



Selfish Routing Congestion Games

Selfish Routing – Motivation

Many agents want to use shared resources

Each of them is selfish and rational
(i.e. maximizes his profit)



Examples: Users of a computer network, drivers on roads

How they are going to behave?

How much is lost by letting agents behave selfishly on their own?

Example: Routing in Computer Networks

Imagine a computer network, i.e., computers connected by links.

There are several users, each user wants to route packets from a *source* computer z_i to a *target* computer t_i . For this, each user i needs to choose a path in the network from z_i to t_i .

We assume that the more agents try to route their messages through the same link, the more the link gets congested and the more costly the transmission is.

Now assume that the users are selfish and try to minimize their cost (typically transmission time). How would they behave?

Atomic Routing Games

The network routing can be formalized using an **atomic routing game** that consists of

- ▶ a directed multi-graph $G = (V, E, \delta)$,

Here V is a set of vertices, E is a set of edges, $\delta : E \rightarrow V \times V$ so that if $\delta(e) = (u, v)$ then e leads from u to v . The multigraph G models the network.

- ▶ n pairs of source-target vertices $(z_1, t_1), \dots, (z_n, t_n)$ where $z_1, \dots, z_n, t_1, \dots, t_n \in V$,


(Each pair (z_i, t_i) corresponds to a user who wants to route from z_i to t_i)

- ▶ for each $e \in E$ a cost function $c_e : \mathbb{N} \rightarrow \mathbb{R}$ such that $c_e(m)$ is the cost of routing through the link e if the amount of traffic through e is m .

Each user i chooses a simple path from z_i to t_i and pays the sum of the costs of the links on the path.

An atomic routing game is **symmetric** if $z_1 = \dots = z_n$ and $t_1 = \dots = t_n$.


Atomic Routing Games



Here we assume at most three users. Each edge is labeled by the cost if one, two, or all three users route through the edge, respectively.

Here we consider a symmetric case with three users, each has the source z and target t .

Atomic Routing Games



Here, e.g., the red user pays $3 + 2 = 5$:

- ▶ 3 for the first step from z (he shares the edge with the blue one)
- ▶ 2 for the second step to t (he is the only user of the edge)

Atomic routing games are usually studied as a special case of so called (*atomic*) *congestion games*.

Congestion Games

A *congestion game* is a tuple $G = (N, R, (S_i)_{i \in N}, (c_r)_{r \in R})$ where

- ▶ $N = \{1, \dots, n\}$ is a set of *players*,
- ▶ R is a set of *resources*,
- ▶ each $S_i \subseteq 2^R \setminus \{\emptyset\}$ is a set of *pure strategies* for player i ,
- ▶ each $c_r : \mathbb{N} \rightarrow \mathbb{R}$ is a *cost function* for a resource $r \in R$.

Notation: $S = S_1 \times \dots \times S_n$ and $c = (c_1, \dots, c_n)$.

Congestion Games

A *congestion game* is a tuple $G = (N, R, (S_i)_{i \in N}, (c_r)_{r \in R})$ where

- ▶ $N = \{1, \dots, n\}$ is a set of *players*,
- ▶ R is a set of *resources*,
- ▶ each $S_i \subseteq 2^R \setminus \{\emptyset\}$ is a set of *pure strategies* for player i ,
- ▶ each $c_r : \mathbb{N} \rightarrow \mathbb{R}$ is a *cost function* for a resource $r \in R$.

Notation: $S = S_1 \times \dots \times S_n$ and $c = (c_1, \dots, c_n)$.

Intuition:

- ▶ Each player allocates a set of resources by playing a pure strategy $s_i \subseteq R$.
- ▶ Then each player "pays" for every allocated resource $r \in s_i$ based on c_r and the number of *other* players who demand the same resource r :
 - ▶ If ℓ players use the resource r , then each of them pays $c_r(\ell)$ for this particular resource r .

Congestion Games: Payoffs and Nash Equilibria

Let $\# : R \times S \rightarrow \mathbb{N}$ be a function defined for $r \in R$ and $s = (s_1, \dots, s_n) \in S$ by $\#(r, s) = |\{i \in N \mid r \in s_i\}|$.

I.e., $\#(r, s)$ is the number of players using the resource r in the strategy profile s .

We define the payoff for player i by

$$u_i(s) = - \sum_{r \in s_i} c_r(\#(r, s)) \quad (28)$$

Intuitively, the more congested a resource $r \in s_i$ is, the more player i has to pay for it.

Congestion Games: Payoffs and Nash Equilibria

Let $\# : R \times S \rightarrow \mathbb{N}$ be a function defined for $r \in R$ and $s = (s_1, \dots, s_n) \in S$ by $\#(r, s) = |\{i \in N \mid r \in s_i\}|$.

I.e., $\#(r, s)$ is the number of players using the resource r in the strategy profile s .

We define the payoff for player i by

$$u_i(s) = - \sum_{r \in s_i} c_r(\#(r, s)) \quad (28)$$

Intuitively, the more congested a resource $r \in s_i$ is, the more player i has to pay for it.

Definition 83

Nash equilibria are defined as usual, a pure strategy profile $(s_1, \dots, s_n) \in S$ is a Nash equilibrium if for every player i and every $s'_i \in S_i$ we have $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$.

Atomic Routing Games and Congestion Games

Given an atomic routing game we may model it as a congestion game $(N, R, (S_i)_{i \in N}, (c_r)_{r \in R})$:

- ▶ Players $N = \{1, \dots, n\}$ correspond to the pairs of source-target vertices $(z_1, t_1), \dots, (z_n, t_n)$,
- ▶ resources are edges in the multigraph G , i.e, $R = E$,
- ▶ the set of pure strategies S_i of player i consists of all simple paths (i.e., sets of edges) in the multigraph G from his source z_i to his target t_i ,
- ▶ the cost function c_e of each edge $e \in E$ has to be determined according to the properties of the network.

Often (but not always) a linear (affine) function $c_e(x) = a_e x + b_e$ is used (here x is the number of players using the edge e).

Now each Nash equilibrium in G corresponds to a stable situation where no user wants to change his behavior.

Solving Congestion Games

We consider the following questions:

- ▶ Are there pure strategy Nash equilibria?
- ▶ Can the agents "learn" to use the network?
- ▶ How difficult is to compute an equilibrium?

Learning: Myopic Best-Response

Given a pure strategy profile $s = (s_1, \dots, s_n)$, suppose that some player i has an alternative strategy s'_i such that $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$. Player i can switch (unilaterally) from s_i to s'_i improving thus his payoff. Iterating such *improvement steps*, we obtain the following:

Learning: Myopic Best-Response

Given a pure strategy profile $s = (s_1, \dots, s_n)$, suppose that some player i has an alternative strategy s'_i such that $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$. Player i can switch (unilaterally) from s_i to s'_i improving thus his payoff. Iterating such *improvement steps*, we obtain the following:

Myopic best response procedure:

- ▶ Start with an arbitrary pure strategy profile $s = (s_1, \dots, s_n)$.
- ▶ While there exists a player i for whom s_i is *not a best response* to s_{-i} do
 - ▶ $s'_i :=$ a best-response by player i to s_{-i}
 - ▶ $s := (s'_i, s_{-i})$
- ▶ return s

By definition, if the myopic best response terminates, the resulting strategy profile s is a Nash equilibrium.

Learning: Myopic Best-Response

Given a pure strategy profile $s = (s_1, \dots, s_n)$, suppose that some player i has an alternative strategy s'_i such that $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$. Player i can switch (unilaterally) from s_i to s'_i improving thus his payoff. Iterating such *improvement steps*, we obtain the following:

Myopic best response procedure:

- ▶ Start with an arbitrary pure strategy profile $s = (s_1, \dots, s_n)$.
- ▶ While there exists a player i for whom s_i is *not a best response* to s_{-i} do
 - ▶ $s'_i :=$ a best-response by player i to s_{-i}
 - ▶ $s := (s'_i, s_{-i})$
- ▶ return s

By definition, if the myopic best response terminates, the resulting strategy profile s is a Nash equilibrium.

There is a strategic-form game where it does not terminate.

Learning: Myopic Best-Response

Given a pure strategy profile $s = (s_1, \dots, s_n)$, suppose that some player i has an alternative strategy s'_i such that $u_i(s'_i, s_{-i}) > u_i(s_i, s_{-i})$. Player i can switch (unilaterally) from s_i to s'_i improving thus his payoff. Iterating such *improvement steps*, we obtain the following:

Myopic best response procedure:

- ▶ Start with an arbitrary pure strategy profile $s = (s_1, \dots, s_n)$.
- ▶ While there exists a player i for whom s_i is *not a best response* to s_{-i} do
 - ▶ $s'_i :=$ a best-response by player i to s_{-i}
 - ▶ $s := (s'_i, s_{-i})$
- ▶ return s

By definition, if the myopic best response terminates, the resulting strategy profile s is a Nash equilibrium.

There is a strategic-form game where it does not terminate.

Theorem 84

For every congestion game, the myopic best response terminates in a Nash equilibrium for an arbitrary starting pure strategy profile.

Complexity of Congestion Games

For concreteness, assume $c_r(j) = a_r \cdot j + b_r$ where a_r, b_r are some non-negative constants.

Myopic best response can be used to compute Nash equilibria but how many steps it makes?

A naive bound would be the number of strategy profiles which is exponential in the number of players.

Complexity of Congestion Games

For concreteness, assume $c_r(j) = a_r \cdot j + b_r$ where a_r, b_r are some non-negative constants.

Myopic best response can be used to compute Nash equilibria but how many steps it makes?

A naive bound would be the number of strategy profiles which is exponential in the number of players.

Assume that the cost functions have values in \mathbb{N} . Then the myopic best response procedure starting in s stops after at most $\sum_{r \in R} \sum_{j=1}^{\#(r,s)} c_r(j)$ steps. This gives a pseudo-polynomial time procedure.

Complexity of Congestion Games

For concreteness, assume $c_r(j) = a_r \cdot j + b_r$ where a_r, b_r are some non-negative constants.

Myopic best response can be used to compute Nash equilibria but how many steps it makes?

A naive bound would be the number of strategy profiles which is exponential in the number of players.

Assume that the cost functions have values in \mathbb{N} . Then the myopic best response procedure starting in s stops after at most $\sum_{r \in R} \sum_{j=1}^{\#(r,s)} c_r(j)$ steps. This gives a pseudo-polynomial time procedure.

How many steps are really needed? On some instances any sequence of improvement steps to NE is of exponential length.

In fact, the problem of computing NE in congestion games is PLS-complete. PLS class (Polynomial Local Search) models the difficulty of finding a locally optimal solution to an optimization problem (e.g. travelling salesman is PLS-complete).

Complexity of Atomic Routing Games

Finding Nash equilibria in Atomic Routing Games is PLS-complete even if the cost functions are linear.

There is a polynomial time algorithm for *symmetric atomic routing games with non-decreasing cost functions* based on a reduction to the *minimum-cost flow problem*.

Here symmetric means that all players have the same source z and the same target t . Hence they also choose among the same simple paths.

Non-Atomic Selfish Routing

- ▶ So far we have considered situations where each player (user, driver) has enough "weight" to explicitly influence payoffs of others (so a deviation of one player causes changes in payoffs of other players).
- ▶ In many applications, especially in the case of highway traffic problems, individual drivers have negligible influence on each other. What matters is a "distribution" of drivers on highways.
- ▶ To model such situations we use *non-atomic routing games* that can be seen as a limiting case of atomic selfish routing with the number of players going to ∞ .

Non-Atomic Routing Games

A *Non-Atomic Routing Game* consists of

- ▶ a directed multigraph $G = (V, E, \delta)$,
- ▶ n source-target pairs $(z_1, t_1), \dots, (z_n, t_n)$,
- ▶ for each $i = 1, \dots, n$, the *amount of traffic* $\mu_i \in \mathbb{R}_{\geq 0}$ from z_i to t_i ,
- ▶ for each $e \in E$ a cost function $c_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ such that $c_e(x)$ is the cost of routing through the link e if the amount of traffic on e is $x \in \mathbb{R}_{\geq 0}$.

For $i = 1, \dots, n$, let \mathcal{P}_i be the set of all simple paths from z_i to t_i .

Intuitively, there are uncountably many players, represented by $[0, \mu_i]$, going from z_i to t_i , each player chooses his path so that his total cost is minimized.

Assume that $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$ for $i \neq j$.

(This also implies that for all $i \neq j$ we have that either $z_i \neq z_j$, or $t_i \neq t_j$.)

Denote by \mathcal{P} the set of all "relevant" simple paths $\bigcup_{i=1}^n \mathcal{P}_i$.

Question: What is a "stable" distribution of the traffic among paths of \mathcal{P} ?

Non-Atomic Routing Games

A *traffic distribution* d is a function $d : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$ such that $\sum_{p \in \mathcal{P}_i} d(p) = \mu_i$. Denote by D the set of all traffic distributions.

Let us fix a traffic distribution $d \in D$.

Given an edge $e \in E$, we denote by $g(d, e)$ the *amount of congestion on the edge e* :

$$g(d, e) = \sum_{p \in \mathcal{P} : e \in p} d(p)$$

Given $p \in \mathcal{P}$, the *payoff for players routing through $p \in \mathcal{P}$* is defined by

$$u(d, p) = - \sum_{e \in p} c_e(g(d, e))$$

Definition 85

A traffic distribution $d \in D$ is a Nash equilibrium if for every $i = 1, \dots, n$ and every path $p \in \mathcal{P}_i$ such that $d(p) > 0$ the following holds:

$$u(d, p) \geq u(d, p') \text{ for all } p' \in \mathcal{P}_i$$

Theorem 86

Every non-atomic routing game has a Nash equilibrium.

Price of Anarchy

Theorem 86

Every non-atomic routing game has a Nash equilibrium.

We define a **social cost** of a traffic distribution d by

$$C(d) = \sum_{p \in \mathcal{P}} d(p) \cdot (-u(d, p)) = \sum_{p \in \mathcal{P}} d(p) \cdot \sum_{e \in p} c_e(g(d, e))$$

Theorem 87

All Nash equilibria in non-atomic routing games have the same social cost.

Price of Anarchy

Theorem 86

Every non-atomic routing game has a Nash equilibrium.

We define a **social cost** of a traffic distribution d by

$$C(d) = \sum_{p \in \mathcal{P}} d(p) \cdot (-u(d, p)) = \sum_{p \in \mathcal{P}} d(p) \cdot \sum_{e \in p} c_e(g(d, e))$$

Theorem 87

All Nash equilibria in non-atomic routing games have the same social cost.

A **price of anarchy** is defined by

$$PoA = \frac{C(d^*)}{\min_d C(d)} \quad \text{where } d^* \text{ is a (arbitrary) Nash equilibrium}$$

Intuitively, PoA is the proportion of additional social cost that is incurred because of agents' self-interested behavior.

Price of Anarchy

Theorem 88 (Roughgarden-Tardos'2000)

For all non-atomic routing games with linear cost functions holds

$$PoA \leq \frac{4}{3}$$

and this bound is tight (e.g. the Pigou's example).

Price of Anarchy

Theorem 88 (Roughgarden-Tardos'2000)

For all non-atomic routing games with linear cost functions holds

$$PoA \leq \frac{4}{3}$$

and this bound is tight (e.g. the Pigou's example).

The price of anarchy can be defined also for atomic routing games:

$$PoA_{atom} := \frac{\max_{s^* \text{ is NE}} \sum_{i=1}^n (-u_i(s^*))}{\min_{s \in S} \sum_{i=1}^n (-u_i(s))}$$

(Intuitively, $\sum_{i=1}^n (-u_i(s))$ is the total amount paid by all players playing the strategy profile s .)

Theorem 89 (Christodoulou-Koutsoupias'2005)

For all atomic routing games with linear cost functions holds

$$PoA_{atom} \leq \frac{5}{2}$$

(which is again tight, just like $\frac{4}{3}$ for non-atomic routing.)

Braess's Paradox

For an example see the green board.

Real-world occurrences (Wikipedia):

- ▶ In Seoul, South Korea, a speeding-up in traffic around the city was seen when a motorway was removed as part of the Cheonggyecheon restoration project.
- ▶ In Stuttgart, Germany after investments into the road network in 1969, the traffic situation did not improve until a section of newly built road was closed for traffic again.
- ▶ In 1990 the closing of 42nd street in New York City reduced the amount of congestion in the area.
- ▶ In 2012, scientists at the Max Planck Institute for Dynamics and Self-Organization demonstrated through computational modeling the potential for this phenomenon to occur in power transmission networks where power generation is decentralized.
- ▶ In 2012, a team of researchers published in Physical Review Letters a paper showing that Braess paradox may occur in mesoscopic electron systems. They showed that adding a path for electrons in a nanoscopic network paradoxically reduced its conductance.