 + c_{a,2}x_a^2 + c_{a,1}x_a$ for $a \in A$, then

$$C(\textsf{NE}) \le \tfrac{5}{4} C(\textsf{SO})$$

Proof Sketch. Bound terms of different degrees independently.

Notation: $C_i(x) := \sum_{a\in A} c_{a,i}\, x_a^{i+1}$ and $C_i^{x'}(x) := \sum_{a\in A} c_{a,i}\, x_a\, (x'_a)^i$

$$(\delta x_a - x_a^{\mathsf{NE}})^2 \ge 0 \Rightarrow \begin{cases} C_1^x(x^{\mathsf{NE}}) & \le C_1(x^{\mathsf{NE}}) + \frac{1}{4}C_1(x) \\ C_2^{x^{\mathrm{NE}}}(x) & \le \frac{1}{2}C_2^x(x^{\mathsf{NE}}) + \frac{1}{2}C_2(x^{\mathsf{NE}}) \end{cases}$$

$$\left(\tfrac{x_a^2}{2} - (x_a^{\mathsf{NE}})^2 + x_a\, x_a^{\mathsf{NE}}\right)^2 \ge 0 \Rightarrow C_3^{x^{\mathrm{NE}}}(x) \le \frac{C_3^x(x^{\mathsf{NE}})}{2} + \frac{C_3(x^{\mathsf{NE}})}{2} + \frac{C_3(x)}{8}$$

Example (cont)

$$
\begin{aligned}
C(x^{\mathsf{NE}}) &\le C^{x^{\mathrm{NE}}}(x) = C_3^{x^{\mathrm{NE}}}(x) + C_2^{x^{\mathrm{NE}}}(x) + C_1^{x^{\mathrm{NE}}}(x) \\
&\le \frac{C_3^x(x^{\mathsf{NE}})}{2} + \frac{C_3(x^{\mathsf{NE}})}{2} + \frac{C_3(x)}{8} \quad + \quad \frac{C_2^x(x^{\mathsf{NE}})}{2} + \frac{C_2(x^{\mathsf{NE}})}{2} \\
&\qquad\qquad + \quad \frac{C_1^x(x^{\mathsf{NE}})}{2} + \frac{C_1(x^{\mathsf{NE}})}{2} + \frac{C_1(x)}{8} \\
&\le \frac{C(x^{\mathsf{NE}})}{2} + \frac{C^x(x^{\mathsf{NE}})}{2} + \frac{C(x)}{8} \\
&\le \tfrac{1}{2}C(x^{\mathsf{NE}}) + \tfrac{5}{8}C(x)
\end{aligned}
$$

Extensions

- Nonseparable Latency Functions
- Nonatomic Congestion Games
- Side Constraints (e.g., Capacities)
- Maximum Latency Objective

Nonseparable Latency Functions

Latencies usually depend on flow on other arcs: $\boxed{\ell_a(f) : \mathbb{R}^A \to \mathbb{R}}$

All works with $\beta(\ell, v) := \max\limits_{x \in \mathbb{R}^A, x \geq 0} \dfrac{\langle \ell(v) - \ell(x), x \rangle}{\langle \ell(v), v \rangle}$

Example: Affine Nonseparable Latencies

Chau & Sim ORL'03

Correa, Schulz & S.-M. IPCO'05

Theorem. In networks with *affine, symmetric* and *nonseparable* latencies (i.e., $\ell(x) = Ax + b$, where $\ell : \mathbb{R}^A \to \mathbb{R}^A$)

$$C(\textbf{NE}) \leq \tfrac{4}{3} C(\textbf{SO})$$

and

$$C(\textbf{NE}) \leq C(\textbf{SO}^{1.25})$$

Perakis has POA results for the *asymmetric* case — Perakis IPCO'04

Her bounds depend on a measure of the asymmetry

Nonatomic Congestion Games

Chau & Sim ORL'03

Roughgarden & Tardos GEB'04

Correa, Schulz & S.-M. IPCO'05

We only used the **VI** and never the network structure

→ Everything holds in the more general NCG

Networks	NCG
arcs	resources
OD pairs	player types
paths	strategies: sets of resources
occupancy	rate of consumption

Side Constraints

- To capture real-world phenomena we can add side constraints to traffic network

  Charnes & Cooper '61
  Jorgensen '63, Hearn '80
  Larsson & Patriksson '95
  . . . , Schulz & S.-M. SODA'03

- It is easy to generalize **SO** (add constraints to math. prog.)

- There are multiple equilibria and **POA is** ∞ even with linear fn

- The **Best NE** is hard to characterize but a good one can be found with same **VI**

- Then: $C(\textbf{best NE}) \leq (1-\beta(\mathcal{L}))^{-1}\, C(\textbf{SO})$ **Price of Stability**

Maximum Latency Objective

Weitz 2001

Roughgarden SODA'04

Correa, Schulz & S.-M. IPCO'04,OR'05

Definition. Let **MM** be a flow that minimizes the maximum latency

Proposition. In s-t-*networks* with latencies drawn from a family of *continuous* latencies $\mathcal{L}$,

$$\boxed{L(\textbf{NE}) \leq (1-\beta(\mathcal{L}))^{-1}\ L(\textbf{MM})}$$

Proof:

$$L(\textbf{NE}) = C(\textbf{NE}) \leq \alpha(\mathcal{L})\, C(\textbf{SO}) \leq \alpha(\mathcal{L})\, C(\textbf{MM}) \leq \alpha(\mathcal{L})\, L(\textbf{MM})$$

Conclusion

- These results apply to games whose equilibria can be given by VIs
- ‘Worst-case’ may be too pessimistic if we don’t restrict instances
- Insights from real-world networks leads to better bounds

Open Problems

- Improved Models: Better ways to characterize “realistic” networks?
- Heavy traffic assumption: General Approach?
  (POA unbounded with high degree polynomials)
- More Applications