Algorithmic Game Theory

(Network) Congestion Games,
Selfish Routing and Potential Games

Vangelis Markakis
markakis@gmail.com

Different Variations

- Congestion Games (CGs): (i) Resources, (ii) Players choose subset of resources, (iii) Payoffs/Costs based on congestion
- Network CGs/Selfish routing: Available subsets of resources form source-destination paths
- Atomic case: k players, each routing a unit weight
- Non-Atomic case: Infinite, "tiny" players form demands
- Splittable case: Players can split their weight
- Example:

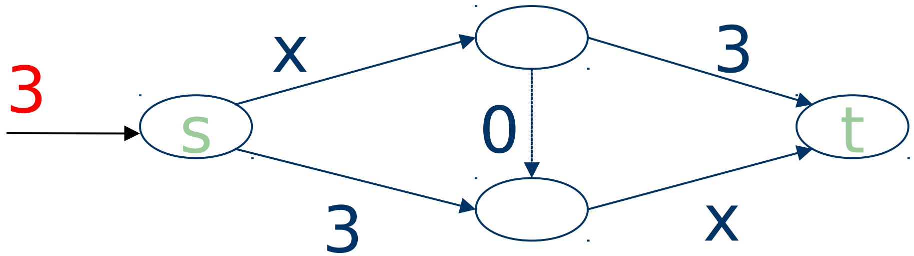

Different Variations

- Congestion Games (CGs): (i) Resources, (ii) Players choose subset of resources, (iii) Payoffs/Costs based on congestion
- Network CGs/Selfish routing: Available subsets of resources form source-destination paths
- Atomic case: k players, each routing a unit weight
- Non-Atomic case: Infinite, “tiny” players form demands
- Splittable case: Players can split their weight
- Example:

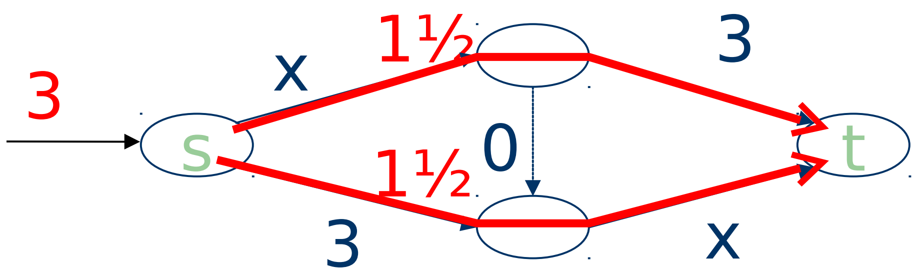

Different Variations

- Congestion Games (CGs): (i) Resources, (ii) Players choose subset of resources, (iii) Payoffs/Costs based on congestion
- Network CGs/Selfish routing: Available subsets of resources form source-destination paths
- Atomic case: k players, each routing a unit weight
- Non-Atomic case: Infinite, "tiny" players form demands
- Splittable case: Players can split their weight
- Example:

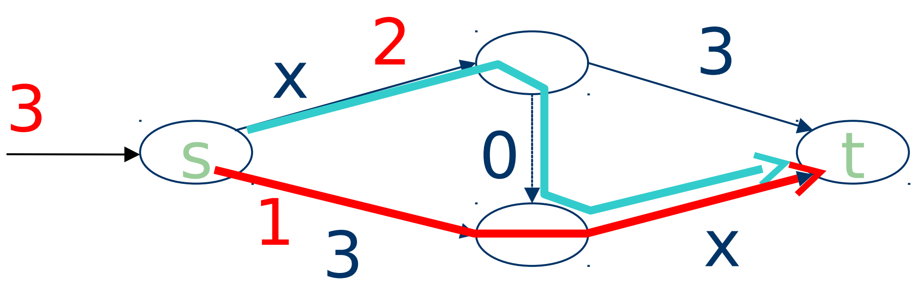

Atomic selfish routing

A model for atomic selfish routing games

Formal description:

- directed graph $G = (V,E)$
- finite number of k players
- Player i has an origin vertex s_i and a destination vertex t_i
- Each player wants to route 1 unit of traffic on a single path from s_i to t_i
  - Similar reasoning applies if we allow the player to split the traffic into different paths from s_i to t_i
- for each edge e, a cost function $c_e()$
  - Assumed nonnegative and nondecreasing
  - Depends on the (integer) number of traffic units crossing edge e

A model for atomic selfish routing games

Consider a feasible flow f

- Let P_i = set of all distinct paths from s_i to t_i
- f can be specified by a vector $(p_1, p_2, \ldots, p_k)$, where for i=1,..., k, the path p_i is the path chosen by player i ($p_i \in P_i$)

Representation as an edge flow vector:

- We can also write f as a vector along edges of the graph
- For every edge e, $f_e = \Sigma_{p:\, e \in p}\, f_p$
- Hence, in this setting, f_e = number of players who selected a path that includes e

Social cost of a flow

$$C(f) = \Sigma_p\, f_p\, c_p(f) = \Sigma_e\, f_e\, c_e(f_e)$$

Equilibrium flows

- When is a flow f an equilibrium flow?
- When no agent has an incentive to switch his unit of traffic to a different path connecting s_i to t_i
- Consider a feasible flow f given by the paths $(p_1, p_2, \ldots, p_k)$
- Given a path $p' \neq p_i$, let
  - $p' \cap p_i$ = set of edges in common between p' and p_i
  - $p' \setminus p_i$ = set of edges in p' but not in p_i
- Definition: A feasible flow $f = (p_1, p_2, \ldots, p_k)$ is a Nash equilibrium flow if for every player i = 1, ..., k, and every $p' \in P_i$,

$$\sum_{e \in p_i} c_e(f_e) \leq \sum_{e \in p' \cap p_i} c_e(f_e) + \sum_{e \in p' \setminus p_i} c_e(f_e + 1)$$

PoA in atomic games

- How bad is selfish routing in atomic games?
- Difference with non-atomic games: here we can have multiple equilibria with different social cost
  - Unlike the non-atomic case, where all equilibria have the same social cost
- For a class of atomic selfish routing games:

$$PoA = \max \max_f C(f)/C(f^*)$$

- The first maximization is w.r.t. to all the games in the class
- The second maximization is w.r.t. all the Nash flows

PoA in atomic games

- What is the effect of atomicity on PoA?
- We can start with linear cost functions

Example:

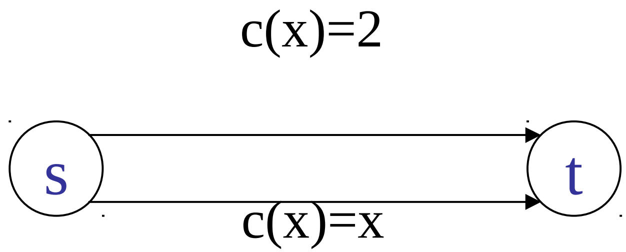

- Consider 2 players, each controlling 1 unit of traffic
- Optimal solution: each player in different edge (also an equilibrium)
- Optimal social cost = 2+1 = 3
- There is a 2^{nd} equilibrium: both agents take the lower edge (why?)
- Hence PoA ≥ 4/3

PoA in atomic games

- Can PoA get higher than 4/3 in atomic games?

Example:

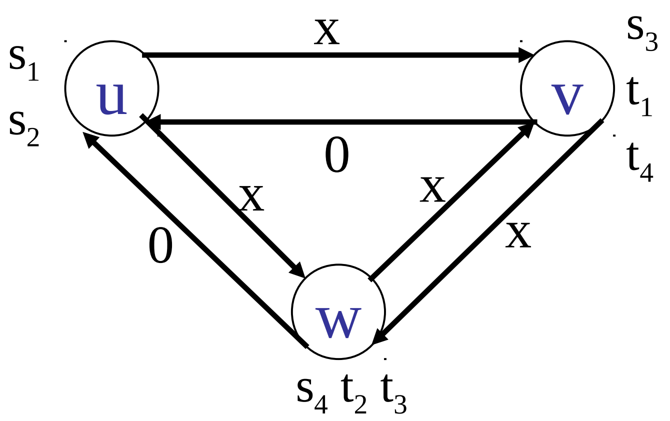

- 4 players, 2 choices per agent, a 2-hop path or a 1-hop path
- Claim: PoA ≥ 5/2 (Homework)

PoA in atomic games

- PoA can be worse in atomic games
- But not much worse for linear cost functions…

Theorem [Christodoulou, Koutsoupias '05]: For atomic selfish routing games with linear cost functions

PoA = 5/2

- [Aland et al. '11]: Generalizations to polynomial cost functions analogous to the non-atomic case
- For polynomials of degree p, PoA upper-bounded by a function of p
  - Exponential in p however
  - Much slower growth in non-atomic case: O(p/ln(p))
- Main conclusion: relatively small PoA for low degree polynomial cost functions

PoA in atomic games

- Proof techniques for atomic games can be applied to analyzing the PoA for a much wider class of games
- Extensive literature over the last 2 decades, since 2000
- Very few PoA proofs do not follow the proof technique of the 5/2 upper bound
- General approach:

$$C(f) \le \alpha C(f^*) + \beta C(f) \Rightarrow PoA \le \alpha/(1-\beta)$$

- Ideas for reducing the PoA in a game
  - Impose restrictions on the strategy space of some players
  - Impose tolls
- Interesting research agenda for transportation engineering
- For more, see Chapters 11-14 of Roughgarden's book

Establishing existence of pure equilibria

- In arbitrary games with multiple (more than 2) players, it is generally hard to argue about existence of pure Nash equilibria
- But in many classes of games derived from application scenarios, one can exploit the structure of the problem
- For selfish routing, pure Nash equilibria exist for both atomic and non-atomic games
- Can we identify properties that can guarantee existence of pure Nash equilibria in other multi-player games as well?

Back to selfish routing

Recall description of atomic games:

- directed graph $G = (V,E)$
- finite number of k players
- Player i has an origin vertex s_i and a destination vertex t_i
- Each player wants to route 1 unit of traffic on a single path from s_i to t_i
- for each edge e, a cost function $c_e()$
  - Assumed nonnegative and nondecreasing
  - Depends on the (integer) number of traffic units crossing edge e
- Strategy space of player i: all the s_i-t_i paths in the graph G
  - Each pure strategy corresponds to a distinct s_i-t_i path

Congestion Games

A generalization:

- A set of players N = {1, 2,..., n}
- A set of m resources, E = {1, 2, ..., m}
- Each resource j has a cost function $c_j(.)$ dependent on the number of players using it
  - $c_j(n_j)$ = cost incurred by resource j when the number of players using j equals n_j
- Strategy space S^i of player i: a collection of subsets of the resources allowable for player i
  - Each pure strategy is a distinct subset of E
- Cost of player i at a strategy profile: sum the resource cost functions over all resources being used by the player

Atomic routing games as congestion games

Resources = edges

- Each edge has a cost function dependent on the number of players using it

Strategy space of player i = all s_i-t_i paths

- Each player selects a subset of the resources that corresponds to a valid s_i-t_i path

Cost of an agent:

- Need to sum over all the cost functions of the edges (resources) being used

Corollary: Routing games is just a special case of congestion games

Pure equilibria and congestion games

Congestion games have been well studied due to their wide applicability in various domains

Theorem [Rosenthal '73]: Every congestion games admits a pure Nash equilibrium

- One of the classic results on congestion games
- Unlike routing games, a cost function does not have to be non-decreasing in a congestion game

Pure equilibria and congestion games

Proof sketch of Rosenthal's theorem

Most important idea:

- Consider a strategy profile s = (s_1, s_2,..., s_n)
- With n_j = number of players using resource j at profile s
- Define the function:

$$\Phi(s) = \sum_{j \in E} \sum_{i=1}^{n_j} c_j(i)$$

- For each player i, let $c_i(s)$ be the cost she experiences
  - $c_i(s) = \Sigma\, c_j(n_j)$ where the sum is over all $j \in s_i$
- Crucial property: for every player i, and every possible deviation s_i'

$$\Phi(s) - \Phi(s_i', s_{-i}) = c_i(s) - c_i(s_i', s_{-i})$$

Analogous proof also for non-atomic games

Generalizing congestion games

- Can we establish existence for a broader class of games?
- The arguments in the proof for congestion games can help us
  - We identified a function Φ that captures improvements by a deviation of a single player
  - The function plays the role of measuring the difference in a player's utility before and after a deviation for any player
  - One single function capturing the deviation gain of every player
- Other games may also possess this property
- This is an example where a proof technique gives rise to a new definition

Potential games

- Definition: a game G is an exact potential game if there exists a function Φ, s.t. for every strategy profile $s = (s_1, s_2, \ldots, s_n)$, for every player i and every strategy s_i' of pl. i:

$$\Phi(s_i, s_{-i}) - \Phi(s_i', s_{-i}) = c_i(s_i, s_{-i}) - c_i(s_i', s_{-i})$$

- a game is an ordinal potential game if there exists a function Φ, s.t. for every $s = (s_1, s_2, \ldots, s_n)$, for every pl. i and every strategy s_i' of pl. i:

$$\Phi(s_i, s_{-i}) > \Phi(s_i', s_{-i}) \text{ iff } c_i(s_i, s_{-i}) > c_i(s_i', s_{-i})$$

- We refer to a game as a potential game if it is either an exact or an ordinal potential game
  - The function F is called a potential function
  - Not always easy to find

Potential games

Examples

- Congestion games
  - Using the potential function in Rosenthal's theorem
- Network cost-sharing games: games that regard the sharing of cost for building a network
  - See chapter 14 in 20LAGT
- Location games: games where players need to decide where to locate a store or a service point
  - See chapter 15 in 20LAGT

Finding a potential function is the only general methodology we know for proving that a multi-player game has pure Nash equilibria

Another Example: The MAX-CUT game

Note: we can adjust the definitions and use the utility function of a player rather than the cost

Example:

- Undirected graph G = (V,E) representing a social group (colleagues, student mates, etc)
- Each player is a node v in V
- An edge (u, v) means u "does not like" v (and vice versa)
- Suppose each player has to choose among 2 suggested excursions (or 2 activities in general)
- For simplicity: strategy of node i: $s_i \in$ {Black,White}
- Utility of node i at profile s: # neighbors of different color

Another Example: The MAX-CUT game

Lemma: for every graph G, the corresponding game is a potential game

Proof:

- Consider a profile s = (s_1, s_2,..., s_n) with $s_i \in \{B, W\}$
- Under the profile s, the players are partitioned in 2 sets, T_1, and T_2
- Define $\Phi(s)$ = size of the cut determined by T_1, and T_2

  = # edges crossing the cut
- Claim: $\Phi(s)$ is a potential function for this game

Potential games

- Theorem: every finite potential game admits a pure Nash equilibrium
- Proof: the profile minimizing Φ is an equilibrium (or maximizing Φ if we use utilities instead of costs)
  - Let $s = (s_1, s_2, \ldots, s_n)$ be a global minimum of Φ
  - Suppose it is not a Nash equilibrium, so some player i can improve by deviating
  - new profile: $s' = (s_i', s_{-i})$
  - $\Phi(s') - \Phi(s) = c_i(s') - c_i(s) < 0$
  - Thus, $\Phi(s') < \Phi(s)$, contradicting that s minimizes Φ

- More generally, the set of pure Nash equilibria is exactly the set of local minima of the potential function
  - Local minimum = no player can improve the potential function by a unilateral deviation

Best Response Dynamics and its variants

Reaching an equilibrium

- Suppose that we have a multi-player game with pure Nash equilibria
- How do we expect the players to find an equilibrium?
- Meaningful question for games that are played repeatedly
- If a player does not have a dominant strategy what would she do?
  - Probably start with some initial strategy
  - As she observes the other players' actions, she can adjust her own in the next rounds
  - Essentially each player is applying some learning algorithm to determine her next move
  - We can observe a dynamic behavior of each player in a sequence of rounds based on her observations for the other players

Best response dynamics

Vanilla version

- Each player starts with some arbitrary strategy
  - Let s = (s_1, s_2,…, s_n) current profile
  - If there exists a player who is not currently playing a best response, switch that player's strategy to his best response
    - If there are multiple such players, pick one arbitrarily
    - If there are multiple best responses for a player pick one arbitrarily
  - Update current profile
- Terminate when no player can improve (thus a Nash equilibrium)

- We can define several variations of the basic version
  - Introduce specific criteria for breaking ties
  - Better response vs best response
  - Synchronous vs asynchronous

Best response dynamics

Visualization of dynamics

- We can think of best response dynamics as a walk in a graph
- Directed graph, G = (V, E)
- V = set of all strategy profiles
- There is an edge from a profile s to s' if there exists a best response move by some player at s that results in s'

Convergence of best response dynamics

•It is not obvious whether this process converges or not

•If the process does not converge, the corresponding graph has a directed cycle (possibly more)

•If the process converges from any initial profile, the graph has no cycles

Examples

Prisoner's dilemma

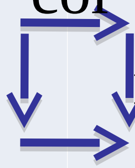

3,3	0,4
4,0	1,1

Battle of the sexes

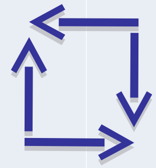

2,1	0,0
0,0	1,2

Matching pennies

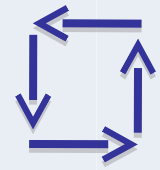

-1,1	1,-1
1,-1	-1,1

- The first two have
- Can we come up with conditions that guarantee convergence of the dynamics?
- For sure, games with no pure equilibria do not converge
- Q: Are the first two games potential games?

Convergence of best response dynamics

Theorem: In a finite potential game, and for any initial strategy profile, best response dynamics converge to a pure equilibrium

- In every iteration, some player makes an improvement move
- Hence, the potential function strictly decreases
- Since the strategy space is finite, the potential function cannot decrease forever
    - It will halt at a local minimum of the potential, i.e., an equilibrium

Back to the MAX-CUT game

Corollary: best response dynamics converge to an equilibrium for the MAX-CUT game

How do we implement best response dynamics here?

- Consider a profile $s = (s_1, s_2, \ldots, s_n)$ with $s_i \in \{B, W\}$
- Under the profile s, the players are partitioned in 2 sets, T_1, and T_2
- If s is not an equilibrium, some player has a better move, i.e., by switching his strategy, he increases the size of the cut between the black and white sets
- Claim: Best response dynamics are equivalent to the greedy ½-approximation for the MAX-CUT problem!

Speed of convergence

- How fast do best response dynamics converge in potential games?
- We measure it with the number of iterations needed
- In worst case, it can be very slow
- A best response move may decrease the potential function only by a tiny amount
- It may require exponentially many (in terms of the number of players) iterations to converge

Approximate best response dynamics

Relax the convergence requirements:

- We can compromise with convergence to an approximate equilibrium
- Convenient version of approximation: for potential games, a profile $s = (s_1, s_2, \ldots, s_n)$ is an ε-equilibrium if for every player i and every deviation s_i':

$$c_i(s_i', s_{-i}) \geq (1-\varepsilon)\, c_i(s)$$

- No deviation can produce a significant drop in the cost

Approximate best response dynamics

- We can also impose only moves that provide significant improvement
- ε-move: a deviation s_i' from the current profile s.t. $c_i(s_i', s_{-i}) < (1-\varepsilon)\, c_i(s)$

ε-best response dynamics (basic version):

- Each player starts with some arbitrary strategy
  - Let $s = (s_1, s_2, \ldots, s_n)$ be the current profile
  - While the current profile is not an ε-equilibrium:
    - Pick a player who has an ε-move
    - Break ties arbitrarily if there are multiple such players or multiple ε-moves
    - Update current profile

Approximate best response dynamics

- A slight adjustment to a better version:

ε-best response dynamics (maximum-gain):

- Each player starts with some arbitrary strategy
  - Let $s = (s_1, s_2, \ldots, s_n)$ be the current profile
  - While the current profile is not an ε-equilibrium:
    - Pick a player who has an ε-move
    - If there are multiple players who have an ε-move, pick the player i who can obtain the largest cost decrease in $c_i(s)$
    - Update current profile

Approximate best response dynamics

- An application to atomic selfish routing games:

Consider an atomic routing games where:

  - All players have the same origin and destination vertex
  - There exists an $\alpha \geq 1$, s.t. for every edge e of the graph, $c_e(f+1) \in [c_e(f), \alpha c_e(f)]$ (α-bounded jump condition)

Then, the max-gain variant of ε-best response dynamics converges to an ε-equilibrium in polynomial time

- In particular, at most $k\alpha/\varepsilon \ln(\Phi(s^0)/\Phi_{min})$ iterations
- s^0 = initial profile
- Φ_{min} = minimum of Φ

Beyond best response dynamics

- Several other variations have been considered
- Rich interaction between machine learning and game theory
- We can think of each player as using a learning algorithm
- This leads to a probabilistic algorithm for each player
- Important example: no-regret dynamics
  - Each player maintains a probability distribution on his pure strategies based on past performance
  - Multiplicative weight updates in each step
  - Players try to bound the regret of their strategy against playing the best pure strategy
  - Average regret -> 0
  - Convergence to a different equilibrium concept (coarse correlated equilibria)

Dynamics and equilibrium concepts

- For more on no-regret dynamics see Chapter 17 in 20LAGT
- At the end, what would be an appropriate stability concept?
- Nash equilibria seem appropriate only for 0-sum games
- But still the driving force behind any other concept
- Considerations for studying alternative equilibrium notions:
  - Computational complexity
  - Convergence of natural learning algorithms (this would mean that players actually have a chance to reach such a state)
- Still a question under investigation…