

Lecture 2

Lecturer: Aaron Roth

Scribe: Aaron Roth

Congestion Games

In general, to represent an n player game in which each player has k actions, we need k^n numbers just to encode the utility functions. Clearly, for even moderately large k and n , nobody could be expected to understand, let alone play rationally, in such a game. Hence, we will generally think about games that have substantially more structure – despite being large, they have a concise description that makes them easy to reason about.

In this lecture, we will think about congestion games. It will be convenient to think of players as having *cost* functions rather than *utility* functions. Players want to minimize their cost, rather than maximize their utility – but if you like, you can define their utility functions to be the negation of their cost functions.

Definition 1 A congestion game is defined by:

1. A set of n players P
2. A set of m facilities F
3. For each player i , a set of actions A_i . Each action $a_i \in A_i$ represents a subset of the facilities: $a_i \subseteq F$.
4. For each facility $j \in F$, a cost function $\ell_j : \{0, \dots, n\} \rightarrow \mathbb{R}_{\geq 0}$. $\ell_j(k)$ represents “the cost of facility j when k players are using it”.

Player costs are then defined as follows. For action profile $a = (a_1, \dots, a_n)$ define $n_j(a) = |\{i : j \in a_i\}|$ to be the number of players using facility j . Then the cost of agent i is:

$$c_i(a) = \sum_{j \in a_i} \ell_j(n_j(a))$$

i.e. the total cost of the facilities she is using.

Example 1 In class we gave examples of a network routing game and a network creation game. Both are defined over a graph $G = (V, E)$, with facilities $F = E$, and each A_i corresponding to some set of s_i t_i paths in the graph. In general, in a network routing game we think of the cost of each edge as increasing with k (i.e. usage causes congestion). In contrast, in a network creation game, we think of the cost of an edge as being split between the users of that edge (who share the cost to construct the edge) – for example, we could have $\ell_j(k) = w_j/k$, where w_j is the construction cost of edge j .

Ok – so congestion games define an interesting class of n player, many action games that nevertheless have a simple structure and concise representation. What can we say about them? Do they have pure strategy Nash equilibria? Can we *find* those equilibria efficiently? Would agents, interacting together in a decentralized way naturally find said equilibria?

Note that the more of these questions we can answer “yes”, the more we can be comfortable with treating “pure strategy Nash equilibria” as reasonable predictions for what rational players should end up doing in a congestion game.

To answer many of these questions, we will consider “Best response dynamics”. We present it as an algorithm, but you could equally well think about it as a natural model for how people would actually behave in a game. The basic idea is this: we start with players playing an arbitrary set of actions. Then, in arbitrary order, they take turns changing their actions so that they are best responding to their opponents. We continue until (if?) this process converges.

We first make a simple observation:

Algorithm 1 Best Response Dynamics

Initialize $a = (a_1, \dots, a_n)$ to be an arbitrary action profile.
while There exists i such that $a_i \notin \arg \min_{a \in A_i} c_i(a, a_{-i})$ **do**
 Set $a_i = \arg \min_{a \in A_i} c_i(a, a_{-i})$
end while
Halt and return a .

Claim 2 *If best response dynamics halts, it returns a pure strategy Nash equilibrium.*

Proof Immediate from halting condition – by definition, every player must be playing a best response.

■ Of course, it won't always halt – consider matching pennies – but what the above claim means is that

to prove the *existence* of pure strategy Nash equilibria in congestion games, it suffices to analyze the above algorithm and prove that it always halts.

Theorem 3 *Best response dynamics always halt in congestion games.*

Corollary 4 *All congestion games have at least one pure strategy Nash equilibrium.*

Proof We will study the following potential function $\phi : A \rightarrow \mathbb{R}$ defined as follows:

$$\phi(a) = \sum_{j=1}^m \sum_{k=1}^{n_j(a)} \ell_j(k)$$

(Note that the potential function is not social welfare). Now consider how the potential function changes in a single round of best response dynamics, when player i switches from playing some action $a_i \in A_i$ to playing $b_i \in A_i$ instead.

First, because this was a step of best response dynamics, we know that the switch decreased player i 's cost:

$$\begin{aligned} \Delta c_i &\equiv c_i(b_i, a_{-i}) - c_i(a_i, a_{-i}) \\ &= \sum_{j \in b_i \setminus a_i} \ell_j(n_j(a) + 1) - \sum_{j \in a_i \setminus b_i} \ell_j(n_j(s)) \\ &< 0 \end{aligned}$$

The change in potential is:

$$\begin{aligned} \Delta \phi &\equiv \phi(b_i, a_{-i}) - \phi(a_i, a_{-i}) \\ &= \sum_{j \in b_i \setminus a_i} \ell_j(n_j(a) + 1) - \sum_{j \in a_i \setminus b_i} \ell_j(n_j(s)) \\ &= \Delta c_i \end{aligned}$$

Hence, we know that $\Delta \phi < 0$. But since ϕ can take on only finitely many different values (why?) and decreases between each round of best response dynamics, best response dynamics must eventually halt (and hence output a pure strategy Nash equilibrium). ■

Of course, we have only proven convergence, not *fast* convergence. It might take a long time, and if it takes an unreasonably long time (say exponentially many rounds in the number of players), then it might not be a reasonable prediction to assert that rational players will play a Nash equilibrium.

Bad news: it might take a really long time for best response dynamics to converge! But we will be able to say that they converge quickly to an *approximate* Nash equilibrium.

Definition 5 An action profile $a \in A$ is an ϵ -approximate pure strategy Nash equilibrium if for every player i , and for every action $a'_i \in A_i$:

$$c_i(a_i, a_{-i}) \leq c_i(a'_i, a_{-i}) + \epsilon$$

i.e. nobody can gain more than ϵ by deviating.

Lets consider a modification of best response dynamics that only has people move if they can decrease their cost by at least ϵ :

Algorithm 2 FindApproxNash(ϵ)

Initialize $a = (a_1, \dots, a_n)$ to be an arbitrary action profile.
while There exists i, a'_i such that $c_i(a'_i, a_{-i}) \leq c_i(a_i, a_{-i}) - \epsilon$ **do**
 Set $a_i = \arg \min_{a \in A_i} c_i(a, a_{-i})$
end while
Halt and return a .

Claim 6 If FindApproxNash(ϵ) halts, it returns an ϵ -approximate pure strategy Nash equilibrium

Proof Immediately, by definition. ■

Theorem 7 In any congestion game, FindApproxNash(ϵ) halts after at most:

$$\frac{n \cdot m \cdot c_{max}}{\epsilon}$$

steps, where $c_{max} = \max_{j,k} \ell_j(k)$ is the maximum facility cost.

Proof We revisit the potential function ϕ . Recall that $\Delta c_i = \Delta \phi$ on any round when player I moves. Observe also that at every round, $\phi \geq 0$, and

$$\phi(a) = \sum_{j=1}^m \sum_{k=1}^{n_j(a)} c_j(k) \leq n \cdot m \cdot c_{max}$$

By definition of the algorithm, we have $\Delta c_i = \Delta \phi \leq -\epsilon$ at every round, and so the theorem follows. ■