
Algorithmic Game Theory and Applications

Lecture 16: Selfish Network Routing, Congestion Games, and the Price of Anarchy

Kousha Etessami

warning, again

- In the few remaining lectures, we will briefly look at several vast subjects at the forefront of AGT research, at the intersection of CS, economics and e-commerce, and “internetology”.
- These are extremely active research topics, and there are entire conferences and workshops dedicated to them. See, e.g., the *ACM conference on Economics and Computation*.
- Supplementary reference reading: relevant chapters of the textbook *Algorithm Game Theory*, edited by N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, Cambridge Press, 2007.

For today’s lecture, Chapter 18 (Roughgarden) and Chapter 19 (Tardos and Wexler), are relevant.

- We will only scratch the surface of these subjects, so please explore them further if they interest you.

games and the internet

- Basic idea: “The internet is a huge experiment in interaction between agents (both human and automated)”.
- Such interactions can profitably be viewed from a game theoretic viewpoint: agents trying to maximize their own payoffs, etc.
- What are the implications of selfish behavior?
- How do we set up the rules of these games to harness “socially optimal” results?

(Selfish) Network Routing as a Game




Figure 1: “The Internet”

- Selfish agent $i = 1, 2, 3$, wants to route packets from source S_i to destination T_i . So, agent i must choose a **directed path** from S_i to T_i .
- The **delay** on each edge of the path is governed by the **congestion** of that edge, i.e., by the total number of agents using that edge in their path.
- Agents may want to change their choice if their path is too congested.
- What is a **Nash Equilibrium** in this game?
- What are the **welfare properties** of such NE?
(Is it socially optimal? If not, how bad can it be?)

Congestion Games ([Rosenthal, 1973])

A **Congestion Game**, $G = (N, R, (Z_i)_{i \in N}, (d_r)_{r \in R})$ has:

- A finite set $N = \{1, \dots, n\}$ of **players**.
- A finite set of $R = \{1, \dots, m\}$ of **resources**.
- For each player, i , a set $Z_i \subseteq 2^R$, of admissible **strategies** for player i . So a pure strategy, $s_i \in Z_i$ is simply a set of resources.
- Each resource $r \in R$ has a **cost function**:

$$d_r : \mathbb{N} \rightarrow \mathbb{Z}$$

Intuitively, $d_r(j)$ is the cost of using resource r if there are j agents simultaneously using r .

- For a pure strategy profile $s = (s_1, \dots, s_n) \in Z_1 \times \dots \times Z_n$, the **congestion** on resource r is:

$$n_r(s) \doteq |\{i \mid r \in s_i\}|$$

- Under strategy profile $s = (s_1, \dots, s_n)$, the **total cost** to player i is:

$$C_i(s) \doteq \sum_{r \in s_i} d_r(n_r(s))$$

- Every player, i , of course want to minimize its own (expected) cost.

Best response dynamics, and pure Nash Equilibria

In a congestion game G , for any pure strategy profile $s = (s_1, \dots, s_n)$, suppose that some player i has a better alternative strategy, $s'_i \in Z_i$, such that $C_i(s_{-i}; s'_i) < C_i(s)$.

Player i can switch (unilaterally) from s_i to s'_i . This takes us from profile s to profile (s_{-i}, s'_i) .

We call this a single (strict) improvement step.

Starting at an arbitrary pure strategy profile s , what happens if the players perform a sequence of such improvement steps?

Theorem: ([Rosenthal'73]) *In any congestion game, every sequence of strict improvement steps is necessarily finite, and terminates in a pure Nash Equilibrium.*

Thus, in particular, every congestion game has a pure strategy Nash Equilibrium.

Proof: Potential functions

Proof: Consider the following potential function:

$$\varphi(s) \doteq \sum_{r \in R} \sum_{i=1}^{n_r(s)} d_r(i) \quad (1)$$

What happens to the value of $\varphi(s)$ if player i switches unilaterally from s_i to s'_i , taking us from profile s to profile $s' := (s_{-i}; s'_i)$?

Claim: $\varphi(s) - \varphi(s') = C_i(s) - C_i(s')$.

Proof: Re-order the players in any arbitrary way, and index them as players $1, 2, \dots, n$. (In particular, a player formerly indexed i could be re-indexed as n .)

Then, for $i' \in \{1, \dots, n\}$, define

$$n_r^{(i')}(s) = |\{i \mid r \in s_i \wedge i \in \{1, \dots, i'\}\}|$$

By exchanging the order of summation in equation (1) for $\varphi(s)$, it can be seen that (check this yourself):

$$\varphi(s) = \sum_{i=1}^n \sum_{r \in s_i} d_r(n_r^{(i)}(s))$$

Now note that $n_r^{(n)}(s) = n_r(s)$. Thus

$$\sum_{r \in s_n} d_r(n_r^{(n)}(s)) = \sum_{r \in s_n} d_r(n_r(s)) = C_n(s)$$

So, if player n switches from strategy s_n to s'_n , leading us from profile s to $s' = (s_{-n}; s'_n)$, then:

$$\varphi(s) - \varphi(s') = C_n(s) - C_n(s').$$

But note that when re-ordering we could have chosen player n to be any player we want! So this holds for every player i . \square

To complete the proof of the Theorem:

Observe that every strict improvement step must decrease the value of the potential function $\varphi(s)$ by at least 1 (the costs $d_r(s)$ are all integers).

Furthermore, there are only finitely many pure strategies s , so there are finite integers:

$$a = \min_s \varphi(s) \text{ and } b = \max_s \varphi(s).$$

Thus, every improvement sequence is finite.

Finally, note that the last profile s in any improvement sequence which can not be further improved is, by definition, a pure Nash equilibrium. \square

The complexity of finding a pure NE in network congestion games

Consider a **network congestion game** where we are given a network with source-sink node pairs (S_i, T_i) , for each player i , and each player must to choose a route (path) from S_i to T_i .

Suppose the cost (**delay**) of an edge, e , under profile s , is defined to be some *linear* function:

$$d_e(n_e(s)) = \alpha_e n_e(s) + \beta_e$$

One obvious way to compute a pure NE is to perform an arbitrary improvement sequence. However, this may conceivably require many improvement steps.


Is there a better algorithm?

It turns out that it is as hard as any **polynomial local search** problem to compute a pure NE for network congestion games:

Theorem: [Fabrikant et.al.'04, Ackermann et.al.'06].
Computing a pure NE for a network congestion game is
PLS-complete,
even when all edge delay functions, d_e , are linear.

So, unfortunately, a P-time algorithm is unlikely.


A flow network game



(from [Roughgarden-Tardos'00])


- A large number n of customers in the network want to go from s to t .
- Each can either take the edge labeled 1, with “latency” 1 (delay of crossing edge), or edge labeled with latency x . Here x represents the “congestion”, given by the ratio of the number of customers that are using that edge divided by the total n .
- Assume n is very large, (basically, $n \rightarrow \infty$).
- What is the delay in Nash Equilibrium?
(NEs in such a setting yield an essentially unique average delay [Beckmann, et. al. '56].)
- What is a “globally optimal” customer routing strategy profile that minimizes average delay?
What is the globally optimal average delay?

a modified game




- What is the NE, and what is the average delay it induces?
- What is the globally optimal average delay?

a different network



- What is the NE, and what is its average delay?
- What is a globally optimal strategy profile and optimal average delay?
- What if an ambitious “network service provider” wanted to build additional “superfast” lines?

Braess's paradox



- What is the NE and its average delay?
- What is the globally optimal average delay?

social welfare and the price of anarchy

Recall that in a strategic game Γ , we may have different measures of the “social welfare” $welfare(x)$ under a particular profile of (mixed) strategies $x \in X$. For example, “utilitarian” social welfare is $welfare(x) := \sum_{i=1}^n U_i(x)$. For a game Γ , let $NE(\Gamma)$ be the set of NE’s of Γ .

For our next definition suppose $welfare(x) > 0$ for all $x \in X$. (In many games we could enforce this by, e.g., “shifting” all payoffs by an additive factor.)

A version of “the price of anarchy” can be defined as: ([Koutsoupias-Papadimitriou’98])

$$PA(\Gamma) := \frac{\max_{x \in X} welfare(x)}{\min_{x \in NE(\Gamma)} welfare(x)}$$

Thus, the “price of anarchy” is the ratio of best “global” outcome to the the worst NE outcome. Note: this ratio is ≥ 1 and larger means “worse”.

It would be comforting to establish that in various situations the “price of anarchy” isn’t too high.

price of anarchy in the flow network game

- For flow f let $welfare(f) := 1/(\text{average s-t-delay})$.
- In Braess's paradox, the price of anarchy is $4/3$: by playing the NE the average delay is 2, but playing half-and-half on the upper and lower route, the average delay is $3/2$ (and that's optimal).
- We have seen that the price of anarchy in network games can be arbitrarily high, when x^d is an edge label.
- Remarkably, [Roughgarden-Tardos'00] showed that in a more general flow network setting (where there can be multiple source-destination pairs (s_j, t_j)), as long as "congestions" labeling edges are linear functions of x , the worst-case price of anarchy is $4/3$.
- In other words, for linear latencies, the Braess's paradox example yields the worst-case scenario.

Back to atomic network congestion games

What is the price of anarchy in atomic network congestion games?

Theorem: [Christodoulou-Koutsoupias'2005]. *The price of anarchy for an NE in atomic network congestion games with linear utilities is*

$$5/2$$

(And this is tight, just like $4/3$ for non-atomic network congestion games.)