**Algorithms, Games, and Networks** January 31, 2013

Lecture 6

*Lecturer: Avrim Blum* *Scribe: Xiaowen Ding*

# 1 Overview

In today's lecture, we have a conference talk[EC09]. We discuss Price of Uncertainty(PoU), which describes the impact that small fluctuations or a small number of incorrectly modeled (Byzantine) players can have on dynamics in systems.

# 2 Perturbation model and Byzantine model

In previous lectures, we implicitly were operating under the assumption that if people are in an equilibrium, then they will stay there. But games are an abstraction, and in this lecture we look at ways in which imperfections in this abstraction could lead to behavior that goes astray. We will be focusing here in particular on *potential games* and on dynamics in which players move one at a time making favorable deviations if they find any (which as we saw last time, is very natural for potential games since (a) it leads to an equilibrium, and (b) it ordinarily cannot cause the cost of the state to increase by more than the maximum gap between potential and cost – more about this below).

Specifically, we consider the following two settings. The first is the *perturbation model* in which we allow for small fluctuations in costs between moves. The other is the *Byzantine model* in which there could be one or few unpredictable players (players whose payoffs are unmodeled).

For example, if you think about traffic flow, the cost of driving along some stretch of road is not just a function of the number of drivers on it, but also depends on whether it is raining, whether there is construction, various chance events etc. So if we look over time, each day the cost function is slightly different. Also, it might be the case that some players act unpredictably, which is the Byzantine model.

If a game has small fluctuations in costs, or a few Byzantine players, could behavior spiral out of control?

Here are a few ways this could happen:

- small changes cause good equilibria to disappear, only bad ones left

- Bad behavior by a few players causes pain for all
- Neither of above, but instead through more subtle interaction with dynamics

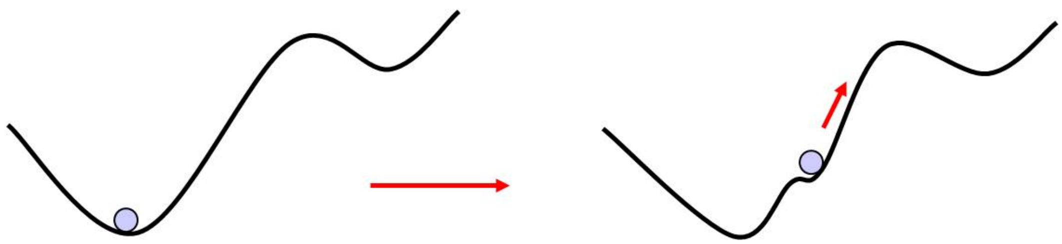

Observe the graph, small fluctuation can cause natural dynamics to get system into a high-cost state.

## 3 Game properties

In this lecture, we are interested in games with the following properties:

1.Potential games with best or better response dynamics.

Potential games have non-negative potential function $\Phi(S)$ such that if any player moves, the change of potential is the same as the change of his cost.

Here, better-response means if there are multiple choices that one can take to reduce his cost, just take one of them. And best-response means he would take the one of minimum cost.

Better-response dynamics will eventually reach equilibrium because the potential is decreasing all the time and there are only finite states. Eventually, the potential could not decrease any more and everyone will be in a stable state.

2.Small gap between potential and social cost.

Remember what we have discussed in previous lecture, in cost-sharing games, the potential can be bounded by cost: $cost(S) \leq \Phi(S) \leq \log(n) cost(S)$.

The maximum gap between $\Phi(S)$ and $cost(S)$ indicates how bad a state can get if there are no fluctuations. Also, single small perturbations can't make dynamics do bad things in this model.

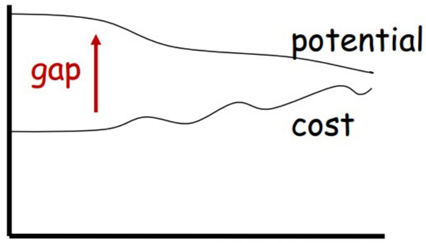


3.No individual player can influence total cost of others by too much.

So, if Byzantine players cause a large effect it must be through a sequence of moves by players rather than, e.g., the Byzantine players exercising a "nuclear option".

# 4 Set-cover games

## 4.1 Definition

Set-cover games are a special case of fair cost-sharing.

In a set-cover game, we have $n$ players and $m$ resources with costs $c_1, c_2, ...c_m$ respectively.

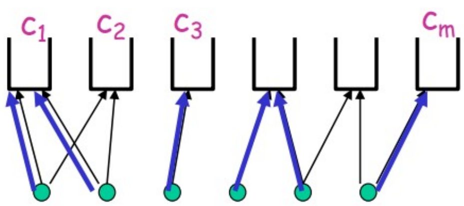


Each player has his personal allowable resources, and has to choose one of those allowable resources. The cost of each resource is split among all players that choose that resource.

So the total cost of the game is $cost(S) = \sum_{e:n_e \geq 1} c_e$, and we can define potential $\Phi(S) = \sum_c \sum_{i=1}^{n_e} \frac{c_e}{i}$ as what we defined in previous lecture.

Also, we can model this game as a special case of the fair-cost-sharing in networks game. Everyone has a single source and has access to some mid-points for free. And then, the cost from the mid-point to the terminal is the cost of a resource. Here is the corresponding graph.

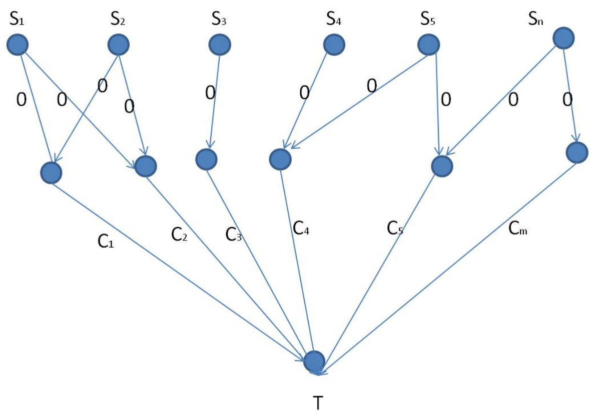


Another interpretation is the set cover problem. Each set is a resource and each node is a player, each player can choose one set from the sets it belongs to.

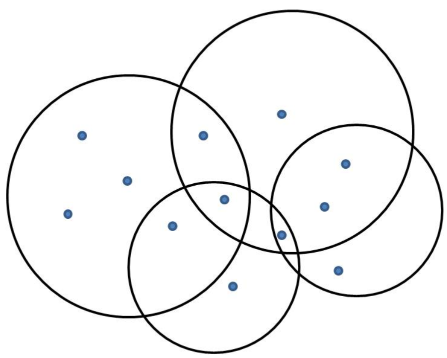

Here, the problem we consider is if:

- players always follow best response dynamics
- we start this game in a low-cost state
- costs of resources can fluctuate between $c_i^t \in [c_i, c_i(1+\varepsilon)]$ or there is one or few Byzantine players

Then, how bad can things get?

**Definition 1 (Price-of Uncertainty(PoU))** *Price-of-Uncertainty($\varepsilon$) of game is the maximum ratio of eventual social cost to initial cost.*

We can model this conceptually as a directed graph, where $s \to t$ indicates that $t$ is reachable from $s$ by a move by one player such that under some perturbation, $t$ has lower cost for that player than $s$ does, which means perturbations can cause BR to move from $s$ to $t$.

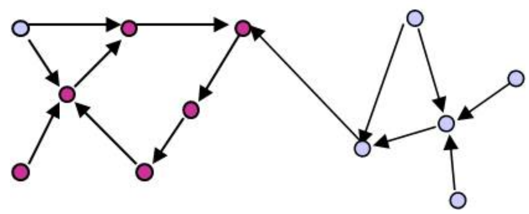

Notice that if $\varepsilon = 0$, we get the best response graph.

### 4.2 Byzantine model

Recall that in previous lecture, we proved that $\forall$ equilibrium $s$, cost of $s$ is at most $n$ times optimal cost. In following picture, two equilibriums have cost $n - \delta$ and 1 separately. And worst equilibrium has cost $n - \delta$ times the optimal cost for any $\delta > 0$.

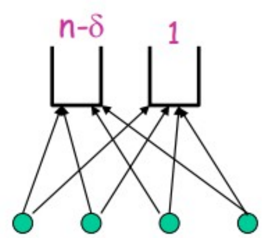


While there is not much mischief a Byzantine player could make in the above game, what we will show is that there exist set-cover game instances such that a single Byzantine player can cause best-response dynamics to move from a Nash equilibrium of cost $O(OPT)$ to one that is a factor $\Omega(n)$ times worse, which in a sense is the worst situation possible.

**Theorem 2** *For set-cover games, a single Byzantine player can cause best-response dynamics to move from a pure Nash equilibrium of cost $O(OPT)$ to a Nash equilibrium of cost $\Omega(n \times OPT)$.*

In the following game, we allow the Byzantine player to also control the order in which players move.

Consider the following set-cover game, there are two kinds of players:

n players of type I, where each player $i$ has two sets to choose from: a common set $s^*$ of cost $n$, and a set $s_i$ of cost $n - \delta$.

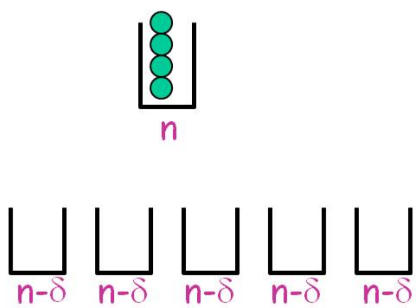


And $n-1$ players of type II, such that player $k$ of type II may either choose any of the sets $s_i$ or else its own set $f_k$ of cost $n/k$.

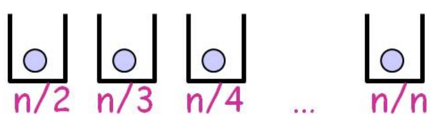


And we have another Byzantine player. First, we are in the best equilibrium. Then, the Byzantine player and type II players slowly lure all type I players into the sets $s_i$. First, Byzantine player moves to $s_1$, every player of type II will move to $s_1$, then player 1 of type I will move to $s_1$. The Byzantine player then sequentially moves to each set $f_{n-1}, ..., f_2$ , causing the players of type II to move to their sets $f_k$ in that order. Specifically, at the time player k of type II moves, the set $s_1$ has a cost to it of $(n-1)/k$, whereas set $f_k$ has cost (with the Byzantine player) of $n/(2k)$.

Now the Byzantine player moves to set $s_2$, causing type II players move to $s_2$. Specifically, at the time player k moves, set $s_2$ has cost $(n-1)/k$ which is lower than the cost $n/k$ of $f_k$. At the end of this step we have the same configuration of type II players as in the initial state, except with $s_2$ rather than $s_1$. The entire process then repeats for player 2 of type I, and so on, until each player i of type I is on its own set $s_i$. Finally, since $s^*$ is now empty, none of the type I players wish to move so we are at an equilibrium.

## 4.3 Perturbation model

Note that in the previous example, if there is no Byzantine player and $\varepsilon = 1$, things could happen just as with the Byzantine player. Type II players can move to $s_i$ sequentially to make the $i$th player of type I moving to $s_i$.

Now, we show that if fluctuations are sufficiently small, then the situation cannot get so out of hand.

**Theorem 3** *In the set-cover game, for any fluctuation $\varepsilon > 0$, the potential of the final set is at most $(1+\varepsilon)^{nm}$ times the original potential.*

Think of players in sets as a stack of chips. View the $i$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 $(1+\varepsilon)$ because fluctuation is at most $(1+\varepsilon)$.

There are at most $mn$ different positions because each player has at most $m$ resources to choose from and each stack has $n$ positions. So, following the path of any chip and removing loops, cost of final set is at most $(1+\varepsilon)^{nm}$ times initial value.

If $\varepsilon = O(1/nm)$, PoU is at most $(1+\varepsilon)^{nm} \times \log n = O(\log n)$ because $\Phi(S) \leq cost(S) \times \log n$.

## 5 Fair cost sharing in general graphs

We now consider fair cost sharing in general graphs, under the assumption that there are many players of each *type*. That is, for each player there are many other players who also have the same starting and ending points. (In game theory, two players are of the same *type* if they have the same strategy space and utility function.)

**Theorem 4** *If number of players of each type is $\Omega(m)$, the PoU of best response is $O(1)$ for any constant $\varepsilon$.*

In this case, it is hard to analyze cost of the state directly. Instead, we track the upper bound $c^*(S_t) = cost(S_0 \cup ... \cup S_t)$.

Call an edge "marked" if it is ever used throughout the best-response process, including those used in the initial state. Then $c^*$ is the total cost of all marked edges. So $c^*$ is always an upper bound on the cost until the current state.

Any time a best-response path for some $(s_i, t_i)$ pair uses an unmarked edge, the total cost of the unmarked edges used in that path is at most $(c^*/n_i)(1+\varepsilon)$ because the average cost of players of this type is at most $c^*/n_i$ before this step, and a player would never deviate to a path of cost more than the average.

Notice that $c^*$ can change at most $m$ times. Each time, the total cost can increase by at most a multiplicative $(1+(1+\varepsilon)/n_i)$ with minimum $n_i$. With $n_i = \Omega(m)$, $(1+(1+\varepsilon)/n_i)^m = O(1)$ when $\varepsilon$ is constant.

## 6 Market sharing model and $\beta$-nice game

Here, we discuss market sharing, in which we replace cost with benefit. Here, each player try to maximize their benefit instead of minimizing their cost.

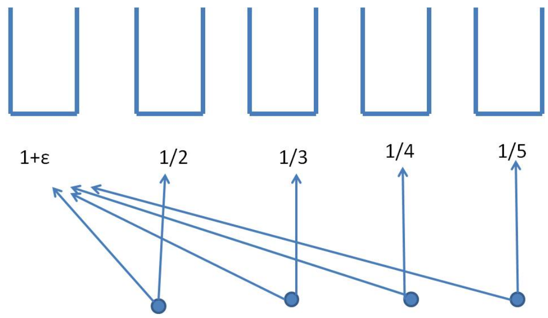


Even if there is no fluctuation, the system might change to a state that is $1/\log(n)$ of optimal by sequentially moving to the resource of $1+\varepsilon$.

In this graph, equilibrium state is a state that the first half players stay in personal their resources, and the second half move to public resource. And the social benefit is about $\log(n/2)$.

One interesting observation is that the social optimal is not an equilibrium. Then, our goal is that equilibrium cannot too far away from optimal.

For each player, we define $\Delta_i(S) = cost(S) - cost(S^i)$ where $S^i$ is $i$th best response, and $\Delta(S) = \sum_i \Delta_i(S)$.

**Definition 5** *An exact potential game $G$ with a potential function $\Phi$ is $\beta$-nice iff for any state $S$, we have $2\Delta(S) \geq cost(S) - \beta OPT$.*

$\beta$-nice games incentive grows stronger as cost gets above $\beta$ times optimal.(Typically $\beta = PoA$)

$\beta$-nice games can at least show state won't get above $\beta$ times optimal, even with substantial perturbation or many Byzantine players.

# 7 Conclusion

In set-cover games,

- a single Byzantine player can make cost $n$ times optimal.
- with perturbation $\varepsilon = 1$, $PoU = \Omega(n)$
- If $\varepsilon = O(1/nm)$, $PoU = O(\log n)$

In general fair-cost-sharing games,

- For each type of players, the number of players is $\Omega(n)$, $PoU = O(1)$ for any constant $\varepsilon > 0$

Subsequent results:

- 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 $\varepsilon$-perturbations
  - Tight bounds on effect of B Byzantine players

Open problems:

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

## References

[EC09] Maria-Florina Balcan, Avrim Blum, Yishay Mansour, The Price of Uncertainty: Safety analysis for multiagent systems