CMU 15-896 Algorithms, Games, & Networks
Spring 2013

01/31/13

Price of Anarchy, Price of Stability, Potential & Congestion Games, Part II: Dynamics, Stability, and Influence

Your guide:

Avrim Blum

The Price of Uncertainty: Safety analysis for multiagent systems

Maria-Florina Balcan Georgia Tech
Avrim Blum Carnegie Mellon
Yishay Mansour Tel-Aviv

[Workshop on Innovations in AGT, 2011][EC'09]

Games are an abstraction

Aim to identify most important decisions of participants and main factors in payoffs.

- But there are always low-order unmodeled factors or variables.
- And some players whose incentives not modeled well at all.

High-level question: How much impact can these issues have on dynamics in the system?

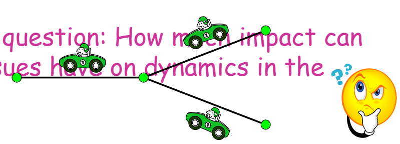

Games are an abstraction

We'd like to think that if we get people into a good equilibrium, and players are selfish, reasonably myopic, etc, then behavior will stay there.

We can safely walk away and be confident the system will stay in a good state

But what if there are small fluctuations in underlying cost/payoff structure? Could they cause natural dynamics to spin out of control?

Games are an abstraction

We'd like to think that if we get people into a good equilibrium, and players are selfish, reasonably myopic, etc, then behavior will stay there.

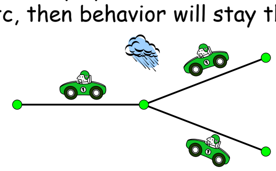

But what if there are small fluctuations in underlying cost/payoff structure? Could they cause natural dynamics to spin out of control?

Games are an abstraction

We'd like to think that if we get people into a good equilibrium, and players are selfish, reasonably myopic, etc, then behavior will stay there.

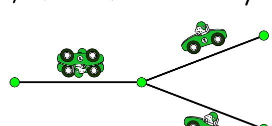

Can small fluctuations or a few unpredictable players cause natural dynamics to get system into a high-cost state?

Or what about a few players acting unpredictably?

High-level question

A few ways this could happen:

- Small changes cause good equilibria to disappear, only bad ones left. (economy?)
- Bad behavior by a few players causes pain for all (nukes)
- Neither of above, but instead through more subtle interaction with dynamics.

Can small fluctuations or a few unpredictable players cause natural dynamics to get system into a high-cost state?

Focus here on this last issue.

Focus of this work

Games with the following properties:

- Potential games, best/better-response dynamics.
  - These games have non-negative potential function $\Phi(S)$ such that if any player moves, say reducing his cost by Δ, then Φ decreases by Δ too.
  - Better-response dynamics *will* reach equilibrium.
  - And the maximum gap between $\Phi(S)$ and cost(S) bounds how bad a state can get if no fluctuations.

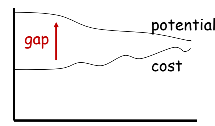

Focus of this work

Games with the following properties:

- Potential games, best/better-response dynamics.
- Small gap between potential and social cost.
  - Without fluctuations, can walk away from good state even if not an equilibrium: gap bounds possible increase.
  - Single perturbation can't make dynamics do bad things.

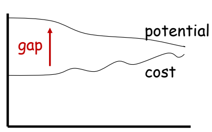

Focus of this work

Games with the following properties:

- Potential games, best/better-response dynamics.
- Small gap between potential and social cost.
- No individual player can influence total cost of others by too much. (With Byzantine players, will define social cost as sum of costs to *others*.)

Focus of this work

Games with the following properties:

- Potential games, best/better-response dynamics.
  - Example: set-cover games.
    - n players, m resources, with costs $c_1,\ldots,c_m$.
    - Each player allowed to use some resources, not others.

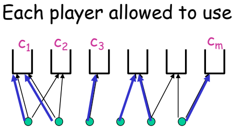

Cost(S) = $\sum_{e: n_e \geq 1} c_e$

Φ(S) = $\sum_e \sum_{i=1}^{n_e} c_e / i$

- Each player chooses some allowable resource.
- Players split cost with all others choosing same one.

$\text{cost}(S) \leq \phi(S) \leq H(n) \cdot \text{cost}(S)$

Model

- Players follow best (or better) response dynamics.
- Costs of resources can fluctuate between moves: $c_i^{\dagger} \in [c_i, c_i(1+\varepsilon)]$
- Alternatively, one/few Byzantine players who move between time steps
- Play begins in a low-cost state.
- How bad can things get?

Price-of-Uncertainty(ε) of game = maximum ratio of eventual social cost to initial cost.

Model

- Players follow best (or better) response dynamics.
- Costs of resources can fluctuate between moves: $c_i^t \in [c_i, c_i(1+\varepsilon)]$

Price-of-Uncertainty(ε) of game = maximum ratio of eventual social cost to initial cost.

Model

- Players follow best (or better) response dynamics.
- Does reachable set contain state of high cost compared to start?
- Costs of resources can fluctuate between moves: $c_i^t \in [c_i, c_i(1+\varepsilon)]$

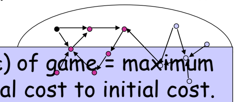

Price-of-Uncertainty(ε) of game = maximum ratio of eventual social cost to initial cost.

One way to look at this:

- Define graph: one node for each state. Edge $u \to v$ if perturbation can cause BR to move from u to v.

Set-cover games

- n players, m resources. Each player chooses one of allowable resources. Players split cost with all others choosing same one.
- Price of anarchy = n

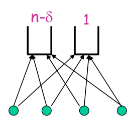

- Potential $\Phi \in$ [cost, cost·log(n)]

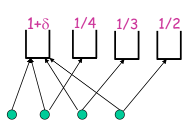

Main results

Set-cover games:

Good news:

- If $\varepsilon = O(1/nm)$ then PoU = O(log n).

Bad news:

- For ε=1, PoU = $\Omega(n)$. For any constant ε, PoU = $\Omega(n^{1/2})$
- Also, a single Byzantine player can take state from a PNE of cost O(OPT) to one of cost Ω(n·OPT).

Upper bounds hold even for **better**-response

Lower bounds hold even for **best**-response

Main results

General fair-cost-sharing games:

- If *many players* for each (s_i,t_i) pair ($n_i = \Omega(m)$), then PoU = O(1) even for constant ε >0.
- Open for general number of players.

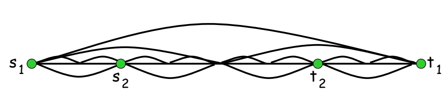

Main results

General fair-cost-sharing games:

- If *many players* for each (s_i,t_i) pair ($n_i = \Omega(m)$), then PoU = O(1) even for constant ε >0.

Matroid congestion games: (strategy sets are bases of matroid. E.g., set-cover where choose k resources)

- If $\varepsilon = O(1/nm)$ then PoU = O(log n) for fair cost-sharing.
- In general, if $\varepsilon = O(1/nm)$ then PoU = O(GAP).

These require **Best**-Response. **Better**-resp not enough

Also results for β-nice games, job scheduling, …

Set-Cover games (lower bound)

A single Byzantine player can take state from a PNE of cost O(OPT) to one of cost $\Omega(n \cdot OPT)$.

Set-Cover games (lower bound)

Two kinds of players:

- n of Class I:
- n-1 of Class II:

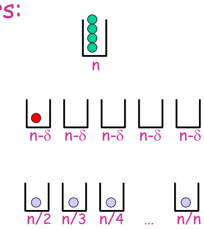

Plus one Byzantine player…

➢ Implies bound for case that costs can fluctuate by factor of 2. More complicated ex for $\varepsilon \to 1/n^{1/2}$.

Set-Cover games (upper bound)

For upper bound, think of players in sets as a stack of chips.

- View i^{th} position in stack j as having cost c_j/i. Load chips with value equal to initial cost.
- When player moves from j to k, move top chip. Cost of position goes up by at most $(1+\varepsilon)$.
- At most mn different positions. So, following the path of any chip and removing loops, cost of final set is at most $(1+\varepsilon)^{nm}$ times its value.

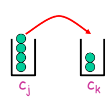

So, if $\varepsilon = O(1/nm)$ then PoU = O(log n).

Matroid games

In matroid games, can think of each player as controlling a *set* of chips.

- Nice property of best response in matroids:
  - Can always order the move so that each individual chip is doing better-response.
- Apply previous argument.
- Fails for better-response.
  - Here, can get player to do kind of binary counting, bad even for exponentially-small ε.

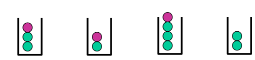

Fair cost sharing in general graphs

If many players of each type, can also show best-response dynamics can't do badly.

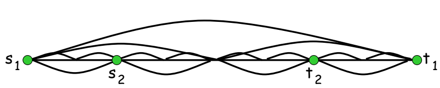

Fair cost sharing in general graphs

If many players of each type, can also show best-response dynamics can't do badly.

Outline of argument:

- Hard to analyze cost of state directly, instead track upper bound $c^*(S_t) = cost(S_0 \cup \ldots \cup S_t)$.
  - c^* changes at most m times.
- Many players of each type $\Rightarrow$ average cost of each is low compared to c^* $\Rightarrow$ each change to c^* is small.

Because it can never be a BR move to switch to something of cost > $(1+\varepsilon)$ times average for type.

PoU versus Price of Anarchy

Main focus: games with both good and bad equilibria.

- Can small fluctuations or single Byzantine player cause behavior to move from good to bad?

Can also have cases where state can get worse even than worst equilibrium.

- Market-sharing: 2 vs log(n) even for best-response.

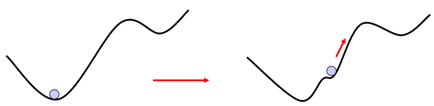

PoU versus Price of Anarchy

β–nice games [AAEMS-EC08] : incentives grow stronger as cost gets above β times optimal (typically β = PoA).

- $2\Delta(S) \geq \mathrm{cost}(S) - \beta\,\mathrm{OPT}$, $\Delta(S) = \sum_i \Delta_i(S)$.
- Here, at least can show state won't get above O(β) times optimal, even with substantial perturbation or many Byzantine players (random order).

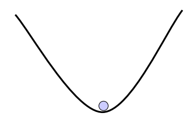

Subsequent results

[Balcan, Constantin, Ehrlich]

- Set-cover games:
  - Upper bound with dependence only on m, not n.
  - Lower bound under random move-ordering.
- Consensus games:
  - Nearly-tight bounds on effect of ε-perturbations
  - Tight bounds on effect of B Byzantine players.

Summary and open problems

Looking at: when can small perturbations or a few bad players lead natural dynamics astray?

- When is it safe to turn your back?
- Upper/lower bounds for a number of classes of games.

Open problems:

- General case of fair cost-sharing games?
- Analyze time to failure for random fluctuations?
- Instance-based analysis?