Optimal incentives for socio-technical systems

Dario Paccagnan

ETH zürich

UC SANTA BARBARA

Acknowledgements

Rahul Chandan

Bryce Ferguson

Jason Marden

FNSNF
FONDS NATIONAL SUISSE
SCHWEIZERISCHER NATIONALFONDS
FONDO NAZIONALE SVIZZERO
SWISS NATIONAL SCIENCE FOUNDATION

NSF

Office of Naval Research
USN ONR
Science & Technology

Multiagent Coordination

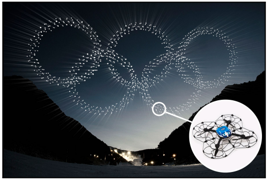

Social Systems ⟵ - - Socio-technical Systems - - ⟶ Engineered Systems

Multiagent Coordination

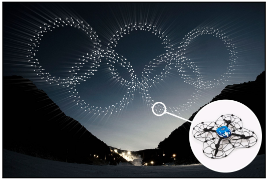

Social Systems ←-- Socio-technical Systems --→ Engineered Systems

- ▷ Traffic
- ▷ Energy Markets

- ▷ Resource Allocation
- ▷ Sensor Coverage

Multiagent Coordination

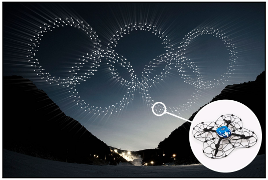

Social Systems ← - - Socio-technical Systems - - → Engineered Systems

- ▷ Traffic
- ▷ Energy Markets

- ▷ Resource Allocation
- ▷ Sensor Coverage

Multiagent Coordination

Central Goal: coordinate socio-technical systems to desirable behaviour

Social Systems ⟵ - -

Socio-technical Systems

- ▷ Traffic
- ▷ Energy Markets

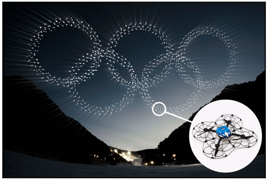

- - ⟶ Engineered Systems

- ▷ Resource Allocation
- ▷ Sensor Coverage

Socio-technical systems

Infrastructure

Socio-technical systems

Infrastructure

+

Users Behavior

Socio-technical systems

Infrastructure

$+$

Users Behavior

$=$

Performance

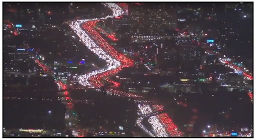

Socio-technical systems

Infrastructure

+

Users Behavior

=

Performance

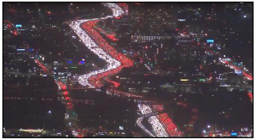

Q: How to incentivize desirable system-level behaviour?

Socio-technical systems

Infrastructure

+

Users Behavior

=

Performance

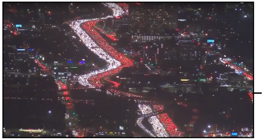

Mechanism

Q: How to incentivize desirable system-level behaviour?

Socio-technical systems are pervasive ...

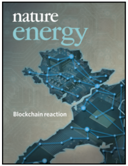

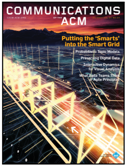

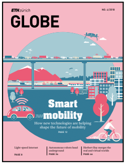

Socio-technical systems are pervasive ...

Paradigm shift: technology now interacts with human users

… and come with key challenges

GOV.UK

Search

Policy paper

The Grand Challenges

Updated 13 September 2019

Contents

The 4 Grand Challenges are focused on the trends which will transform our future:

- Artificial Intelligence and data
- Ageing society
- Clean growth
- Future of mobility

… and come with key challenges

GOV.UK

Search

Policy paper

The Grand Challenges

Updated 13 September 2019

Contents

- Artificial Intelligence and data
- Ageing society
- Clean growth
- Future of mobility

The 4 Grand Challenges are focused on the trends which will transform our future:

- Artificial Intelligence and data
- Ageing society
- Clean growth
- Future of mobility

… and come with key challenges

GOV.UK   Search

Policy paper

The Grand Challenges

Updated 13 September 2019

Contents

The 4 Grand Challenges are focused on the trends which will transform our future:

- Artificial Intelligence and data
- Ageing society
- Clean growth
- Future of mobility

⇝ **an interdisciplinary endeavour**:
computer science, control theory, optimization, economics, social sciences, urban planning, …

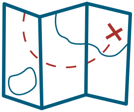

ROADMAP

1. Mechanisms for smart mobility: congestion pricing

$\rightsquigarrow$ optimal tolling mechanisms

2. Outlook and opportunities

Congestion is soaring...

New York

London

Beijing

Nairobi

Congestion is soaring...

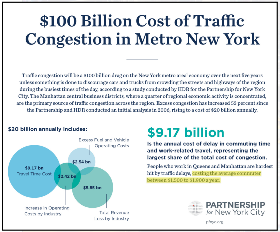

$100 Billion Cost of Traffic Congestion in Metro New York

Traffic congestion will be a $100 billion drag on the New York metro area[1] economy over the next five years unless something is done to discourage cars and trucks from crowding the streets and highways of the region during the busiest times of the day, according to a study conducted by HDR for the Partnership for New York City. The Manhattan central business districts, where a quarter of regional economic activity is concentrated, are the primary source of traffic congestion across the region. Excess congestion has increased 53 percent since the Partnership and HDR conducted an initial analysis in 2006, rising to a cost of $20 billion annually.

$20 billion annually includes:

$9.17 billion

Is the annual cost of delay in commuting time and work-related travel, representing the largest share of the total cost of congestion.

People who work in Queens and Manhattan are hardest hit by traffic delays, costing the average commuter between $1,500 to $1,900 a year.

PARTNERSHIP for New York City

pfnyc.org

CHINADAILY 中国日报网 .COM.CN   US | EUROPE | AFRICA | ASIA | 中文   July 9, 2019

Traffic jams cost average Beijinger $1,126 annually

By Meng Jing (chinadaily.com.cn)

Updated: 2016-01-20 16:44

The Fourth Ring Road in Beijing in a morning rush hour. City authorities are working on solutions to ease the pressures on traffic in the city. [Zhuo Ensen / For China Daily]

The Telegraph   News Politics Sport Business Money Opinion Tech Life & Style Travel Culture

Traffic jams cost the average motorist more than £1,000 a year

Motorists typically faced costs of £1,168 each CREDIT: ANDREW MILLIGAN/PA

New York

London

Beijing

Nairobi

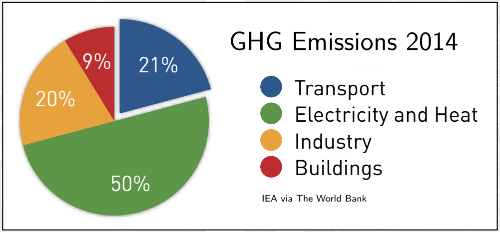

…and tolls being proposed to alleviate the issue


The New York Times

Over $10 to Drive in Manhattan? What We Know About the Congestion Pricing Plan

The price of entering Manhattan could reach as much as $25 for some drivers.
Dave Sanders for The New York Times


Forbes

Most Cities Will Have To Introduce Congestion Charging, Say Experts At Global Transit Conference

Carlton Reid Contributor

Traffic jam and queues on the highway E4 outside Stockholm. GETTY

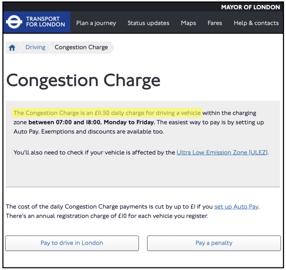
MAYOR OF LONDON

TRANSPORT FOR LONDON | Plan a journey | Status updates | Maps | Fares | Help & contacts

Driving > Congestion Charge

Congestion Charge

The Congestion Charge is an £11.50 daily charge for driving a vehicle within the charging zone **between 07:00 and 18:00, Monday to Friday.** The easiest way to pay is by setting up Auto Pay. Exemptions and discounts are available too.

You'll also need to check if your vehicle is affected by the Ultra Low Emission Zone (ULEZ).

The cost of the daily Congestion Charge payments is cut by up to £1 if you set up Auto Pay. There's an annual registration charge of £10 for each vehicle you register.

Pay to drive in London | Pay a penalty

...and tolls being proposed to alleviate the issue

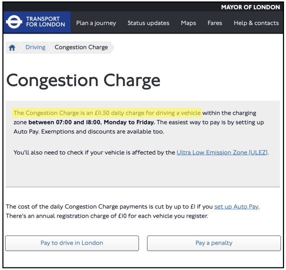

▷ Current: blunt policies

…and tolls being proposed to alleviate the issue


The New York Times

Over $10 to Drive in Manhattan? What We Know About the Congestion Pricing Plan

The price of entering Manhattan could reach as much as $25 for some drivers.
Dave Sanders for The New York Times


Forbes

Most Cities Will Have To Introduce Congestion Charging, Say Experts At Global Transit Conference

Carlton Reid Contributor

Traffic jam and queues on the highway E4 outside Stockholm. GETTY

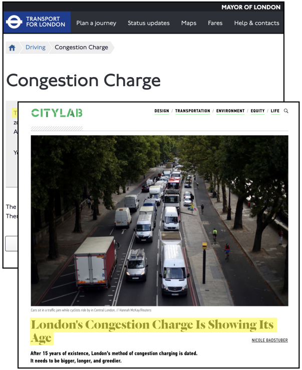
MAYOR OF LONDON

TRANSPORT FOR LONDON

Plan a journey Status updates Maps Fares Help & contacts

Driving Congestion Charge

Congestion Charge

CITYLAB

DESIGN / TRANSPORTATION / ENVIRONMENT / EQUITY / LIFE

London's Congestion Charge Is Showing Its Age

NICOLE BADSTUBER

After 15 years of existence, London's method of congestion charging is dated. It needs to be bigger, longer, and greedier.

- ▷ Current: blunt policies
- ▷ Future: fine grained + adaptive pricing using location data

How do we design fine grained and adaptive congestion pricing?

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

Congestion pricing in a nutshell

▷ **Problem:** collective behaviour of selfish agents is often inefficient

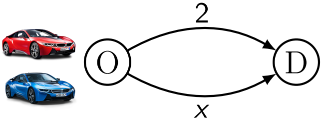

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

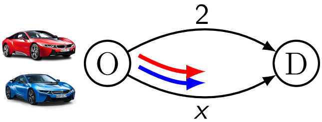

Selfish routing
System cost: $2 + 2 = 4$

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

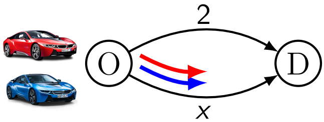

Selfish routing
System cost: $2 + 2 = 4$

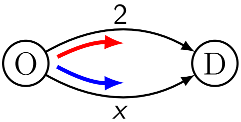

Optimal routing
System cost: $2 + 1 = 3$

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

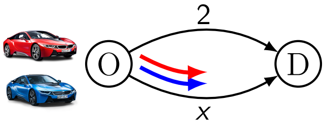

Selfish routing
System cost: $2 + 2 = 4$

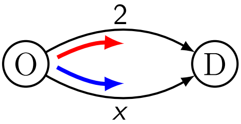

Optimal routing
System cost: $2 + 1 = 3$

- **Congestion pricing:** influence behavior to minimize total traveltime

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

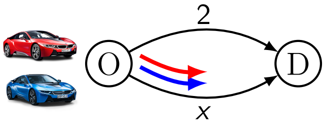

Selfish routing
System cost: $2 + 2 = 4$

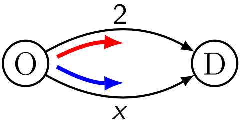

Optimal routing
System cost: $2 + 1 = 3$

- **Congestion pricing:** influence behavior to minimize total traveltime

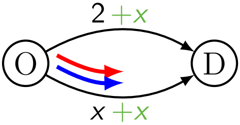

Selfish routing + tolls $\tau(x) = x$

Congestion pricing in a nutshell

▷ **Problem:** collective behaviour of selfish agents is often inefficient

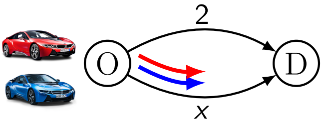

Selfish routing
System cost: $2 + 2 = 4$

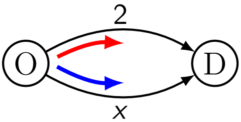

Optimal routing
System cost: $2 + 1 = 3$

▷ **Congestion pricing:** influence behavior to minimize total traveltime

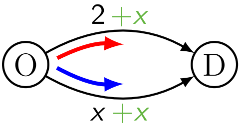

Selfish routing + tolls $\tau(x) = x$

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

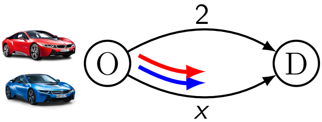

Selfish routing
System cost: $2 + 2 = 4$

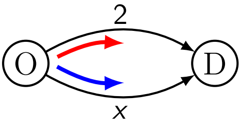

Optimal routing
System cost: $2 + 1 = 3$

- **Congestion pricing:** influence behavior to minimize total traveltime

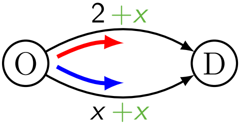

Selfish routing + tolls $\tau(x) = x$
System cost back to 3

- Robustness?
- Information?

Congestion pricing in a nutshell

- **Problem:** collective behaviour of selfish agents is often inefficient

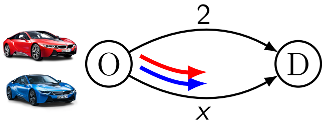

Selfish routing
System cost: $2 + 2 = 4$

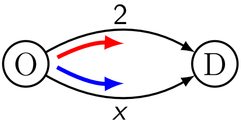

Optimal routing
System cost: $2 + 1 = 3$

- **Congestion pricing:** influence behavior to minimize total traveltime

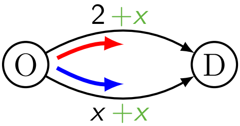

Selfish routing + tolls $\tau(x) = x$
System cost back to 3

- Robustness?
- Information?

Q: how to compute "optimal" tolls?

Routing on a network - model

Routing on a network - model

- graph

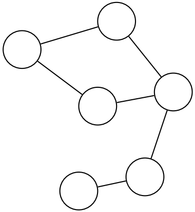

Routing on a network - model

- graph

- agent i, $\{O_i, D_i\}$

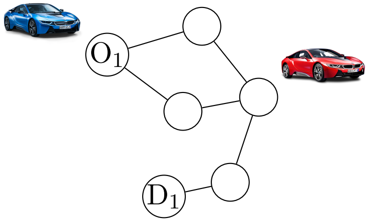

Routing on a network - model

- graph

- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$

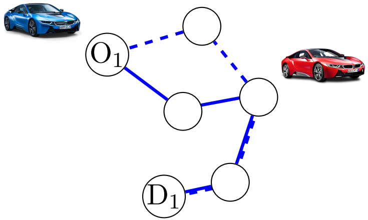

Routing on a network - model

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

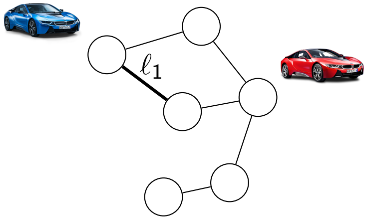

Routing on a network - model

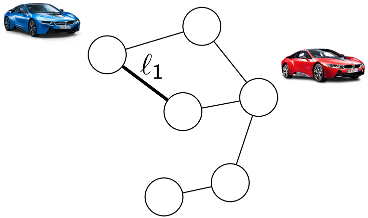

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e)$$

Routing on a network - model

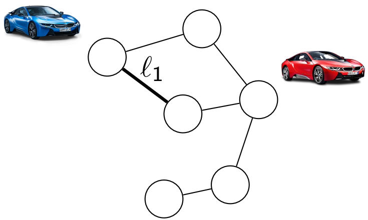

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

Routing on a network - model

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

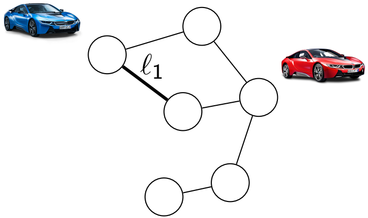

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Inefficiency} = \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

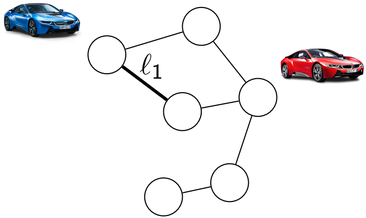

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Inefficiency} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

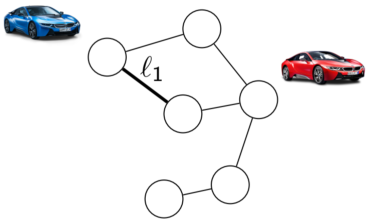

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Price of Anarchy} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

- graph

- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$

- latency functions $\ell_e(|p|_e)$

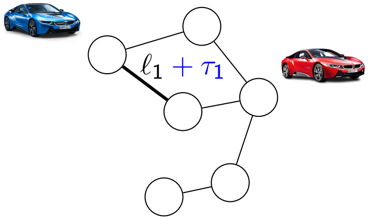

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e) + \tau_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Price of Anarchy} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

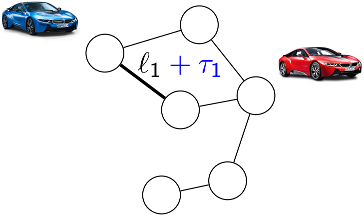

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e) + \tau_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Price of Anarchy} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

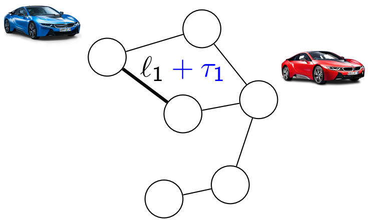

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e) + \tau_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Price of Anarchy} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Routing on a network - model

- graph
- agent i, $\{O_i, D_i\} \Rightarrow$ set of paths $\mathcal{P}_i$
- latency functions $\ell_e(|p|_e)$

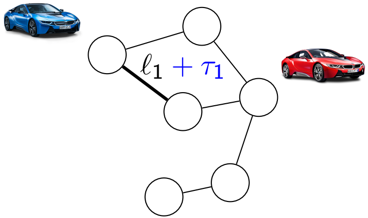

agents' costs

$$C_i(p) = \sum_{e \in p_i} \ell_e(|p|_e) + \tau_e(|p|_e)$$

total traveltime

$$TT(p) = \sum_{e \in E} |p|_e \ell_e(|p|_e)$$

$$\text{Price of Anarchy} = \max_{\text{set of instances}} \frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}}$$

Goal: design tolls that *minimize* price of anarchy

How much information to compute each toll?

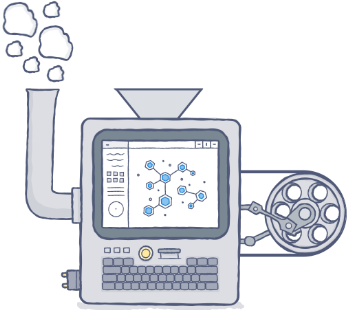

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

How much information to compute each toll?

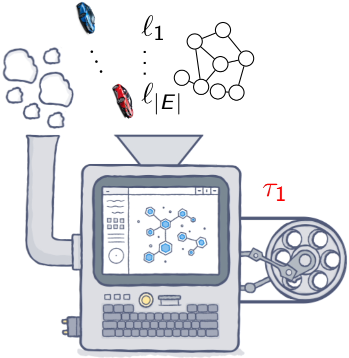

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

How much information to compute each toll?

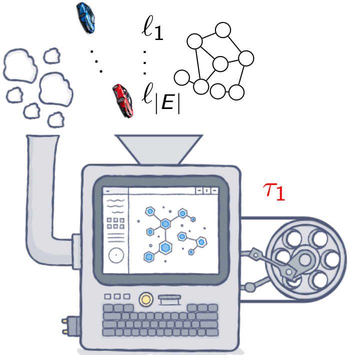

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

local info: $\tau_e = T(\ell_e)$

How much information to compute each toll?

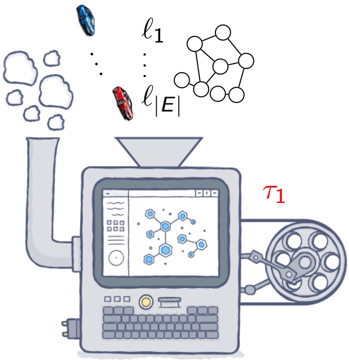

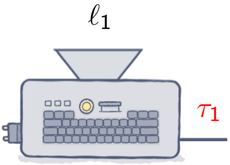

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

local info: $\tau_e = T(\ell_e)$

How much information to compute each toll?

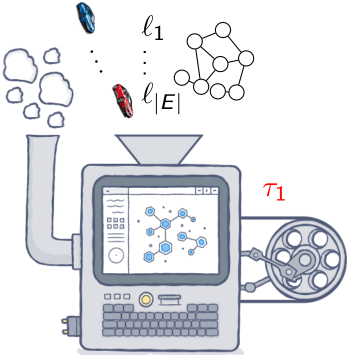

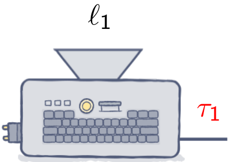

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

- $+$ more performance
- $-$ requires more computation
- $-$ not robust

local info: $\tau_e = T(\ell_e)$

How much information to compute each toll?

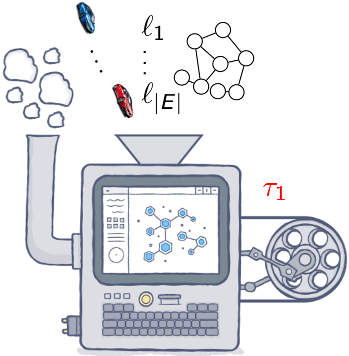

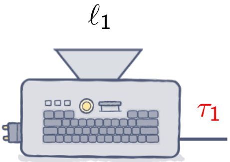

full info: $\tau_e = T(\{O_i, D_i\}, \{\ell_e\}, \text{graph})$

- \+ more performance
- − requires more computation
- − not robust

local info: $\tau_e = T(\ell_e)$

- − less performance
- \+ simpler computation
- \+ robust

History and related works

▷ **Congestion games (Rosenthal 1973)**

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

- **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

Worst-Case Equilibria

Elias Koutsoupias[1] and Christos Papadimitriou[2]

[1] Univ of California, Los Angeles
elias@cs.ucla.edu
[2] Univ of California, Berkeley
christos@cs.berkeley.edu

Abstract. In a system in which noncooperative agents share a common resource, we propose the ratio between the worst possible Nash equilibrium and the social optimum as a measure of the effectiveness of the system. Deriving upper and lower bounds for this ratio in a model in which several agents share a very simple network leads to some interesting mathematics, results, and open problems.

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

- **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

How Bad Is Selfish Routing?

TIM ROUGHGARDEN AND ÉVA TARDOS

Cornell University, Ithaca, New York

Abstract. We consider the problem of routing traffic to optimize the performance of a congested network. We are given a network, a rate of traffic between each pair of nodes, and a latency function for each edge specifying the time needed to traverse the edge given its congestion; the objective is to route traffic such that the sum of all travel times—the total latency—is minimized.

In many settings, it may be expensive or impossible to regulate network traffic so as to implement an optimal assignment of routes. In the absence of regulation by some central authority, we assume that each network user routes its traffic on the minimum-latency path available to it, given the network congestion caused by the other users. In general such a "selfishly motivated" assignment of traffic to paths will not minimize the total latency; hence, this lack of regulation carries the cost of decreased network performance.

In this article, we quantify the degradation in network performance due to unregulated traffic. We

Worst-Case Equilibria

...utsoupias[1] and Christos Papadimitriou[2]

[1] Univ of California, Los Angeles
elias@cs.ucla.edu
[2] Univ of California, Berkeley
christos@cs.berkeley.edu

Abstract. In a system in which noncooperative agents share a common resource, we propose the ratio between the worst possible Nash equilibrium and the social optimum as a measure of the effectiveness of the system. Deriving upper and lower bounds for this ratio in a model in which several agents share a very simple network leads to some interesting mathematics, results, and open problems.

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

History and related works

▷ **Congestion games (Rosenthal 1973)**

applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

▷ **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

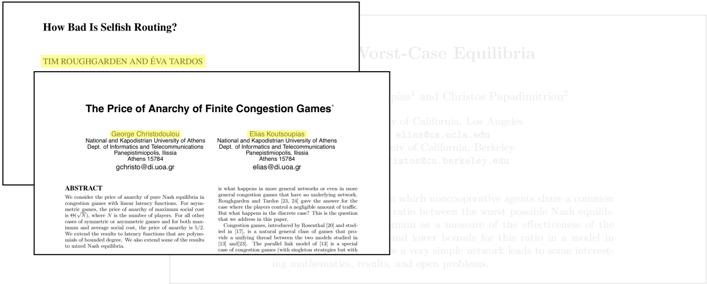
How Bad Is Selfish Routing?

TIM ROUGHGARDEN AND ÉVA TARDOS

Worst-Case Equilibria

...pias[1] and Christos Papadimitriou[2]

...y of California, Los Angeles
elias@cs.ucla.edu
...iv of California, Berkeley
...ristos@cs.berkeley.edu

...n which noncooperative agents share a common ... ratio between the worst possible Nash equilib... ...mum as a measure of the effectiveness of the ... and lower bounds for this ratio in a model in ... a very simple network leads to some interest- ing mathematics, results, and open problems.

The Price of Anarchy of Finite Congestion Games*

George Christodoulou
National and Kapodistrian University of Athens
Dept. of Informatics and Telecommunications
Panepistimiopolis, Ilissia
Athens 15784
gchristo@di.uoa.gr

Elias Koutsoupias
National and Kapodistrian University of Athens
Dept. of Informatics and Telecommunications
Panepistimiopolis, Ilissia
Athens 15784
elias@di.uoa.gr

ABSTRACT

We consider the price of anarchy of pure Nash equilibria in congestion games with linear latency functions. For asymmetric games, the price of anarchy of maximum social cost is $\Theta(\sqrt{N})$, where N is the number of players. For all other cases of symmetric or asymmetric games and for both maximum and average social cost, the price of anarchy is 5/2. We extend the results to latency functions that are polynomials of bounded degree. We also extend some of the results to mixed Nash equilibria.

is what happens in more general networks or even in more general congestion games that have no underlying network. Roughgarden and Tardos [23, 24] gave the answer for the case where the players control a negligible amount of traffic. But what happens in the discrete case? This is the question that we address in this paper.

Congestion games, introduced by Rosenthal [20] and studied in [17], is a natural general class of games that provide a unifying thread between the two models studied in [13] and[23]. The parallel link model of [13] is a special case of congestion games (with singleton strategies but with

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

History and related works

▷ **Congestion games (Rosenthal 1973)**
applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

▷ **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

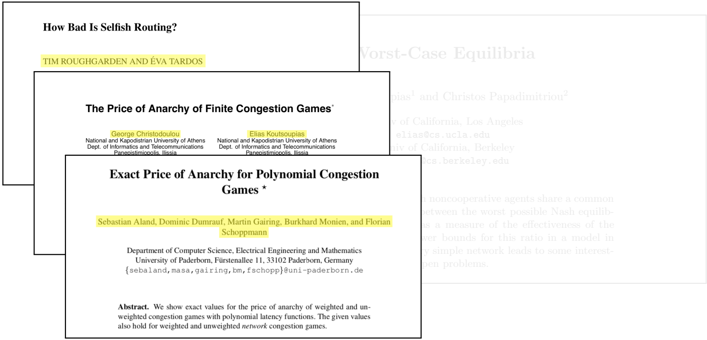
How Bad Is Selfish Routing?

TIM ROUGHGARDEN AND ÉVA TARDOS

The Price of Anarchy of Finite Congestion Games*

George Christodoulou, National and Kapodistrian University of Athens, Dept. of Informatics and Telecommunications, Panepistimiopolis, Ilissia

Elias Koutsoupias, National and Kapodistrian University of Athens, Dept. of Informatics and Telecommunications, Panepistimiopolis, Ilissia

Exact Price of Anarchy for Polynomial Congestion Games *

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{sebaland,masa,gairing,bm,fschopp}@uni-paderborn.de

Abstract. We show exact values for the price of anarchy of weighted and unweighted congestion games with polynomial latency functions. The given values also hold for weighted and unweighted *network* congestion games.

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

- **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

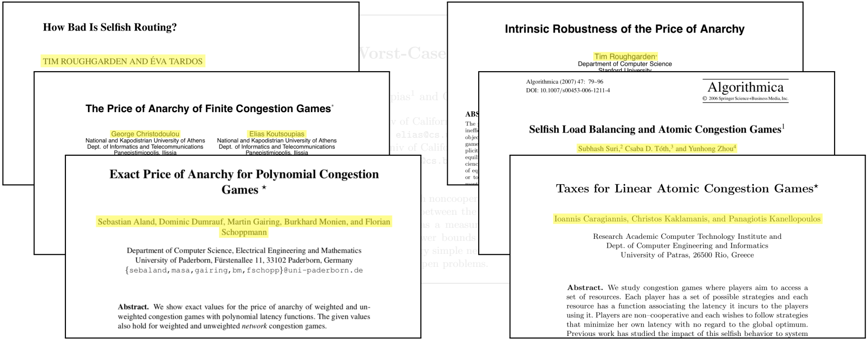

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

- **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

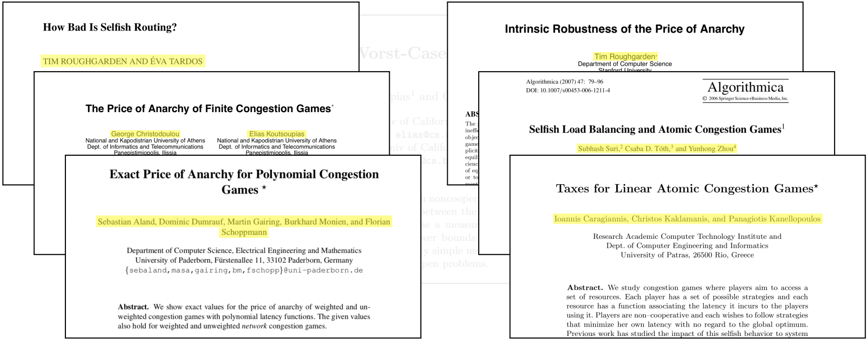

How Bad Is Selfish Routing?

TIM ROUGHGARDEN AND ÉVA TARDOS

The Price of Anarchy of Finite Congestion Games

George Christodoulou
National and Kapodistrian University of Athens
Dept. of Informatics and Telecommunications
Panepistimiopolis, Ilissia

Elias Koutsoupias
National and Kapodistrian University of Athens
Dept. of Informatics and Telecommunications
Panepistimiopolis, Ilissia

Exact Price of Anarchy for Polynomial Congestion Games ⋆

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{sebaland,masa,gairing,bm,fschopp}@uni-paderborn.de

Abstract. We show exact values for the price of anarchy of weighted and unweighted congestion games with polynomial latency functions. The given values also hold for weighted and unweighted *network* congestion games.

Intrinsic Robustness of the Price of Anarchy

Tim Roughgarden
Department of Computer Science

Algorithmica (2007) 47: 79–96
DOI: 10.1007/s00453-006-1211-4

Algorithmica
© 2006 Springer Science+Business Media, Inc.

Selfish Load Balancing and Atomic Congestion Games[1]

Subhash Suri,[2] Csaba D. Tóth,[3] and Yunhong Zhou[4]

Taxes for Linear Atomic Congestion Games⋆

Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos

Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece

Abstract. We study congestion games where players aim to access a set of resources. Each player has a set of possible strategies and each resource has a function associating the latency it incurs to the players using it. Players are non–cooperative and each wishes to follow strategies that minimize her own latency with no regard to the global optimum. Previous work has studied the impact of this selfish behavior to system

$\geq$ 6000 citations

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

History and related works

- **Congestion games (Rosenthal 1973)**
  applied to: road-traffic, electricity markets, load balancing, network design, sensor allocation, wireless data networks

- **Price of anarchy (Koutsoupias, Papadimitriou 1999)**

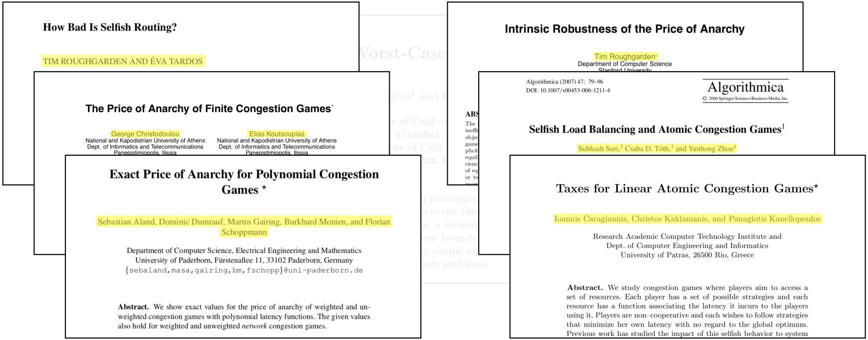

How Bad Is Selfish Routing?

TIM ROUGHGARDEN AND ÉVA TARDOS

Intrinsic Robustness of the Price of Anarchy

Tim Roughgarden
Department of Computer Science

The Price of Anarchy of Finite Congestion Games

George Christodoulou, National and Kapodistrian University of Athens, Dept. of Informatics and Telecommunications, Panepistimiopolis, Ilissia

Elias Koutsoupias, National and Kapodistrian University of Athens, Dept. of Informatics and Telecommunications, Panepistimiopolis, Ilissia

Algorithmica (2007) 47: 79–96
DOI: 10.1007/s00453-006-1211-4

Algorithmica
© 2006 Springer Science+Business Media, Inc.

Selfish Load Balancing and Atomic Congestion Games[1]

Subhash Suri,[2] Csaba D. Tóth,[3] and Yunhong Zhou[4]

Exact Price of Anarchy for Polynomial Congestion Games ⋆

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann

Department of Computer Science, Electrical Engineering and Mathematics
University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany
{sebaland,masa,gairing,bm,fschopp}@uni-paderborn.de

Abstract. We show exact values for the price of anarchy of weighted and unweighted congestion games with polynomial latency functions. The given values also hold for weighted and unweighted *network* congestion games.

Taxes for Linear Atomic Congestion Games⋆

Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos

Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece

Abstract. We study congestion games where players aim to access a set of resources. Each player has a set of possible strategies and each resource has a function associating the latency it incurs to the players using it. Players are non–cooperative and each wishes to follow strategies that minimize her own latency with no regard to the global optimum. Previous work has studied the impact of this selfish behavior to system

$\geq$ 6000 citations

⇝ quantification: Papadimitriou, Tardos, Roughgarden, Nisan, Suri, Vazirani, Stier-Moses, Anshelevich, Christodoulou, Aland, Gairing, ...

⇝ optimization: Wierman, Roughgarden, Marden, Caragiannis, Gairing, Biló, ...

Preview of the results

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

d	Untolled [1, 2, 3, ...]			
1	2.50			
2	9.58			
3	41.54			
4	267.64			
5	1513.57			

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

d	Untolled [1, 2, 3, . . .]			
1	2.50			
2	9.58			
3	41.54			
4	267.64			
5	1513.57			

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

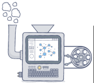

d	Untolled [1, 2, 3, ...]	Lower bound full info [4, 5]		
1	2.50	2		
2	9.58	5		
3	41.54	15		
4	267.64	52		
5	1513.57	203		

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15
[4] Caragiannis, Kaklamanis, Kanellopoulos, ESA 06; [5] Bilò, Vinci, EC 16 and TEAC 19

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

d	Untolled [1, 2, 3, ...]	Lower bound full info [4, 5] 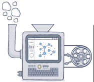	Optimal toll local info 
1	2.50	2	2.01
2	9.58	5	5.10
3	41.54	15	15.55
4	267.64	52	55.45
5	1513.57	203	220.40

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15
[4] Caragiannis, Kaklamanis, Kanellopoulos, ESA 06; [5] Bilò, Vinci, EC 16 and TEAC 19

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

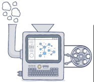

d	Untolled [1, 2, 3, ...]	Lower bound full info [4, 5]	Optimal toll local info
1	2.50	2	2.01
2	9.58	5	5.10
3	41.54	15	15.55
4	267.64	52	55.45
5	1513.57	203	220.40

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results

⇝ Tolls based on local info ≈ tolls with full info

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15
[4] Caragiannis, Kaklamanis, Kanellopoulos, ESA 06; [5] Bilò, Vinci, EC 16 and TEAC 19

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

d	Untolled [1, 2, 3, ...]	Lower bound full info [4, 5] 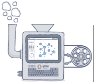	Optimal toll local info 	Optimal toll local info & constant
1	2.50	2	2.01	2.15
2	9.58	5	5.10	5.33
3	41.54	15	15.55	18.36
4	267.64	52	55.45	89.41
5	1513.57	203	220.40	469.74

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results

⇝ Tolls based on local info ≈ tolls with full info

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15
[4] Caragiannis, Kaklamanis, Kanellopoulos, ESA 06; [5] Bilò, Vinci, EC 16 and TEAC 19

Preview of the results

Main result: first solution to design of optimal tolls in congestion games (via linear programming)

▷ **Example**: prices of anarchy for polynomial latencies of degree d

d	Untolled [1, 2, 3, ...]	Lower bound full info [4, 5] 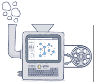	Optimal toll local info 	Optimal toll local info & constant
1	2.50	2	2.01	2.15
2	9.58	5	5.10	5.33
3	41.54	15	15.55	18.36
4	267.64	52	55.45	89.41
5	1513.57	203	220.40	469.74

⇝ Approach recovers altogether [1, 2, 3, ...] + produces novel results
⇝ Tolls based on local info ≈ tolls with full info
⇝ Tolls based on local info & constant do not lose much

[1] Christodoulou, Koutsoupias, STOC 05; [2] Aland et al., STACS 06; [3] Roughgarden, STOC09 and JACM 15
[4] Caragiannis, Kaklamanis, Kanellopoulos, ESA 06; [5] Bilò, Vinci, EC 16 and TEAC 19

How did we obtain this result?

1. Structure of optimal tolls: optimal tolls are linear

2. LP to characterize efficiency of linear tolls

3. LP to compute optimal tolls

Optimal local tolls are linear maps

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs (O_i, D_i)

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs $(\mathsf{O}_i, \mathsf{D}_i)$, any # of agents $|N| \leq n$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs (O_i, D_i), any # of agents $|N| \le n$
- any latency $\ell \in \mathcal{L}$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs (O_i, D_i), any # of agents $|N| \leq n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \geq 0\}$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs $(\mathrm{O}_i, \mathrm{D}_i)$, any # of agents $|N| \leq n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \geq 0\}$
  for given bases in $B = \{b_1(x), \ldots, b_m(x)\}$

Optimal local tolls are linear maps

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs $(\mathsf{O}_i, \mathsf{D}_i)$, any # of agents $|N| \le n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \ge 0\}$
  for given bases in $B = \{b_1(x), \ldots, b_m(x)\}$

Local tolling scheme: $\tau_e = T(\ell_e)$

Optimal local tolls are linear maps

$$\mathrm{PoA}(B, n, T) = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs $(\mathrm{O}_i, \mathrm{D}_i)$, any # of agents $|N| \leq n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \geq 0\}$
  for given bases in $B = \{b_1(x), \ldots, b_m(x)\}$

Local tolling scheme: $\tau_e = T(\ell_e)$

Optimal local tolls are linear maps

$$\mathrm{PoA}(B, n, T) = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs $(\mathrm{O}_i, \mathrm{D}_i)$, any # of agents $|N| \leq n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \geq 0\}$
  for given bases in $B = \{b_1(x), \dots, b_m(x)\}$

Local tolling scheme: $\tau_e = T(\ell_e)$

Claim: There exists a local optimal tolling T^{opt} that is linear, i.e.,

$$T^{\mathrm{opt}}(\ell_e) = T^{\mathrm{opt}}\left(\sum_j \alpha_j^e \cdot b_j\right) = \sum_j \alpha_j^e \cdot T^{\mathrm{opt}}(b_j)$$

Optimal local tolls are linear maps

$$\mathrm{PoA}(B, n, T) = \sup_{\text{set of instances}} \left(\frac{\text{total traveltime in worst equilibrium}}{\text{minimum total traveltime}} \right)$$

Set of instances:

- any graph, any pairs (O_i, D_i), any # of agents $|N| \leq n$
- any latency $\ell \in \mathcal{L} = \{\sum_{j=1}^{m} \alpha_j \cdot b_j(x), \quad \alpha_j \geq 0\}$
  for given bases in $B = \{b_1(x), \ldots, b_m(x)\}$

Local tolling scheme: $\tau_e = T(\ell_e)$

Claim: There exists a local optimal tolling T^{opt} that is linear, i.e.,

$$T^{\mathrm{opt}}(\ell_e) = T^{\mathrm{opt}}\left(\sum_j \alpha_j^e \cdot b_j\right) = \sum_j \alpha_j^e \cdot T^{\mathrm{opt}}(b_j)$$

▷ finding $T^{\mathrm{opt}}(b_j)$ is sufficient!

Computing efficiency of given linear tolls

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \dots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

s.t. $b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu\left[f_j(x+y)y - f_j(x+y+1)z\right] \geq 0$

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

$$\text{s.t. } b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu \left[f_j(x+y)y - f_j(x+y+1)z\right] \geq 0$$

$$\forall j \in \{1, \ldots, m\}$$

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

s.t. $b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu\,[f_j(x+y)y - f_j(x+y+1)z] \geq 0$

$$\forall\, j \in \{1, \ldots, m\}, \quad \forall\, (x, y, z) \in \mathbb{N}^3 \text{ with } x + y + z \leq n$$

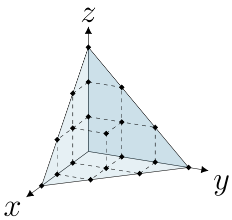

tuples (x, y, z)

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \dots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

$$\text{s.t. } b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu \left[f_j(x+y)y - f_j(x+y+1)z\right] \geq 0$$

$$\forall\, j \in \{1, \dots, m\}, \quad \forall\, (x, y, z) \in \mathbb{N}^3 \ \text{ with } \ x + y + z \leq n$$

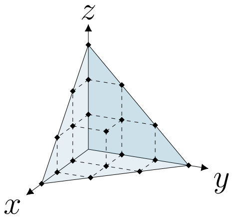

tuples (x, y, z)

▷ $|\mathcal{I}| = \mathcal{O}(n^3)$, but suffices $\mathcal{O}(n^2)$

Computing efficiency of given linear tolls

[**Paccagnan**, et al.]

Theorem: given $b_1(x), \ldots, b_m(x)$, and linear tolls T, let $f_j = b_j + T(b_j)$.

$$\mathrm{PoA}(B, n, T) = 1/C^\star$$

$$C^\star = \max_{\nu \in \mathbb{R}_{\geq 0},\, \rho \in \mathbb{R}} \rho$$

$$\text{s.t. } b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + \nu\,[f_j(x+y)y - f_j(x+y+1)z] \geq 0$$

$$\forall j \in \{1, \ldots, m\}, \quad \forall (x, y, z) \in \mathbb{N}^3 \text{ with } x + y + z \leq n$$

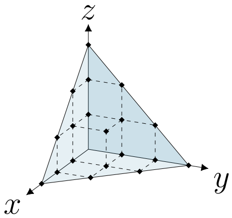

tuples (x, y, z)

▷ $|\mathcal{I}| = \mathcal{O}(n^3)$, but suffices $\mathcal{O}(n^2)$

▷ gives worst-case instance

The main idea

$$\mathrm{PoA} = \sup_{\text{set of instances}} \left(\frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}} \right)$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\text{total travel time in worst equilibrium}}{\text{minimum total travel time}} \right)$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\text{minimum total travel time}} \right)$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad C_i(e_i, e_{-i}) \leq C_i(o_i, e_{-i}) \;\; \forall i$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad C_i(e_i, e_{-i}) \leq C_i(o_i, e_{-i}) \;\; \forall i$$

2. relax the previous program

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad \sum_i C_i(e_i, e_{-i}) \leq \sum_i C_i(o_i, e_{-i})$$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad C_i(e_i, e_{-i}) \leq C_i(o_i, e_{-i}) \quad \forall i$$

2. relax the previous program

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad \sum_i C_i(e_i, e_{-i}) \leq \sum_i C_i(o_i, e_{-i})$$

3. how to parametrize an instance $G \in \tilde{\mathcal{G}}$?

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad C_i(e_i, e_{-i}) \leq C_i(o_i, e_{-i}) \quad \forall i$$

2. relax the previous program

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad \sum_i C_i(e_i, e_{-i}) \leq \sum_i C_i(o_i, e_{-i})$$

3. how to parametrize an instance $G \in \tilde{\mathcal{G}}$?
   $\mathcal{O}(n^3)$ variables suffice for $TT(e)$, $TT(o)$, $\sum_i C_i(e) - C_i(o_i, e_{-i})$

The main idea

$$\mathrm{PoA} = \sup_{G \in \mathcal{G}} \left(\frac{\max_{p \in \mathrm{NE}(G)} TT(p)}{\min_{p \in \mathcal{P}} TT(p)} \right)$$

Key idea: transform the definition of PoA itself into a linear program

1. reduce to only two allocations: optimal o and worst-equilibrium e

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad C_i(e_i, e_{-i}) \leq C_i(o_i, e_{-i}) \quad \forall i$$

2. relax the previous program

$$\mathrm{PoA} = \sup_{G \in \tilde{\mathcal{G}}} \left(\frac{TT(e)}{TT(o)} \right) \quad \text{s.t.} \quad \sum_i C_i(e_i, e_{-i}) \leq \sum_i C_i(o_i, e_{-i})$$

3. how to parametrize an instance $G \in \tilde{\mathcal{G}}$?

   $\mathcal{O}(n^3)$ variables suffice for $TT(e)$, $TT(o)$, $\sum_i C_i(e) - C_i(o_i, e_{-i})$

4. massage and take the dual □

Designing optimal tolls

Designing optimal tolls

[**Paccagnan**, et al.]

Theorem: Given $b_1(x), \ldots, b_m(x)$, an optimal local toll is given by

$$T^{\mathrm{opt}}(\ell_e) = \sum_j \alpha_j^e \cdot T^{\mathrm{opt}}(b_j) \quad \text{where} \quad T^{\mathrm{opt}}(b_j) = f_j^{\mathrm{opt}} - b_j$$

Designing optimal tolls

[**Paccagnan**, et al.]

Theorem: Given $b_1(x), \ldots, b_m(x)$, an optimal local toll is given by

$$T^{\mathrm{opt}}(\ell_e) = \sum_j \alpha_j^e \cdot T^{\mathrm{opt}}(b_j) \quad \text{where} \quad T^{\mathrm{opt}}(b_j) = f_j^{\mathrm{opt}} - b_j$$

$$f_j^{\mathrm{opt}} \in \underset{f_j \in \mathbb{R}^n,\, \rho \in \mathbb{R}}{\arg\max}\ \rho$$

$$\text{s.t.} \quad b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + f_j(x+y)y - f_j(x+y+1)z \geq 0 \quad \forall (x,y,z) \in \mathcal{I}$$

Designing optimal tolls

[**Paccagnan**, et al.]

Theorem: Given $b_1(x), \ldots, b_m(x)$, an optimal local toll is given by

$$T^{\rm opt}(\ell_e) = \sum_j \alpha_j^e \cdot T^{\rm opt}(b_j) \quad \text{where} \quad T^{\rm opt}(b_j) = f_j^{\rm opt} - b_j$$

$$f_j^{\rm opt} \in \arg\max_{f_j \in \mathbb{R}^n,\, \rho \in \mathbb{R}} \rho$$

$$\text{s.t.} \quad b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + f_j(x+y)y - f_j(x+y+1)z \geq 0 \quad \forall (x,y,z) \in \mathcal{I}$$

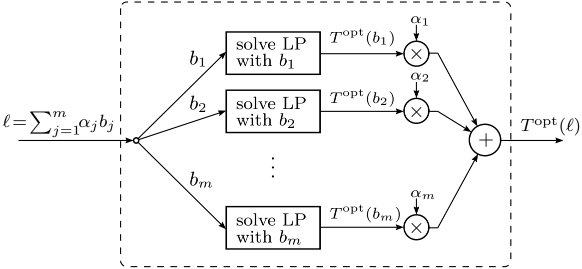

Designing optimal tolls

[**Paccagnan**, et al.]

Theorem: Given $b_1(x), \ldots, b_m(x)$, an optimal local toll is given by

$$T^{\mathrm{opt}}(\ell_e) = \sum_j \alpha_j^e \cdot T^{\mathrm{opt}}(b_j) \quad \text{where} \quad T^{\mathrm{opt}}(b_j) = f_j^{\mathrm{opt}} - b_j$$

$$f_j^{\mathrm{opt}} \in \underset{f_j \in \mathbb{R}^n,\, \rho \in \mathbb{R}}{\arg\max}\ \rho$$

$$\text{s.t.}\quad b_j(x+z)(x+z) - \rho b_j(x+y)(x+y) + f_j(x+y)y - f_j(x+y+1)z \geq 0 \quad \forall (x,y,z) \in \mathcal{I}$$

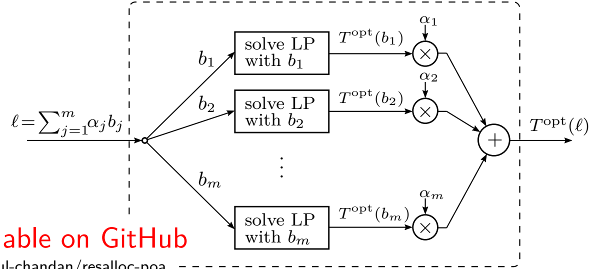

⇝ code available on GitHub

github.com/rahul-chandan/resalloc-poa

Summary of results and opportunities

Congestion-pricing as mechanism for road-traffic routing

▷ first solution to design of optimal tolls in congestion games

[TAC19a], [TAC19b] **Paccagnan**, Chandan, Marden, "Distributed resource allocation through utility design - Part I and II", TAC 19
[ArXiv19] **Paccagnan**, Chandan, Ferguson, Marden, "Incentivizing efficient use of shared infrastructure", ArXiv 19

Summary of results and opportunities

Congestion-pricing as mechanism for road-traffic routing

▷ first solution to design of optimal tolls in congestion games

▷ computing/optimizing efficiency of equilibria is a tractable LP

[TAC19a], [TAC19b] **Paccagnan**, Chandan, Marden, "Distributed resource allocation through utility design - Part I and II", TAC 19

[ArXiv19] **Paccagnan**, Chandan, Ferguson, Marden, "Incentivizing efficient use of shared infrastructure", ArXiv 19

Summary of results and opportunities

Congestion-pricing as mechanism for road-traffic routing

▷ first solution to design of optimal tolls in congestion games

▷ computing/optimizing efficiency of equilibria is a tractable LP

Consequences:

⇝ recovers/generalizes existing results altogether

⇝ local information as good as full information

⇝ simple mechanisms do not loose performance

[TAC19a], [TAC19b] **Paccagnan**, Chandan, Marden, "Distributed resource allocation through utility design - Part I and II", TAC 19

[ArXiv19] **Paccagnan**, Chandan, Ferguson, Marden, "Incentivizing efficient use of shared infrastructure", ArXiv 19

Summary of results and opportunities

Congestion-pricing as mechanism for road-traffic routing

▷ first solution to design of optimal tolls in congestion games

▷ computing/optimizing efficiency of equilibria is a tractable LP

Consequences:

⇝ recovers/generalizes existing results altogether

⇝ local information as good as full information

⇝ simple mechanisms do not loose performance

Opportunities: novel framework is still unexplored

[TAC19a], [TAC19b] **Paccagnan**, Chandan, Marden, "Distributed resource allocation through utility design - Part I and II", TAC 19

[ArXiv19] **Paccagnan**, Chandan, Ferguson, Marden, "Incentivizing efficient use of shared infrastructure", ArXiv 19

Summary of results and opportunities

Congestion-pricing as mechanism for road-traffic routing

- ▷ first solution to design of optimal tolls in congestion games
- ▷ computing/optimizing efficiency of equilibria is a tractable LP

Consequences:

⇝ recovers/generalizes existing results altogether
⇝ local information as good as full information
⇝ simple mechanisms do not loose performance

Opportunities: novel framework is still unexplored

- constraints on tolls
- carrots vs sticks
- knowledge on the latency functions
- ...

[TAC19a], [TAC19b] **Paccagnan**, Chandan, Marden, "Distributed resource allocation through utility design - Part I and II", TAC 19
[ArXiv19] **Paccagnan**, Chandan, Ferguson, Marden, "Incentivizing efficient use of shared infrastructure", ArXiv 19