Algoritmi Distribuiti e Reti Complesse (modulo II)

Luciano Gualà
guala@mat.uniroma2.it
www.mat.uniroma2.it/~guala

Algorithmic Game Theory

Algorithmic Issues in Non-cooperative (i.e., strategic) Distributed Systems

Two Research Traditions

- Theory of Algorithms: computational issues
  - What can be feasibly computed?
  - How much does it take to compute a solution?
  - Which is the quality of a computed solution?
  - Centralized or distributed computational models
- Game Theory: interaction between self-interested individuals
  - What is the outcome of the interaction?
  - Which social goals are compatible with selfishness?

one of the foremost mathematicians of the 20th century

John von Neumann
(1903-1957)

Games and Economic Behavior
(1944, with O. Morgenstern)

Different Assumptions

- Theory of Algorithms (in distributed systems):
  - Processors are *obedient*, *faulty*, or *adversarial*
  - Large systems, limited computational resources

- Game Theory:
  - Players are *strategic* (selfish)
  - Small systems, unlimited computational resources

The Internet World

- Agents often autonomous (users)
  - Users have their own individual goals
  - Network components owned by providers
- Often involve "Internet" scales
  - Massive systems
  - Limited communication/computational resources

⇨ *Both strategic and computational issues!*

Algorithmic Game Theory = Theory of Algorithms + Game Theory

- Game Theory provides a bunch of tools useful for addressing computational problems in non-cooperative scenarios
  - networks used by self-interested users
- Theory of algorithms sheds light on results of Game Theory
  - for several results on the existence of equilibria/mechanisms we have that such an equilibium/mechanism cannot be found/implemented efficiently

Basics of Game Theory

- A game consists of:
  - A set of players
  - A set of rules of encounter: Who should act when, and what are the possible actions (strategies)
  - A specification of payoffs for each combination of strategies
  - A set of outcomes
- **Game Theory** attempts to predict the outcome of the game (**solution**) by taking into account the individual behavior of the players (**agents**)

Equilibrium

- Among the possible outcomes of a game, equilibria play a fundamental role.
- Informally, an equilibrium is a strategy combination in which individuals are **not willing** to change their state.
- When a player does not want to change his state? In the Homo Economicus model, when he has selected a strategy that maximizes his *individual* payoff, knowing that other players are also doing the same.

FIRST PART:
(Nash) Equilibria

A famous one-shot game: the Prisoner's Dilemma

…the story of two strange and dangerous fellows…

A famous one-shot game: the Prisoner's Dilemma

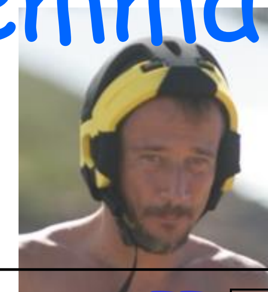

Prisoner I

Prisoner II

		Prisoner II: Don't Implicate	Prisoner II: Implicate
Prisoner I	Don't Implicate	2, 2	5, 1
	Implicate	1, 5	4, 4

Strategy Set (Prisoner II)

Payoffs

Strategy Set (Prisoner I)

Prisoner I's decision

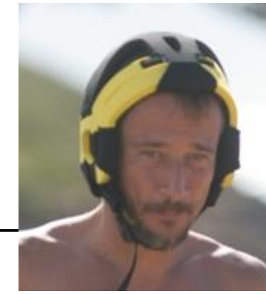

		Prisoner II	
		Don't Implicate	Implicate
Prisoner I	Don't Implicate	2, 2	5, 1
	Implicate	1, 5	4, 4

- Prisoner I's decision:
  - If II chooses Don't Implicate then it is best to Implicate
  - If II chooses Implicate then it is best to Implicate
  - It is best to Implicate for I, **regardless** of what II does: *Dominant Strategy*

Prisoner II's decision

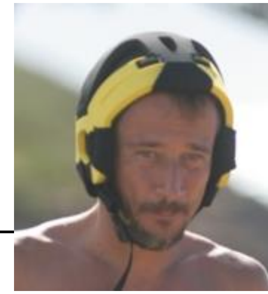

		Prisoner II	
		Don't Implicate	Implicate
Prisoner I	Don't Implicate	2, 2	5, 1
	Implicate	1, 5	4, 4

- Prisoner II's decision:
  - If I chooses Don't Implicate then it is best to Implicate
  - If I chooses Implicate then it is best to Implicate
  - It is best to Implicate for II, **regardless** of what I does: *Dominant Strategy*

Hence…

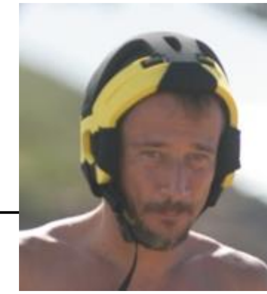

Prisoner I		Prisoner II: Don't Implicate	Prisoner II: Implicate
	Don't Implicate	2, 2	5, 1
	Implicate	1, 5	4, 4

- It is best for both to implicate **regardless** of what the other one does
- **Implicate** is a Dominant Strategy for both
- (Implicate, Implicate) becomes the **Dominant Strategy Equilibrium**
- Note: If they might collude, then it's beneficial for both to Not Implicate, but it's not an equilibrium as both have incentive to deviate

A network game

C, S: peering points

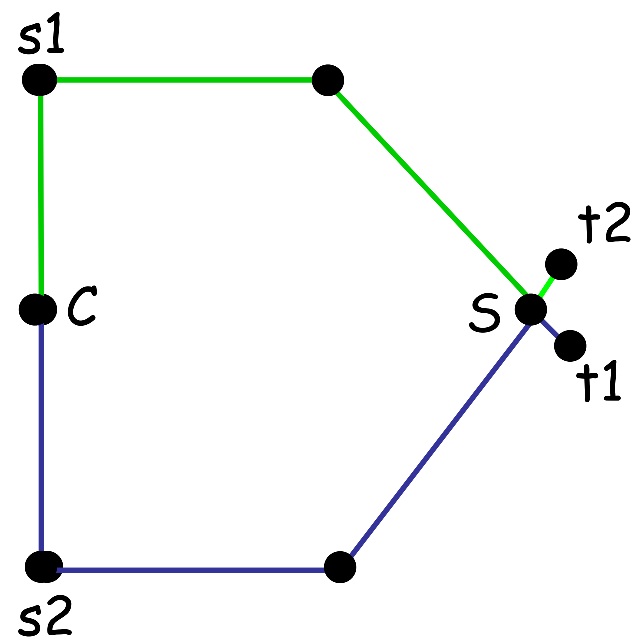

two Internet Service Providers (ISP):
ISP1 e ISP2

ISP1 wants to send traffic from s1 to t1

ISP2 wants to send traffic from s2 to t2

(long) links have cost 1
(for ISP owning the link)

Each ISPi can use two paths: the one passing through C o the one passing through S

A network game

C, S: peering points

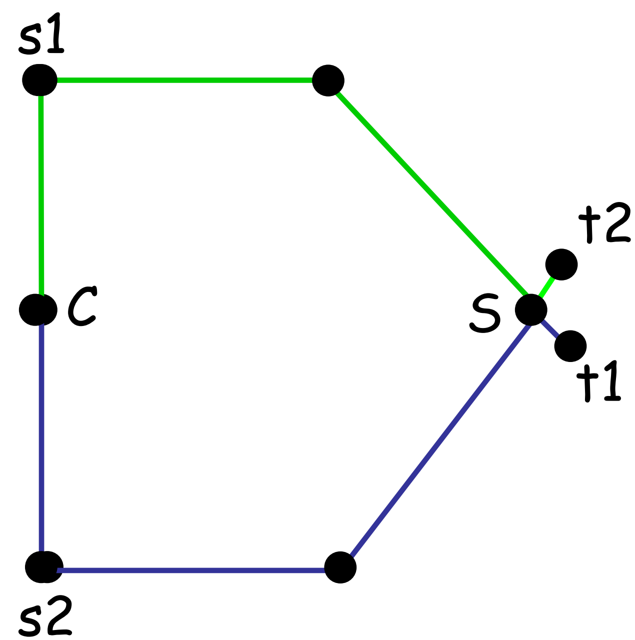

Cost Matrix

		ISP2 through S	ISP2 through C
ISP1	through S	2, 2	5, 1
	through C	1, 5	4, 4

Formal representation of a game: Normal Form

- N **rational** players
- S_i =Strategy set of player i
- The strategy combination $(s_1, s_2, \ldots, s_N)$ gives payoff $(p_1, p_2, \ldots, p_N)$ to the N players

$$\Rightarrow \quad S_1 \times S_2 \times \ldots \times S_N \text{ payoff matrix}$$

Dominant Strategy Equilibrium

- **Dominant Strategy Equilibrium:** is a strategy combination $s^* = (s_1^*, s_2^*, \ldots, s_N^*)$, such that s_i^* is a **dominant strategy** for each i, namely, for any possible alternative strategy profile $s = (s_1, s_2, \ldots, s_i, \ldots, s_N)$:
  - if p_i is a utility, then $p_i(s_1, s_2, \ldots, s_i^*, \ldots, s_N) \geq p_i(s_1, s_2, \ldots, s_i, \ldots, s_N)$
  - if p_i is a cost, then $p_i(s_1, s_2, \ldots, s_i^*, \ldots, s_N) \leq p_i(s_1, s_2, \ldots, s_i, \ldots, s_N)$
- Dominant Strategy is the *best* response to *any* strategy of other players
- If a game has a DSE, then players will immediately converge to it
- Of course, **not all** games (only **very few** in the practice!) have a dominant strategy equilibrium

A more relaxed solution concept: Nash Equilibrium [1951]

- **Nash Equilibrium:** is a strategy combination $s^* = (s_1^*, s_2^*, \ldots, s_N^*)$ such that for each i, s_i^* is a best response to $(s_1^*, \ldots, s_{i-1}^*, s_{i+1}^*, \ldots, s_N^*)$, namely, for any possible alternative strategy s_i of player i
  - if p_i is a utility, then $p_i(s_1^*, s_2^*, \ldots, s_i^*, \ldots, s_N^*) \geq p_i(s_1^*, s_2^*, \ldots, s_i, \ldots, s_N^*)$
  - if p_i is a cost, then $p_i(s_1^*, s_2^*, \ldots, s_i^*, \ldots, s_N^*) \leq p_i(s_1^*, s_2^*, \ldots, s_i, \ldots, s_N^*)$

Nash Equilibrium

- In a NE no agent can unilaterally deviate from its strategy given others' strategies as fixed
- Agent has to deliberate about the strategies of the other agents
- If the game is played **repeatedly** and players converge to a solution, then it has to be a NE
- Dominant Strategy Equilibrium $\Rightarrow$ Nash Equilibrium (but the converse is not true)

Nash Equilibrium: The Battle of the Sexes (coordination game)

		Woman: Stadium	Woman: Cinema
Man	Stadium	(6, 5)	2, 2
	Cinema	1, 1	(5, 6)

- (Stadium, Stadium) is a NE: Best responses to each other
- (Cinema, Cinema) is a NE: Best responses to each other

☹ but they are not Dominant Strategy Equilibria … are we really sure they will eventually go out together????

A similar game: routing congestion game

two traffic streams originated at node O need to be routed to the rest of the network

Costs without congestion:

c(O,A)=1 c(O,B)=2

Costs with congestion:

c(O,A)=5 c(O,B)=6

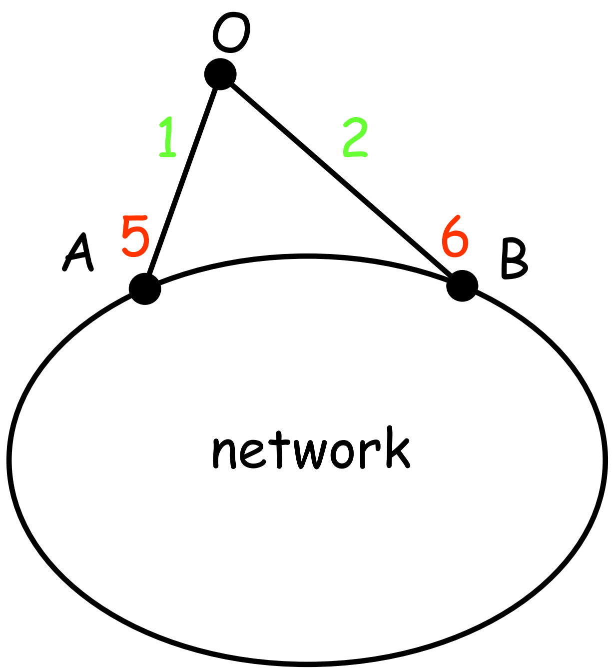

Each stream can use two paths: the one passing through A o the one passing through B

A similar game: routing congestion game

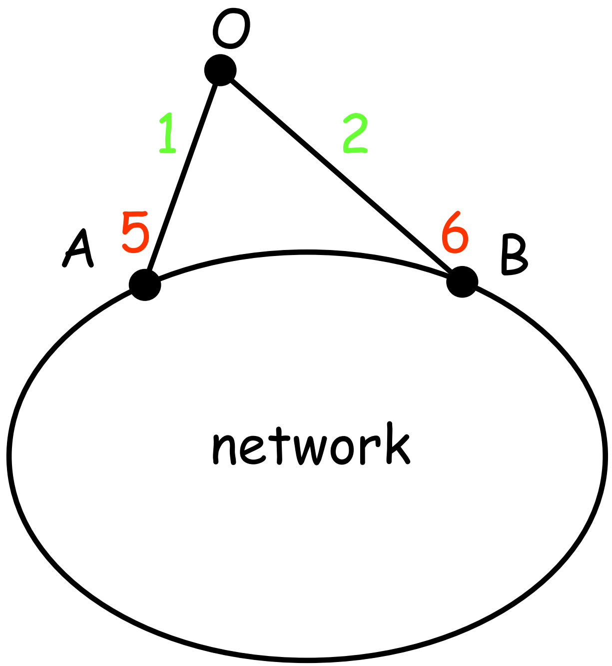

Cost Matrix

		stream 2	
		through A	through B
stream 1	through A	5, 5	1, 2
	through B	2, 1	6, 6

A big game theoretic issue: the existence of a NE

- Unfortunately, for pure strategies games (as those seen so far), it is easy to see that we cannot have a general result of existence
- In other words, there may be no, one, or many NE, depending on the game

A conflictual game: Matching pennies

		Player II	
		Head	Tail
Player I	Head	1,-1	-1,1
	Tail	-1,1	1,-1

- Player I (row) prefers to do what Player II does, while Player II prefer to do the **opposite** of what Player I does!

⇒ In any configuration, one of the players prefers to change his strategy, and so on and so forth…thus, there are no NE!

On the existence of a NE (2)

- However, when a player can select his strategy randomly by using a probability distribution over his set of possible strategies (mixed strategy), then the following general result holds:
- **Theorem (Nash, 1951):** *Any game with a finite set of players and a finite set of strategies has a NE of mixed strategies (i.e., the expected payoff cannot be improved by changing unilaterally the selected probability distribution).*
- **Head or Tail game**: if each player sets p(Head)=p(Tail)=1/2, then the expected payoff of each player is 0, and this is a NE, since no player can improve on this by choosing a different randomization!

Three big computational issues

1. **Finding** a NE, once it does exist
2. Establishing the **quality** of a NE, as compared to a cooperative system, i.e., a system in which agents can cooperate
3. In a repeated game, establishing whether and in how many steps the system will eventually **converge** to a NE

On the quality of a NE

- How inefficient is a NE in comparison to an ideal situation in which the players would strive to collaborate with the common goal of choosing the best outcome? Best outcome w.r.t. what?

- we need a social-choice function C mapping strategy profiles into real numbers
  - C measures the overall quality of an outcome S
  - e.g. C(S): sum of all players' costs/utilities

A worst-case perspective: the Price of Anarchy (PoA)

- **Definition (Koutsopias & Papadimitriou, 1999):** Given a game G and a social-choice function C, let S be the set of all NE. If the payoff represents a cost (resp., a utility) for a player, let OPT be the outcome of G minimizing (resp., maximizing) C. Then, the *Price of Anarchy* (PoA) of G w.r.t. C is

$$PoA_G(C) = \sup_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \left(\text{resp.}, \inf_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \right)$$

The price of stability (PoS)

- **Definition** **(Schulz & Moses, 2003)**: Given a game G and a social-choice function C, let S be the set of all NE. If the payoff represents a cost (resp., a utility) for a player, let OPT be the outcome of G minimizing (resp., maximizing) C. Then, the *Price of Stability (PoS)* of G w.r.t. C is:

$$PoS_{\mathrm{G}}(C) = \inf_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \left(\mathrm{resp.}, \sup_{s \in S} \frac{C(s)}{C(\mathrm{OPT})} \right)$$

Some remarks

- PoA and PoS are
  - ≥ 1 for minimization problems
  - ≤ 1 for maximization problems
- PoA and PoS are small when they are close to 1
- PoS is at least as close to 1 than PoA
- In a game with a unique equilibrium PoA=PoS
- PoA is similar to the concept of approximation ratio of a heuristic
- a bound on the PoS provides a significantly weaker guarantee than a bound on the PoA
- Why to study the PoS?
  - sometimes a nontrivial bound is possible only for PoS
  - PoS quantifies the necessary degradation in quality under the game-theoretic constraint of stability

An example: Selfih Routing


Which way?
The fastest one!


Which way? The fastest one!
Which way? The fastest one!
Which way? The fastest one!
Which way? The fastest one!
Which way? The fastest one!

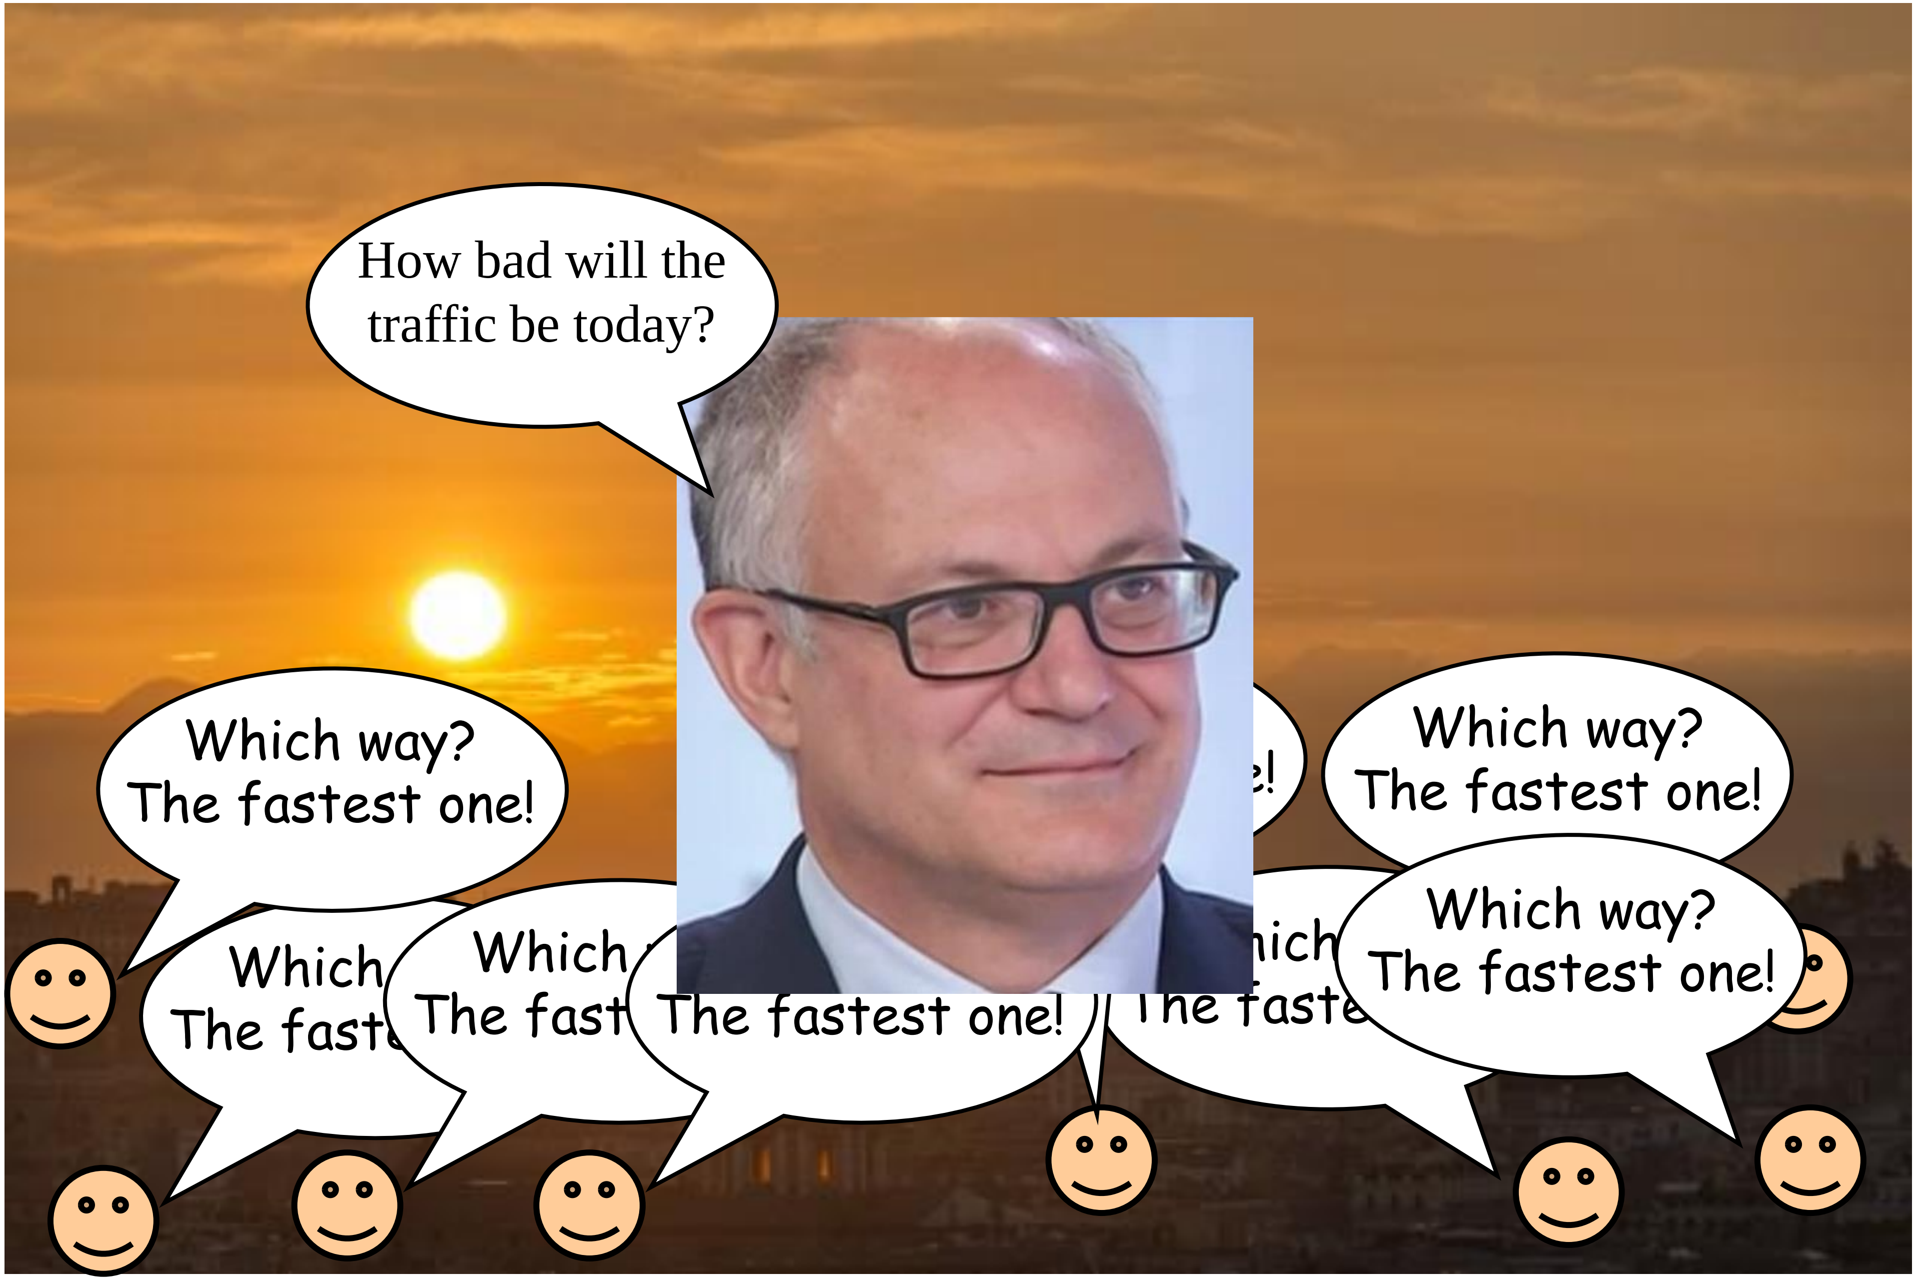
How bad will the traffic be today?
Which way? The fastest one!
The fastest one!
Which way? The fastest one!
Which way? The fastest one!

selfish routing

A large network can be modelled by using game theory

players	$\longrightarrow$	users
strategies	$\longrightarrow$	paths over which users can route their traffic

Non-atomic Selfish Routing:

- there is a large number of (selfish) users
- every user controls a tiny fraction of the traffic
- each edge has a cost function measuring the travel time as function of amount of traffic on the edge
- every user tries to minimize his travel time
- social-choice function (to minimize): average travel time incurred by players

Example: Pigou's game [1920]

One unit of traffic

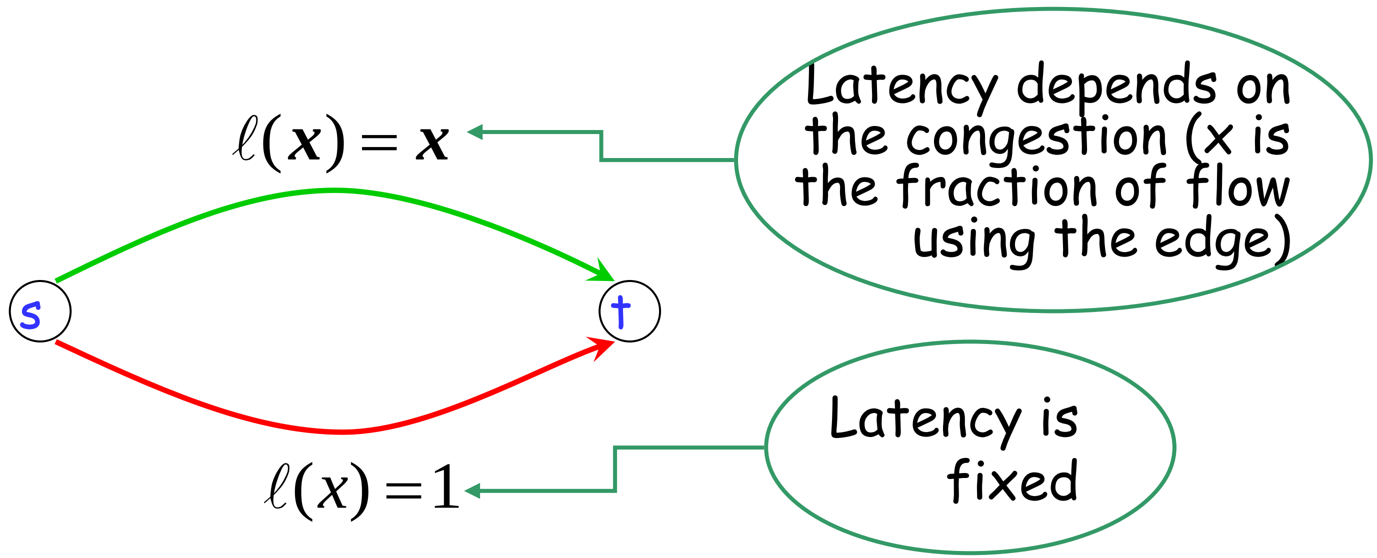

- What is the NE of this game?
- Trivial: all the fraction of flow tends to travel on the upper edge ⇒ the cost of the flow is 1·1 +0·1 =1
- How bad is this NE?
- The optimal solution is the minimum of C(x)=x·x +(1-x)·1 ⇒ C '(x)=2x-1 ⇒ OPT=1/2 ⇒ C(OPT)=1/2·1/2+(1-1/2)·1=0.75

ratio between the two costs (NE vs Opt) = 1/0.75 = **4/3**

Do we have to take into account selfish behaviour of the users when we design a network?

The Braess's paradox

One unit of traffic

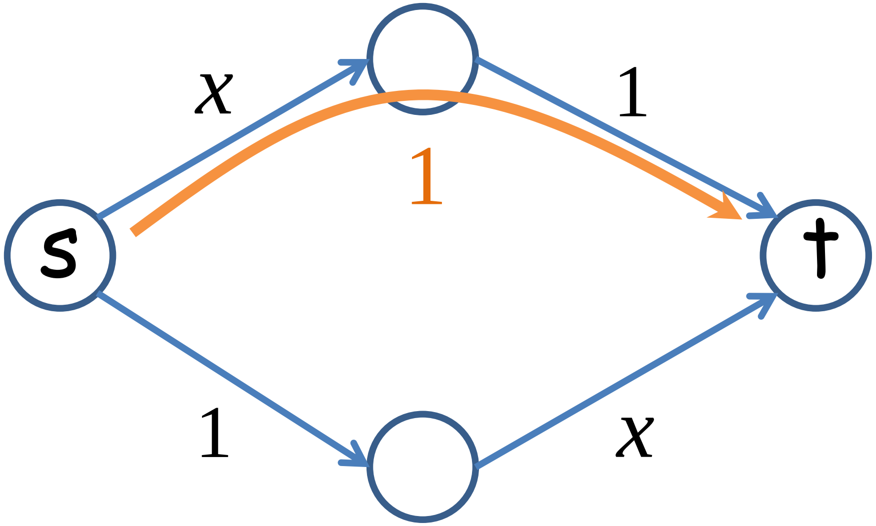

average travel time = 2

is it a NE?

...no!

The Braess's paradox

One unit of traffic

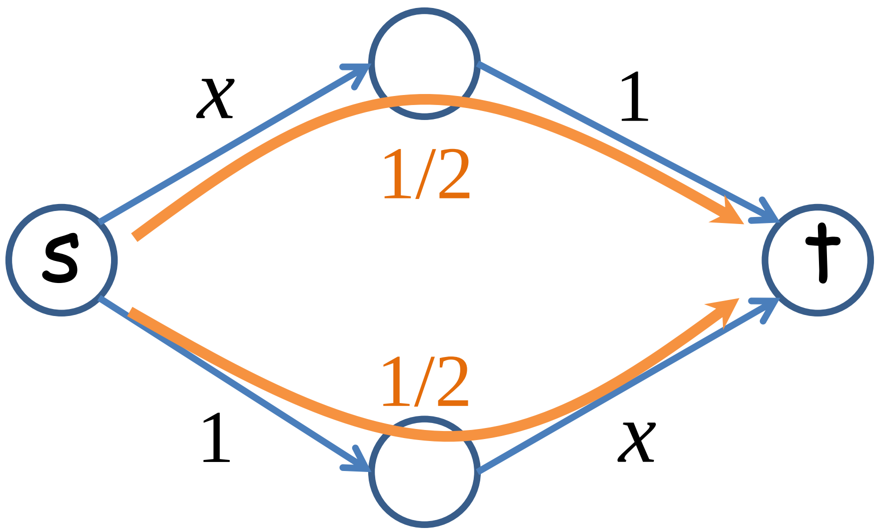

average travel time = 1.5

the only NE

Notice: this is also the optimal outcome.

The Braess's paradox

To reduce the traffic, I will build a new road.

average travel time = 2

One unit of traffic

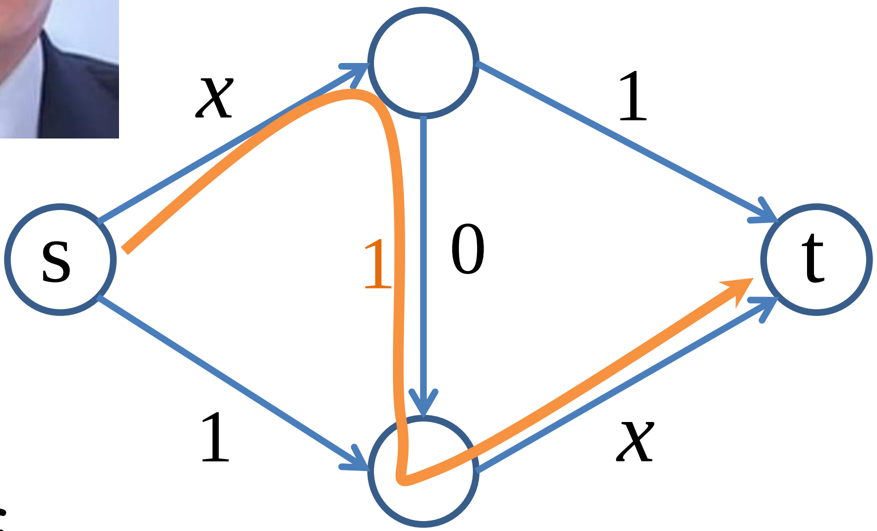

the only NE

The Braess's paradox

One unit of traffic

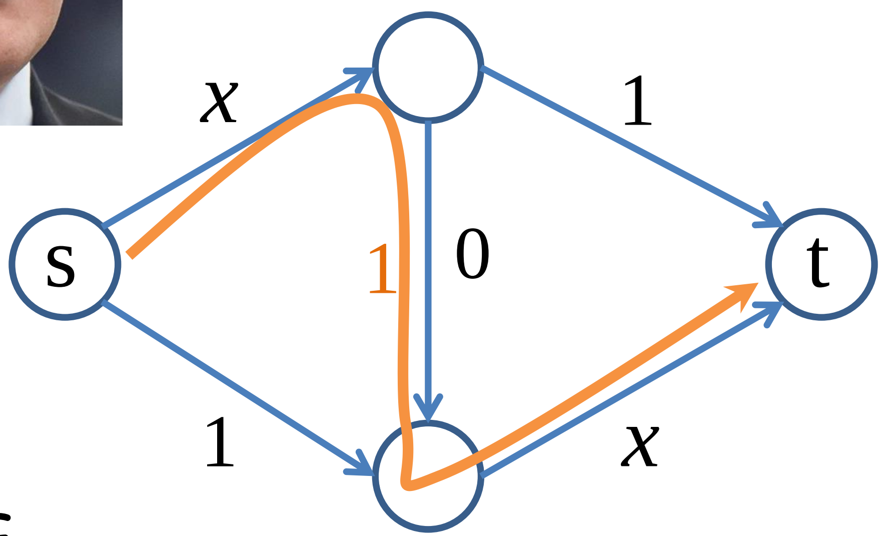

the only NE

average travel time = 2

Notice:

- the optimal outcome as before:
- ½ up& ½ down
- average travel time of 1.5

4/3 as in the Pigou's example

ratio between the two costs (NE vs Opt) = 2/1.5 = **4/3**

Theorem (Roughgarden&Tardos 2000)

The Price of Anarchy of the Selfish Routing Game with linear latency function is at most 4/3

Pollution game

There are *n* countries. Each country faces the choice of either passing legislation to control pollution or not. Assume that pollution control has a cost of 3 for the country, but each country that pollutes adds 1 of all countries (in term of added health costs).
The cost of controlling pollution is 3.

...notice that the cost of controlling pollution is considerably larger than the cost a country pays for being socially irresponsible...

can we bound the PoA?
And the PoS?

Tragedy of commons

There are *n* players. Each player wants to send information along a shared channel of known maximum capacity 1. Player i's strategy is to send x_i units of flow along the channel, for some $x_i \in [0,1]$.
Each player would like to have a large fraction of the bandwidth but the quality of the channel deteriorates as the total assigned bandwidth increases. More precisely, the value of a player i is $x_i(1- \Sigma_j x_j)$.

can we bound the PoA?
And the PoS?