Algorithmic Game Theory, Summer 2019 | Lecture 1 (4 pages)

Introduction to Congestion Games

Thomas Kesselheim | Last Update: April 1, 2019

In this lecture, we get to know *congestion games*, which will be our running example for many concepts in game theory. Before coming to the formal definition, let us consider the following example.

We are given the following directed graph; there are three players, who each want to reach their respective destination node from their start node. Edge labels indicate the cost *each* player incurs if this edge is used by one, two, or all three players. So, if the edge label is a, b, c and the edge is used by two players, then each player has cost b for this edge.

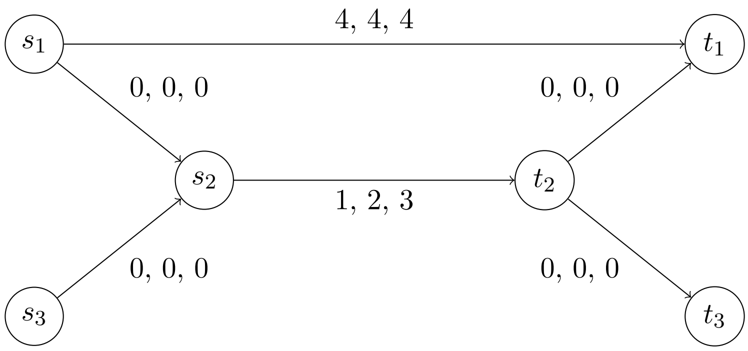

Players 2 and 3 do not have any choice, but player 1 has. He can either use the direct edge or go via s_2 and t_2. That is, we have the following two states.

State A:

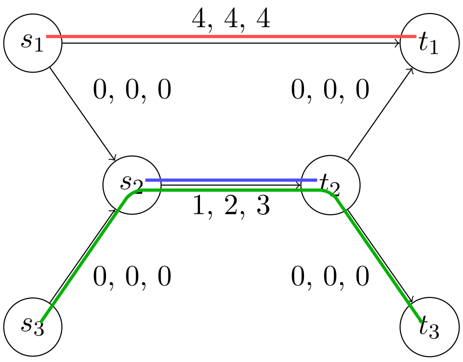

social cost: $4 + 2 + 2 = 8$

State B:

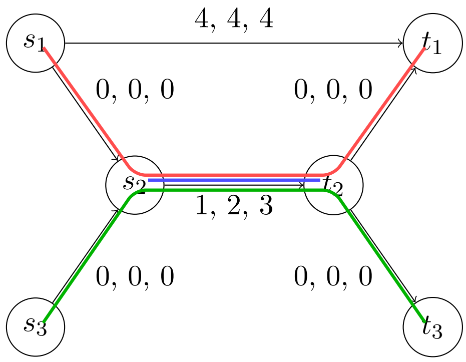

social cost: $3 + 3 + 3 = 9$

We observe that State A has a smaller *social cost* than State B. However, player 1 prefers State B because his *individual cost* is smaller there. In contrast to State A, State B is stable because every player is happy with his choice; it is an equilibrium.

We will introduce a general model that allows us to capture these effects. We will ask questions such as: Are there equilibria? How can these equilibria be found? How much performance is lost due to selfishness?

1 Formal Definition

Definition 1.1 (Congestion Game (Rosenthal 1973))**.** *A* congestion game *is a tuple* $\Gamma = (\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_r)_{r \in \mathcal{R}})$. *The set* $\mathcal{N} = \{1, \ldots, n\}$ *is a set of players; the set* $\mathcal{R}$, $|\mathcal{R}| = m$ *is a set of resources. For each player* $i \in \mathcal{N}$, $\Sigma_i \subseteq 2^{\mathcal{R}}$ *denotes the strategy space of player* i. *Every resource* $r \in \mathcal{R}$ *has delay function* $d_r \colon \{1, \ldots, n\} \to \mathbb{Z}$.

We have already seen one way to construct a congestion game by using a graph.

Example 1.2 (Network Congestion Game)**.** *In a* network *congestion game, there is a graph* $G = (V, E)$*. The resource set* $\mathcal{R}$ *corresponds to the set of edges* E*. For each player* $i \in \mathcal{N}$*, there is a dedicated source-sink pair* (s_i, t_i) *such that* Σ_i *is the set of paths from* s_i *to* t_i*.*

In particular, in the above example

$$\mathcal{N} = \{1, 2, 3\} \qquad \text{and} \qquad \mathcal{R} = \{(s_1, t_1), (s_1, s_2), (s_2, t_2), (s_3, s_2), (t_2, t_1), (t_2, t_3)\} \ .$$

Player 1's strategy set is given by $\Sigma_1 = \{\{(s_1, t_1)\}, \{(s_1, s_2), (s_2, t_2), (t_2, t_1)\}\}$*. These are two strategies: The first one uses only a single resource/edge, the second one uses three. Players 2 and 3 only have one strategy each.*

The delay function of the resource/edge (s_2, t_2) *is* $d_{(s_2,t_2)}(x) = x$ *for all* x*.*

Next, we have to add semantics by formalizing the notion of an individual player's cost.

Definition 1.3. *For any state* $S = (S_1, \ldots, S_n) \in \Sigma_1 \times \cdots \Sigma_n$*, let* $n_r(S) = |\{i \in \mathcal{N} \mid r \in S_i\}|$ *denote the number of players with* $r \in S_i$*, that is, who use resource* r *in* S*. The delay of resource* r *in state* S *is given by* $d_r(n_r(S))$*. Player* i*'s cost,* $i \in \mathcal{N}$*, is defined to be* $c_i(S) = \sum_{r \in S_i} d_r(n_r(S))$*. That is, it is the sum of delays of the resources the player uses.*

Example 1.4. *In the above example, there are two different states. We have* $n_{(s_2,t_2)}(A) = 2$ *and* $n_{(s_2,t_2)}(B) = 3$*.*

Player 1's cost can be computed as $c_1(A) = d_{(s_1,t_1)}(n_{(s_1,t_1)}(A)) = 4$ *in state* A *and* $c_1(B) = d_{(s_1,s_2)}(n_{(s_1,s_2)}(B)) + d_{(s_2,t_2)}(n_{(s_2,t_2)}(B)) + d_{(t_2,t_1)}(n_{(t_2,t_1)}(B)) = 0 + 3 + 0 = 3$*.*

Now, we are ready for the main definition. Consider a player $i \in \mathcal{N}$ and any *fixed* choice of strategies of the other players. The strategies that player i can choose from usually yield different costs. One or multiple minimize the cost. These are called best responses. A pure Nash equilibrium is a state in which each player is choosing such a best response.

Definition 1.5. *A strategy* S_i *is called a* best response *for player* $i \in \mathcal{N}$ *against a profile of strategies* $S_{-i} := (S_1, \ldots, S_{i-1}, S_{i+1}, \ldots, S_n)$ *if* $c_i(S_i, S_{-i}) \leq c_i(S'_i, S_{-i})$ *for all* $S'_i \in \Sigma_i$*. A state* $S \in \Sigma_1 \times \cdots \times \Sigma_n$ *is called a* pure Nash equilibrium *if* S_i *is a best response against the other strategies* S_{-i} *for every player* $i \in \mathcal{N}$*.*

So, in other words, a pure Nash equilibrium is a state in which no player can unilaterally decrease his cost by deviating to a different strategy. It is possible, however, that other strategies have the same cost. Also, equilibria need not be unique.

2 Existence of Pure Nash Equilibria

As our first result, we will show every congestion game has a pure Nash equilibrium. We will talk about *improvement steps.* The pair of states (S, S') is an improvement step if there is some player $i \in \mathcal{N}$ such that $c_i(S') < c_i(S)$ and $S'_{-i} = S_{-i}$.

Example 1.6. *A sequence of (best response) improvement steps:*

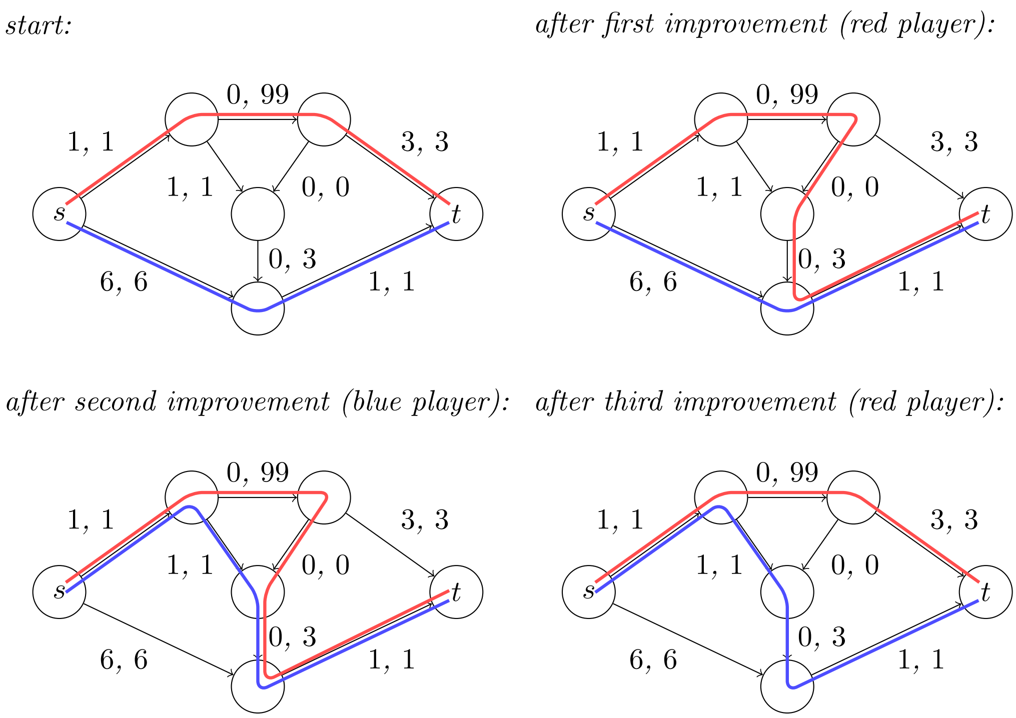

We will show the following theorem.

Theorem 1.7 (Rosenthal 1973)**.** *For every congestion game, every sequence of improvement steps is finite.*

This property is sometimes also called *finite improvement property*. It immediately implies the following corollary.

Corollary 1.8. *Every congestion game has at least one pure Nash equilibrium.*

The reason is as follows: Start from an arbitrary state $S^{(0)}$ and generate an improvement sequence $S^{(0)}, S^{(1)}, \ldots$. If there is no improvement step $(S^{(t)}, S')$, then $S^{(t)}$ is a pure Nash equilibrium. Otherwise, there is improvement step $(S^{(t)}, S')$ and we can set $S^{(t+1)} = S'$. After only finitely many steps, we have to have reached a pure Nash equilibrium, otherwise we would be generating in infinite sequence of improvement steps.

Proof of Theorem 1.7. Rosenthal's analysis is based on a potential function argument. For every state S, let

$$\Phi(S) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{n_r(S)} d_r(k) \enspace .$$

This function is called *Rosenthal's potential function.*

Lemma 1.9. *Let S be any state and let S_i' be an alternative strategy for player i. Then $\Phi(S_i', S_{-i}) - \Phi(S) = c_i(S_i', S_{-i}) - c_i(S)$.*

Proof. We give two different proofs, one is more intuitive, the other one is more algebraic and formal.

The potential $\Phi(S)$ can be calculated by inserting the players one after the other in any order, and summing the delays of the players at the point of time at their insertion.

Without loss of generality player i is the last player that we insert when calculating $\Phi(S)$. Then the potential accounted for player i corresponds to $c_i(S)$, that is, the cost of player i in state S. When calculating $\Phi(S_i', S_{-i})$, everything is the same before inserting player i. Now,

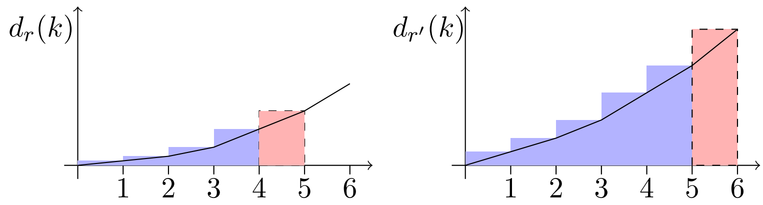

Figure 1: Proof of Lemma 1.9: The contribution of two resources r and r' to the potential is the shaded area. If a player changes from r' to r, his delay changes exactly as the potential value (difference of red areas).

the potential accounted for player i is $c_i(S'_i, S_{-i})$. So, the difference is exactly $c_i(S'_i, S_{-i}) - c_i(S)$ (see Figure 1 for an example).

For the second proof, let's observe how the potential differs. We can reorder the sum to get

$$\Phi(S'_i, S_{-i}) - \Phi(S) = \sum_{r \in \mathcal{R}} \left(\sum_{k=1}^{n_r(S'_i, S_{-i})} d_r(k) - \sum_{k=1}^{n_r(S)} d_r(k) \right) .$$

What is the value of $\Delta_r := \sum_{k=1}^{n_r(S'_i,S_{-i})} d_r(k) - \sum_{k=1}^{n_r(S)} d_r(k)$?

There are four cases. In the first case, resource r is used by player i in both S_i and S'_i. In this case, $n_r(S'_i, S_{-i}) = n_r(S)$ and $\Delta_r = 0$.

In the second case, resource r is neither used in S_i nor S'_i. Again, $n_r(S'_i, S_{-i}) = n_r(S)$ and $\Delta_r = 0$.

In the third case, $r \in S'_i \setminus S_i$. In this case $n_r(S'_i, S_{-i}) = n_r(S) + 1$, so $\Delta_r = d_r(n_r(S'_i, S_{-i}))$.

Finally, the fourth case if $r \in S_i \setminus S'_i$. Now, $n_r(S'_i, S_{-i}) = n_r(S) - 1$, and therefore $\Delta_r = -d_r(n_r(S))$.

We compare this to $c_i(S'_i, S_{-i}) - c_i(S)$, which can be simplified by reordering the sums

$$c_i(S'_i, S_{-i}) - c_i(S) = \sum_{r \in S'_i} d_r(n_r(S'_i, S_{-i})) - \sum_{r \in S_i} d_r(n_r(S)) = \sum_{r \in S'_i \setminus S_i} d_r(n_r(S')) - \sum_{r \in S_i \setminus S'_i} d_r(n_r(S)) = \sum_{r \in \mathcal{R}} \Delta_r .$$

□

The lemma shows that Φ is a so-called *exact potential*, i.e., if a single player changes its strategy, making the cost change by a value of Δ, then Φ decreases by exactly the same amount.

Further observe that

(i) the delay values are integers so that, for every improvement step, $c_i(S'_i, S_{-i}) - c_i(S) \leq -1$,

(ii) for every state S, $\Phi(S) \leq \sum_{r \in \mathcal{R}} \sum_{i=1}^{n} |d_r(i)|$,

(iii) for every state S, $\Phi(S) \geq -\sum_{r \in \mathcal{R}} \sum_{i=1}^{n} |d_r(i)|$.

Consequently, the number of improvements is upper-bounded by $2 \cdot \sum_{r \in \mathcal{R}} \sum_{i=1}^{n} |d_r(i)|$ and hence finite. □