

Lecture 2 – Congestion Games

Instructor: *Xi Chen*Scribes: *Delia Wang*

1 Overview

Last week we looked at a global network connection game, and used the potential function method to show the existence of a global potential function and a pure equilibrium for that game. This week we will broaden our focus to the generalization of this game, which is the category of the Congestion Games. These games are of interest because they model many network, routing and resource allocation scenarios. We will show that it is possible to find pure equilibria for any congestion game at a reasonable level of computational complexity, given that we are allowed a few relaxations of the problem. We will finish the lecture by studying an example of a paradoxical congestion game which has an unintuitive feature.

2 Congestion Game Definition

Definition 1. A *congestion game* is a game, which when expressed in its normal form, consists of the following elements

1. A finite set of k players $\{1, \dots, k\}$.
2. A set of resources $E = \{e_1, \dots, e_m\}$. For example, these resources could be roads or network connections.
3. For each player i , a set A_i of strategies where a strategy $S_i \in A_i$ is a subset of E . That is, a strategy is a set of resources which the player could use to fulfil his task. Continuing with the above example, a strategy might be a collection of roads which the player will use to travel between two points.
4. A strategy vector $\langle S_1, \dots, S_k \rangle$ where S_i is the strategy played by player i .
5. For every resource $e \in E$, a cost function C_e which maps $\{1, 2, \dots, k\}$ to \mathbb{N} and where the variable of C_e denotes the number of players using e in their strategy. An example of a cost function is $C_e(x) = x$, where the cost for each player using the resource grows linearly with the number of other players that are also using the resource. Such cost functions could model latency in a network, or congestion in traffic.
6. For every player i , a price function C_i defined on the set of strategy vectors: $C_i(\langle S_1, \dots, S_k \rangle) = \sum_{e \in S_i} C_e(k_e)$ where k_e is the number of players using resource e .

Note: When a congestion game is given as the input of a computational problem, we can assume that every function C_e is given in the form of a sequence of integers, and not as a mathematical expression. For example, if $C_e(x) = 2^{2^x}$ then the representation of C_e in this fashion would require exponentially many bits in x .

3 Existence of an Exact Potential Function and a Pure Equilibrium For Every Congestion Game

Recall: A potential function is any function which maps possible outcomes of the game to real numbers, such that the equilibria of the game are mapped to the local optima of the function. An exact potential function is a potential function in which the change in the function's value induced by a player changing his strategy is exactly equal to the change in that player's price function.

Theorem 2. *Every congestion game has an exact potential function.*

Proof. The function $f(S) = \sum_{e \in E} \sum_{t=1}^{k_e} C_e(t)$ is an exact potential function.

It follows that a pure equilibrium exists, because we can iteratively allow a player to choose an alternate strategies that lowers his cost until the potential function hits a local minima. When it does, the game will be in a pure equilibrium.

However, in some cases this iterative process may take an exponential number of steps. As it turns out, the problem is PLS-complete. Thus we must allow relaxations of the problem.

4 Relaxations of the Problem

We will impose the following four conditions:

1. For every resource e , its cost function C_e is non-decreasing.
2. For every resource e , the growth of its cost function C_e is bounded by a factor of α for a small constant value of α :

$$\frac{C_e(t+1)}{C_e(t)} \leq \alpha$$

This is known as the “ α -bounded jump condition”.

3. The game is symmetric across players; that is, all players have the same strategy set:

$$A_1 = A_2 = \dots = A_k$$

4. Rather than finding a true pure equilibrium we will relax the notion to “ ε -approximate pure equilibrium”, which is defined as follows: $\langle S_1^*, \dots, S_k^* \rangle$ is an ε -approximate pure equilibrium means a state in which a player can only reduce his cost by at most an ε fraction of the original cost:

$$C_i(\langle S_1^*, \dots, S'_i, \dots, S_k^* \rangle) \geq (1 - \varepsilon)C_i(\langle S_1^*, \dots, S_k^* \rangle)$$

for all players i and their strategies S'_i .

Even with the first three conditions, the problem remains PLS-complete. The critical condition that makes the problem tractable is condition 4.

5 The ‘ ε –best response dynamics’ Algorithm

To obtain an ε –approximate pure equilibrium we will use “ ε –best response dynamics”. The basic idea is the following: Start with an arbitrary strategy vector. If it is an ε –approximate pure equilibrium then we are done. Otherwise: At least one player has an alternate strategy that will lower his cost by at least an ε –fraction. Choose the player who can reduce his cost by the most, and change his strategy. That is, choose player i and strategy S_i such that we maximize the expression

$$\frac{\mathcal{C}_i(\langle S_1, \dots, S_k \rangle) - \mathcal{C}_i(\langle S_1, \dots, S'_i, \dots, S_k \rangle)}{\mathcal{C}_i(\langle S_1, \dots, S_k \rangle)}$$

We check the new strategy vector to see if it is an ε –approximate pure equilibrium, and continue in this fashion until we find one. Now we need to show that this algorithm will terminate quickly with a correct solution.

6 The Proof

Notation: $\mathbb{S}_0, \dots, \mathbb{S}_l$ will denote a sequence of states corresponding to the strategy vectors that we hit while executing the response dynamics. \mathbb{S}_0 is the state corresponding to the original arbitrary strategy vector, and \mathbb{S}_l is the state corresponding to an ε –approximate pure equilibrium.

The high level outline of the proof is as follows:

1. For every congestion game there exists a potential function f . After every iteration of the ε –best response dynamics, a player improves his outcome, thus decreasing the value of the potential function. This implies that the values of the potential function $f(\mathbb{S}_0), \dots, f(\mathbb{S}_l)$ is a decreasing sequence.
2. In Lemma 4 we will see that the chosen player contributes a large fraction of the total cost, implying that the potential function is reduced by a non-negligible fraction of a value $\beta \simeq \frac{1}{k}$ (later we see that $\beta = \frac{\varepsilon}{\alpha k}$).
3. Accordingly, after l many iterations of reducing the potential function by a β –fraction, we know that $f(\mathbb{S}_l) \leq (1 - \beta)^l f(\mathbb{S}_0)$.

Theorem 3 (The Main Theorem). *We can find an ε –approximate pure equilibrium in $\frac{\alpha k}{\varepsilon} \log\left(\frac{f_{max}}{f_{min}}\right)$ steps, where f_{max} and f_{min} are respectively the maximum and minimum values of the potential function.*

Proof of the theorem involves the following important lemma:

Lemma 4. *Let the next player chosen to change his strategy be player i . Then the cost for player i is at least as large as an α –fraction of the cost for any other player j : $\mathcal{C}_i(\mathbb{S}) \geq \frac{1}{\alpha} \mathcal{C}_j(\mathbb{S})$.*

Proof. Let the current state be \mathbb{S} . Let the state induced by allowing i to change his strategy be \mathbb{S}' . Now consider the following alternate strategy: Instead of choosing player i , we choose player j to play strategy S'_i (this is possible because of the symmetry condition we imposed). Let us denote this alternate state by \mathbb{S}'' . That is, we have:

$$\begin{aligned}
\mathbb{S} &= \langle S_1, \dots, S_i, \dots, S_j, \dots, S_k \rangle \\
\mathbb{S}' &= \langle S_1, \dots, S'_i, \dots, S_j, \dots, S_k \rangle \\
\mathbb{S}'' &= \langle S_1, \dots, S_i, \dots, S'_i, \dots, S_k \rangle
\end{aligned}$$

By the ε -best response dynamics, choosing player i maximizes the expression

$$\frac{\mathcal{C}_i(\mathbb{S}) - \mathcal{C}_i(\mathbb{S}')}{\mathcal{C}_i(\mathbb{S})}$$

and so we know that the reduction in cost in moving to \mathbb{S}'' cannot be more than the reduction in moving to \mathbb{S}' , and hence we have:

$$\begin{aligned}
\frac{\mathcal{C}_i(\mathbb{S}) - \mathcal{C}_i(\mathbb{S}')}{\mathcal{C}_i(\mathbb{S})} \geq \frac{\mathcal{C}_j(\mathbb{S}) - \mathcal{C}_j(\mathbb{S}'')}{\mathcal{C}_j(\mathbb{S})} &\implies \mathcal{C}_i(\mathbb{S})\mathcal{C}_j(\mathbb{S}) - \mathcal{C}_i(\mathbb{S}')\mathcal{C}_j(\mathbb{S}) \geq \mathcal{C}_j(\mathbb{S})\mathcal{C}_i(\mathbb{S}) - \mathcal{C}_j(\mathbb{S}'')\mathcal{C}_i(\mathbb{S}) \\
&\implies \mathcal{C}_i(\mathbb{S}')\mathcal{C}_j(\mathbb{S}) \leq \mathcal{C}_j(\mathbb{S}'')\mathcal{C}_i(\mathbb{S}) \\
&\implies \frac{\mathcal{C}_j(\mathbb{S})}{\mathcal{C}_i(\mathbb{S})} \leq \frac{\mathcal{C}_j(\mathbb{S}'')}{\mathcal{C}_i(\mathbb{S}')}
\end{aligned}$$

To prove the bound in the lemma it suffices to show that $\frac{\mathcal{C}_j(\mathbb{S}'')}{\mathcal{C}_i(\mathbb{S}')} \leq \alpha$.

Let us examine the expressions for the numerator and denominator in this ratio. Comparing the cost for player i under \mathbb{S}' with the cost for player j under \mathbb{S}'' we have:

$$\begin{aligned}
\mathcal{C}_j(\mathbb{S}'') &= \sum_{e \in \mathcal{S}_j} C_e(k''_e) \text{ where } k''_e \text{ is the number of players in } \mathbb{S}'' \text{ using resource } e \\
\mathcal{C}_i(\mathbb{S}') &= \sum_{e \in \mathcal{S}_i} C_e(k'_e) \text{ where } k'_e \text{ is the number of players in } \mathbb{S}' \text{ using resource } e \\
\frac{\mathcal{C}_j(\mathbb{S}'')}{\mathcal{C}_i(\mathbb{S}')} &= \frac{\sum_{e \in \mathcal{S}_j} C_e(k''_e)}{\sum_{e \in \mathcal{S}_i} C_e(k'_e)}
\end{aligned}$$

Notice that \mathbb{S}' and \mathbb{S}'' differ only in S_i in \mathbb{S}'' and S_j in \mathbb{S}' :

$$\begin{aligned}
\mathbb{S}' &= \langle S_1, \dots, S'_i, \dots, S_j, \dots, S_k \rangle \\
\mathbb{S}'' &= \langle S_1, \dots, S_i, \dots, S'_i, \dots, S_k \rangle
\end{aligned}$$

For any resource e , S_j can contribute its cost C_e to \mathbb{S}' zero or one times; similarly, S_i can contribute cost C_e to \mathbb{S}'' zero or one times. Thus, the difference between the number of users of e in \mathbb{S}'' and in \mathbb{S}' : $|k''_e - k'_e|$ is at most 1. From the α -bounded jump condition, we have

$$C_e(k''_e) \leq \alpha C_e(k'_e)$$

Substituting this expression into our original ratio gives:


$$\frac{\sum_{e \in S_i} C_e(k''_e)}{\sum_{e \in S_i} C_e(k'_e)} \leq \frac{\alpha \sum_{e \in S_i} C_e(k'_e)}{\sum_{e \in S_i} C_e(k'_e)} = \alpha$$

as required. \square

Proof of Main Theorem. We have from the lemma that $C_i(\mathbb{S}) \geq \frac{1}{\alpha} \cdot C_j(\mathbb{S})$ and since $\sum_j C_j(\mathbb{S}) \geq f(\mathbb{S})$ we also have $C_i(\mathbb{S}) \geq \frac{1}{\alpha k} f(\mathbb{S})$. Going from \mathbb{S} to a new state \mathbb{S}' reduces the cost by at least ε times the original cost (otherwise we would be in an ε -approximate pure equilibrium), so $f(\mathbb{S}) - f(\mathbb{S}') > \varepsilon C_i(\mathbb{S}) \geq \frac{\varepsilon}{\alpha k} f(\mathbb{S})$. The potential function is non-negative and integer-valued, so \mathbb{S} can only change at most $\frac{\alpha k}{\varepsilon} \log\left(\frac{f_{max}}{f_{min}}\right)$ times. \square

7 A Paradoxical Game Example

Consider an 100-player congestion game on the following network, where players want to build paths from S to T :



The edge cost functions are defined as follows:

$$\begin{aligned} S \rightarrow A, B \rightarrow T &: \text{always } 100 + \varepsilon \\ S \rightarrow B, A \rightarrow T &: \text{cost equals number of players using edge} \end{aligned}$$

Suppose that we place a new edge $A \leftrightarrow B$ with cost 0. One might intuitively guess that this addition will help improve the overall cost to players, but this is wrong, it actually makes things worse. In fact, with this addition, some players will do worse, and no player will improve his situation.

To see why this is, consider the equilibrium situation in the original game. In this situation, 50 players will choose the path $S \rightarrow A \rightarrow T$ and 50 players will choose the path $S \rightarrow B \rightarrow T$, for a total cost of 15,000. When the new road is added, the equilibrium situation is that 100 players will choose $S \rightarrow B \rightarrow A \rightarrow T$, for a total cost of 20,000.