Congestion Games

Kevin Leyton-Brown

What you'll learn in this talk

- What is a congestion game?
  - ...And, why would you use a representation other than the normal form to encode a simultaneous-move game?
- What sorts of interactions do they model?
- What good theoretical properties do they have?
- What are potential games, and how are they related to congestion games?

Definition

Each player chooses some subset from a set of resources, and the cost of each resource depends on the number of other agents who select it.

Definition

Each player chooses some subset from a set of resources, and the cost of each resource depends on the number of other agents who select it.

Definition (Congestion game)

A congestion game is a tuple (N, R, A, c), where

- N is a set of n agents;
- R is a set of r resources;
- $A = A_1 \times \cdots \times A_n$, where $A_i \subseteq 2^R \setminus \{\emptyset\}$ is the set of actions for agent i;
- $c = (c_1, \ldots, c_r)$, where $c_k : \mathbb{N} \to \mathbb{R}$ is a cost function for resource $k \in R$.

Definition

Each player chooses some subset from a set of resources, and the cost of each resource depends on the number of other agents who select it.

Definition (Congestion game)

A congestion game is a tuple (N, R, A, c), where

- N is a set of n agents;
- R is a set of r resources;
- $A = A_1 \times \cdots \times A_n$, where $A_i \subseteq 2^R \setminus \{\emptyset\}$ is the set of actions for agent i;
- $c = (c_1, \ldots, c_r)$, where $c_k : \mathbb{N} \to \mathbb{R}$ is a cost function for resource $k \in R$.

Utility functions:

- Define $\# : R \times A \to \mathbb{N}$ as a function that counts the number of players who took any action that involves resource r under action profile a.
- For each resource k, define a cost function $c_k : \mathbb{N} \to \mathbb{R}$.
- Given an action profile $a = (a_i, a_{-i})$,

$$u_i(a) = -\sum_{r \in R | r \in a_i} c_r(\#(r, a)).$$

Motivating Example: Selfish Routing

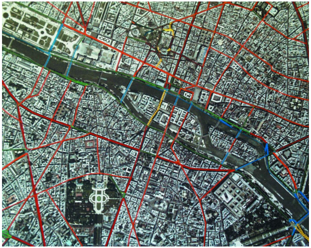

Agents trying to choose uncongested paths in a graph.

- Each edge connecting two nodes is a resource
- Actions are paths in the graph that connect a given user's source and target nodes
  - Stream a video in a computer network
  - Travel along a road network
- The cost function for each resource expresses the latency on each link as a function of its congestion
  - an increasing (possibly nonlinear) function

Motivating Problem: Santa Fe (“El Farol”) Bar Problem

- Each of a set of people independently selects whether or not to go to the bar
- Utility for attending:
  - number of people attending, if less than or equal to 6;
  - 6 minus the number of people attending, if greater than 6
- Utility for not attending: 0
- Note: nonmonotonic cost functions

Motivating Problem: Santa Fe ("El Farol") Bar Problem

- Each of a set of people independently selects whether or not to go to the bar
- Utility for attending:
  - number of people attending, if less than or equal to 6;
  - 6 minus the number of people attending, if greater than 6
- Utility for not attending: 0
- Note: nonmonotonic cost functions

Play the game (raising your hands)

Motivating Problem: Santa Fe ("El Farol") Bar Problem

- Each of a set of people independently selects whether or not to go to the bar
- Utility for attending:
  - number of people attending, if less than or equal to 6;
  - 6 minus the number of people attending, if greater than 6
- Utility for not attending: 0
- Note: nonmonotonic cost functions

Play again. Change your action if you like.

Represent the Santa Fe Bar Problem as a Congestion Game

- n agents
- $n+1$ resources: a seat in the bar, a seat at each agent i's own home
- Actions for i: consume a seat in the bar, consume a seat in i's home
- cost function for seat in own home: 0
- cost function for bar: - number of people attending if this number is leq 6; otherwise number of people attending - 6
- utility function is negative of cost function

Why care about congestion games?

Theorem

Every congestion game has a ("pure-strategy") Nash equilibrium.

Theorem

*A simple procedure (*MyopicBestResponse*) is guaranteed to find a pure-strategy Nash equilibrium of a congestion game.*

Myopic Best Response

- Start with an arbitrary action profile a
- While there exists an agent i for whom a_i is not a best response to a_{-i}
  - $a'_i \leftarrow$ some best response by i to a_{-i}
  - $a \leftarrow (a'_i, a_{-i})$
- Return a

By the definition of equilibrium, MyopicBestResponse returns a pure-strategy Nash equilibrium if it terminates.

Analyzing MBR

In general games MyopicBestResponse can get caught in a cycle, even when a pure-strategy Nash equilibrium exists.

	L	C	R
U	$-1, 1$	$1, -1$	$-2, -2$
M	$1, -1$	$-1, 1$	$-2, -2$
D	$-2, -2$	$-2, -2$	$2, 2$

Can you find a cycle?

Analyzing MBR

In general games MyopicBestResponse can get caught in a cycle, even when a pure-strategy Nash equilibrium exists.

	L	C	R
U	$-1, 1$	$1, -1$	$-2, -2$
M	$1, -1$	$-1, 1$	$-2, -2$
D	$-2, -2$	$-2, -2$	$2, 2$

This game has one pure-strategy Nash equilibrium, (D, R). However, if we run MyopicBestResponse with $a = (L, U)$ the procedure will cycle forever.

Potential Games

Definition (Potential game)

A game $G = (N, A, u)$ is a potential game if there exists a function $P : A \rightarrow \mathbb{R}$ such that, for all $i \in N$, all $a_{-i} \in A_{-i}$ and $a_i, a'_i \in A_i$, $u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i}) = P(a_i, a_{-i}) - P(a'_i, a_{-i})$.

Potential Games

Definition (Potential game)

A game $G = (N, A, u)$ is a potential game if there exists a function $P : A \rightarrow \mathbb{R}$ such that, for all $i \in N$, all $a_{-i} \in A_{-i}$ and $a_i, a_i' \in A_i$, $u_i(a_i, a_{-i}) - u_i(a_i', a_{-i}) = P(a_i, a_{-i}) - P(a_i', a_{-i})$.

Theorem

Every potential game has a pure-strategy Nash equilibrium.

Potential Games

Definition (Potential game)

A game $G = (N, A, u)$ is a potential game if there exists a function $P : A \to \mathbb{R}$ such that, for all $i \in N$, all $a_{-i} \in A_{-i}$ and $a_i, a'_i \in A_i$, $u_i(a_i, a_{-i}) - u_i(a'_i, a_{-i}) = P(a_i, a_{-i}) - P(a'_i, a_{-i})$.

Theorem

Every potential game has a pure-strategy Nash equilibrium.

Proof.

Let $a^* = \arg\max_{a \in A} P(a)$. Clearly for any other action profile a', $P(a^*) \geq P(a')$. Thus by the definition of a potential function, for any agent i who can change the action profile from a^* to a' by changing his own action, $u_i(a^*) \geq u_i(a')$.

Congestion Games have Pure-Strategy Equilibria

Theorem

Every congestion game is a potential game.

- Every congestion game has the potential function

$$P(a) = \sum_{r \in R} \sum_{j=1}^{\#(r,a)} c_r(j)$$

  - There's a proof in the book
  - Main intuition: utility functions are linear combinations of cost functions, so most of the terms in this expansion cancel out when we take the difference between the potential values for two similar action profiles
- It also turns out that every potential game is a congestion game (harder to show)

Myopic Best Response Converges

Theorem

The MyopicBestResponse *procedure is guaranteed to find a pure-strategy Nash equilibrium of a congestion game.*

Proof.

It is sufficient to show that MyopicBestResponse finds a pure-strategy Nash equilibrium of any potential game. With every step of the while loop, $P(a)$ strictly increases, because by construction $u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i})$, and thus by the definition of a potential function $P(a'_i, a_{-i}) > P(a_i, a_{-i})$. Since there are only a finite number of action profiles, the algorithm must terminate.

Myopic Best Response: Analysis

MyopicBestResponse converges for CGs regardless of:

- the cost functions (e.g., they do not need to be monotonic)
- the action profile with which the algorithm is initialized
- which agent best responds (when there's a choice)
- And even if we change best response to "better response"

Myopic Best Response: Analysis

MyopicBestResponse converges for CGs regardless of:

- the cost functions (e.g., they do not need to be monotonic)
- the action profile with which the algorithm is initialized
- which agent best responds (when there's a choice)
- And even if we change best response to "better response"

Complexity considerations:

- the problem of finding a pure Nash equilibrium in a congestion game is PLS-complete
  - as hard to find as any other object whose existence is guaranteed by a potential function argument
  - intuitively, as hard as finding a local minimum in a traveling salesman problem using local search
- We thus expect MyopicBestResponse to be inefficient in the worst case

Conclusions

- Congestion games are a compact and intuitive way of representing interactions in which agents care about the number of others who choose a given resource, and their utility decomposes additively across these resources
- Potential games are a less-intuitive but analytically useful characterization equivalent to congestion games
  - potential function: a single function that captures any player's utility change from deviating
- These games always have pure-strategy Nash equilibria
- MyopicBestResponse always converges to a pure-strategy Nash equilibrium in congestion/potential games.
- They're very widely studied in the literature, particularly in CS
  - realistic model
  - pure-strategy equilibria are otherwise fairly rare
  - nice computational story

References

1. *A class of games possessing pure-strategy Nash equilibria*, RW Rosenthal, International Journal of Game Theory, 1973.
2. *Potential games*, D Monderer, LS Shapley, Games and Economic Behavior, 1996.
3. *Bounding the inefficiency of equilibria in nonatomic congestion games*, T Roughgarden, É Tardos, Games and Economic Behavior, 2004.
4. *"The El Farol Bar Problem"*, Wikipedia, `http://en.wikipedia.org/wiki/El_Farol_Bar_problem`, accessed 1/15/2014.
5. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, Y Shoham, K Leyton-Brown, Cambridge University Press, 2009, §6.4.