Algorithmic Game Theory March 16, 2016

Lecture 3

Lecturer: Michal Feldman *Scribe: Chen Ziv, Doron Tiferet, Tzvika Geft*

Lesson overview

In the previous lessons we've seen games where there was always a Nash equilibrium. Today we will see an example of other types of games and equilibria. Topics on the agenda:

1. Potential games
2. Solution concepts
3. Cost sharing games

1 Potential Games

Definition 1 *A game is a* ***potential game*** *if there exists a potential function ϕ such that for each profile s and for each unilateral deviation s_i' such that $s' = (s_i', s_-i)$, the following holds for each player i:*

$$\phi(s') - \phi(s) = C_i(s') - C_i(s)$$

Theorem 2 *Every potential game has a pure Nash equilibrium.*

Proof: Let s be a profile that brings ϕ, the game's potential function, to a minimum (such s exists because we assume the number of profiles is finite). By the equality in ϕ's definition, s has to be a Nash equilibrium: If it's not, there is a unilateral deviation by a certain player that reduces his cost function. But, the equality means that the potential function decreases as well due to the deviation, contrary to the fact that s is a minimum. ■

Theorem 3 *Every selfish routing game has a pure Nash equilibrium.*

Proof: Notations:

- f - The players' flow.

- f_e - The number of players who use edge e as part of f.
- $c_e(i)$ - The edge cost e when used by i players.

We show that a selfish routing game is potential game with the following potential function:

$$\phi(f) = \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i)$$

For a given flow f let P_i be the flow according to f from s_i to t_i for player i and let $\hat{P}_i$ be a different flow from s_i to t_i. Let $\hat{f}$ the flow after player i transitioned from P_i to $\hat{P}_i$ Examine what happened to the potential function during the transition:

$$\phi(\hat{f}) - \phi(f) = \sum_{e \in \hat{P}_i \setminus P_i} C_e(f_e + 1) - \sum_{e \in P_i \setminus \hat{P}_i} C_e(f_e) \tag{1}$$

The capacity of every edge that belongs both to P_i and $\hat{P}_i$ doesn't change, because the number of players that went through the edge doesn't change, and thus the cost of those edges cancels each other in the difference.

For each edge $e \in \hat{P}_i$ and $e \notin P_i$ the number of players using e is f_e+1, and thus the difference between $\phi(\hat{f})$ and $\phi(f)$ is $C_e(f_e+1)$. Following a similar principle we can conclude that for each edge $e \in P_i$ and $e \notin \hat{P}_i$ the capacity of e is reduced by 1, and thus the last element of the second operand is not canceled out and remains in the difference.

Examining the change for a specific player:

Recall that $C_i(f) = \sum_{e \in P_i} C_e(f_e)$, so:

$$C_i(\hat{f}) - C_i(f) = \sum_{e \in \hat{P}_i \setminus P_i} C_e(f_e + 1) - \sum_{e \in P_i \setminus \hat{P}_i} C_e(f_e) \tag{2}$$

From the equivalence of (1) and (2) we conclude that the difference in ϕ is equal to the change in the cost function of player i. Hence, it satisfies the conditions of a potential function. ■

Remarks

1. Notice we did not assume anything about the cost function $C_e(\cdot)$ for the proof.
2. We can expand the theorem to congestion games.

Definition 4 *Congestion games are games which have:*

- *A finite number of players.*
- *A finite number of resources E.*
- *Each player i has a collection of possible strategies $S_i \subseteq 2^E$.*
- *Each resource $e \in E$ has a cost functions $C_e(S_e)$ where S_e is the number of players using e in profile S.*

Theorem 5 *Every congestion game has a pure Nash equilibrium.*

Proof: We define a potential function for a congestion game as follows:

$$\phi(S) = \sum_{e \in E} \sum_{i=1}^{S_e} C_e(i)$$

where $C_e(i)$ is the cost for a player to use edge e with capacity i. The rest of the proof is similar to the proof of theorem 3. ■

2 Solution concepts

In the previous lessons we saw the following definitions (all of which refer to minimization games):

Definition 6 *A strategy profile $s = (s_i, s_{-i})$ is a* ***pure Nash equilibrium*** *if for every player i and for every unilateral deviation $s'_i \in S_i$ we have:*

$$C_i(s) \leq C_i(s'_i, s_{-i})$$

Definition 7 *The distributions $\sigma_1, ..., \sigma_k$ over $S_1, ..., S_k$ are a* ***mixed equilibrium*** *if for every player i and for every unilateral deviation $s'_i \in S_i$ we have:*

$$E_{s \sim \sigma}[C_i(s)] \leq E_{s \sim \sigma}[C_i(s'_i, s_{-i})]$$

where σ is the product distribution $\sigma_1 \times ... \times \sigma_k$.

Now we define broader concepts of equilibria:

Definition 8 *The distribution σ over $S_1 \times ... \times S_k$ is a* ***correlated equilibrium (CE)*** *if for every player i and for every $s_i, s'_i \in S_i$ we have:*

$$E_{s \sim \sigma}[C_i(s)|s_i] \leq E_{s \sim \sigma}[C_i(s'_i, s_{-i})|s_i]$$

Each player knows σ in addition to his own strategy s_i.

Example 1 ***Traffic light game***. *Let's describe the following two player maximization game: Each player wants to pass the junction first but if both players cross it together an accident will occur. Each player can decide whether to stop or to cross:*

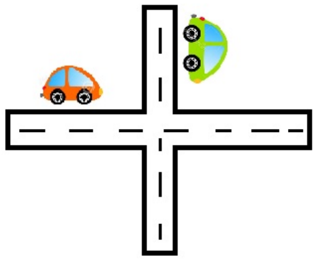

		Player 2	
		Stop	Go
Player 1	Stop	$0, 0$	**0,1**
	Go	**1,0**	$-5, -5$

We have a pure equilibrium in either of two bold cells where one player stops and the other passes. In such an equilibrium one player will never get a chance to cross the junction. A correlated equilibrium gives both players a chance to cross while guaranteeing no accident, like a traffic light. The corresponding distribution looks like this:

		Player 2	
		Stop	Go
Player 1	Stop		0.5
	Go	0.5	

Any change of strategy by a single player will lower the expectation since the change either causes a collision or causes the junction not be used. Note that instead of a 50-50 distribution we can have p and $1-p$ (representing a traffic light with different durations for different directions).

Definition 9 *The distribution σ over $S_1 \times ... \times S_k$ is a* ***coarse correlated equilibrium (CCE)*** *if for every player i and for every unilateral deviation $s_i' \in S_i$ we have:*

$$E_{s\sim\sigma}[C_i(s)] \leq E_{s\sim\sigma}[C_i(s_i', s_{-i})]$$

Remark This equilibrium is the same as the mixed equilibrium except that the distribution is now a joint one over all the players' strategies instead of each player having his own distribution.

The definitions we went over form a hierarchy of equilibria:

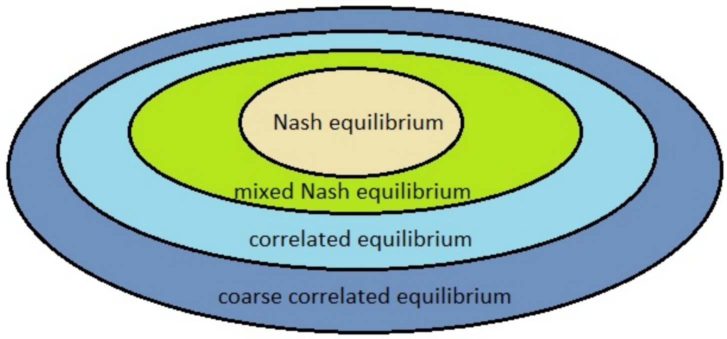

Note that a pure Nash equilibrium doesn't always exist while a mixed one always does. Computationally though, it's hard to find a mixed equilibrium.

Examples

We will now look at a selfish routing game and present different types of equilibria, applying the definitions above. The game has 4 players, each wanting to get from s to t. There are 6 parallel edges between s and t, each having $c_e(x) = x$ as the cost function:

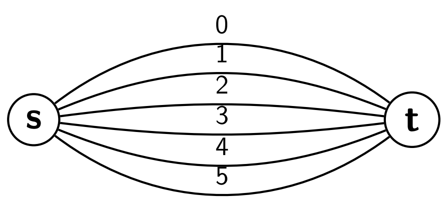

Pure equilibrium: Each player chooses his own edge. There are $\binom{6}{4}$ such equilibria.

Mixed equilibrium: Each player randomly (uniformly) chooses an edge.

CE: A uniform distribution over all the strategy profiles where there is one edge used by two players and two edges with one player each.

CCE: The uniform distribution will be similar to the CE case except that now the set of edges will be either $\{0, 1, 2\}$ or $\{3, 4, 5\}$. Note how this is not a CE since when a player "is told" to take edge 3 for example, a better strategy would be to take one of the edges $\{0, 1, 2\}$, which are guaranteed to be free according to the distribution.

We can see how the overall cost of the equilibrium gets worse as it becomes more general. In the pure equilibrium case the cost is 4 which is optimal, while in the CE and CCE cases the cost goes up to 6 (and can be increased further). This leads us to continue the discussion on price of anarchy we started previous lesson.

Price of Anarchy Generalization

Last class we proved that the price of anarchy for atomic selfish routing games with edges having affine costs is bounded by 2.5. The proof was done with respect to pure equilibria but we would like to be able to talk about the more general equilibria introduced this class. Just before we go in that direction, we will recap the stages in our previous proof:

Stage 1: We used the Nash equilibrium definition to get the following inequality for the ith player:

$$C_i(s) \leq C_i(s_i^*, s_{-i})$$

Where s is the profile of some pure Nash equilibrium and s^* is the optimal profile.

Stage 2: Summing up the above inequality over all the players allowed us to bound the cost of the equilibrium:

$$cost(s) = \sum_i C_i(s) \leq \sum_i C_i(s_i^*, s_{-i})$$

Stage 3: We then used a lemma that allowed us to separate the costs of the optimal profile and the equilibrium:

$$\sum_i C_i(s_i^*, s_{-i}) \leq \frac{5}{3} cost(s^*) + \frac{1}{3} cost(s)$$

Rearranging the inequalities then gave the desired result:

$$\frac{cost(s)}{cost(s^*)} \leq 2.5$$

We would like to generalize the process above. Let's look at a definition that will help us:

Definition 10 *A minimization game is called* $(\lambda, \mu)-$*smooth if*

$$\sum_i C_i(s_i^*, s_{-i}) \leq \lambda cost(s^*) + \mu cost(s) \tag{3}$$

for every strategy profile s *and for some optimal profile* s^* *when*

$$cost(s) \leq \sum_i C_i(s). \tag{4}$$

Claim 11 *For every minimization game that is* $(\lambda, \mu)-$*smooth we have* $PoA \leq \frac{\lambda}{1-\mu}$, *with respect to pure Nash equilibrium.*

Proof: We essentially repeat our steps above (substituting λ and μ) to get the desired results. Let s be a pure NE:

$$cost(s) \leq \sum_i C_i(s) \leq \sum_i C_i(s_i^*, s_{-i}) \leq \lambda cost(s^*) + \mu cost(s)$$
$$(1-\mu)cost(s) \leq \lambda cost(s^*)$$
$$PoA = \frac{cost(s)}{cost(s^*)} \leq \frac{\lambda}{1-\mu}$$

■

The claim gives us very little progress since it only refers to pure equilibria, which don't always exist. Note that even extending it to mixed equilibria (whose existence is guaranteed) won't help us much in practice since finding them is computationally hard. We therefore strengthen it to the most general equilibrium learned so far:

Theorem 12 *Under the same conditions $PoA \leq \frac{\lambda}{1-\mu}$ with respect to CCE.*

Proof: Let σ be a CCE and s^* be an optimal profile. The proof resembles the steps we discussed above except that we are now working with the expected cost:

$$\begin{aligned} E_{s\sim\sigma}[cost(s)] &\underset{(4)}{\leq} E_{s\sim\sigma}\big[\sum_i C_i(s)\big] \\ &= \sum_i E_{s\sim\sigma}[C_i(s)] \underset{CCE}{\leq} \sum_i E_{s\sim\sigma}[C_i(s_i^*, s_{-i})] \\ &= E_{s\sim\sigma}\big[\sum_i C_i(s_i^*, s_{-i})\big] \underset{(3)}{\leq} E_{s\sim\sigma}[\lambda cost(s^*) + \mu cost(s)] \\ &= \lambda cost(s^*) + \mu E_{s\sim\sigma}[cost(s)] \end{aligned}$$

As in stage 3 we now rearrange to have:

$$PoA = \frac{E_{s\sim\sigma}[cost(s)]}{cost(s^*)} \leq \frac{\lambda}{1-\mu}$$

■

Remark Going back to the definition of $(\lambda, \mu)-$smoothness, note that we require the inequality (3) to hold for every strategy profile s. In our pure NE proof we only care about equilibrium profiles, so we don't need such a strict requirement. In contrast, the theorem we just proved relies heavily on it holding for other profiles since a profile s which is drawn from σ is not necessarily a pure equilibrium.

3 Cost-Sharing Games

A cost-sharing game for k players consists of a directed graph $G = (V, E)$, where each edge $e \in E$ has a constant cost γ_e. Each player i has a source vertex s_i and a target vertex t_i. The strategy space of player i consists of the set of all paths from s_i to t_i in G. We define the strategy profile as the vector $p = (p^1, p^2, ..., p^k)$, where p^i denotes player i's strategy. The corresponding graph for strategy profile p is $\left(V,\ \bigcup_{i=1}^{k} p^i\right)$.

In **fair** cost-sharing games, the cost of using an edge e under strategy profile p is evenly distributed among all players that use the edge (the number of such players being p_e):

$$c_e\left(p_e\right) = \tfrac{\gamma_e}{p_e}$$

The cost each player i pays under profile p is given by $c_i\left(p\right) = \sum_{e \in p^i} c_e(p_e)$, and the collective cost of the profile is defined as the sum of all players costs:

$$cost(p) = \sum_{i=1}^{k} c_i(p) = \sum_{i=1}^{k} \sum_{e \in p^i} c_e(p_e) = \sum_{i=1}^{k} \sum_{e \in p^i} \frac{\gamma_e}{p_e} = \sum_{e \in E,\ p_e \geq 1} \gamma_e$$

Note that cost-sharing games are potential games with the following potential function:

$$\phi(p) = \sum_{e \in E} \sum_{i=1}^{p_e} c_e(i) = \sum_{e \in E} \sum_{i=1}^{p_e} \frac{\gamma_e}{i}$$

As we have seen in Theorem 2, all local minima of the potential function are Nash equilibria.

Claim 13 *There exists a cost-sharing game where the price of anarchy, with respect to pure NE, is the number of players.*

Proof: Consider the following game where all players share the source node s and target node t:

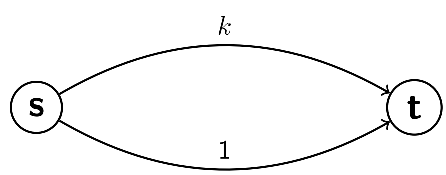

The optimal strategy profile p^* has all the players using the bottom edge, with total cost being $cost(p^*) = \sum_{e \in E,\ p^*_e \geq 1} \gamma_e = 1$.

Consider the strategy profile p where all the players use the top edge. The cost each player pays is given by $c_i(p^*) = \sum_{e \in (p^*)^i} \frac{\gamma_e}{p^*_e} = \frac{\gamma_{e_1}}{p^*_{e_1}} = \frac{k}{k} = 1$. The only deviation that can happen is a player taking the bottom edge and paying $c_i(p) = \sum_{e \in p^i} \frac{\gamma_e}{p_e} = \frac{\gamma_{e_2}}{p_{e_2}} = \frac{1}{1} = 1$, which has no added value. This makes p a pure Nash equilibrium, with $cost(p) = \sum_{e \in (p^*)^i,\ p^*_e \geq 1} \gamma_e = k$.

The price of anarchy for this game is then $PoA = \frac{cost(p)}{cost(p^*)} = \frac{k}{1} = k$ ■

By the lemma, we get that cost-sharing games' PoA could not be bounded by a constant. At first, this result seems to imply that in such games the players' selfish behavior will likely lead to high overall cost. But we can define a different measure: instead of referring to the worst price equilibrium, we examine the price of the best price equilibrium. By setting the initial state to the best equilibrium, we can have a system that doesn't necessarily have a very high overall cost and still maintain stability.

Definition 14 *The price of stability, denoted PoS, is defined as follows:*

$$PoS = \frac{cost(\text{best equilibrium})}{cost(OPT)}$$

For the aforementioned game, the PoS will be 1, since the optimal state is also an equilibrium. This encouraging result leads us to believe that the PoS for all cost-sharing games could be upper-bounded by a constant.

Claim 15 *The maximal PoA for all cost-sharing games could not be upper-bounded by a constant.*

Proof: Consider the following k-player game:

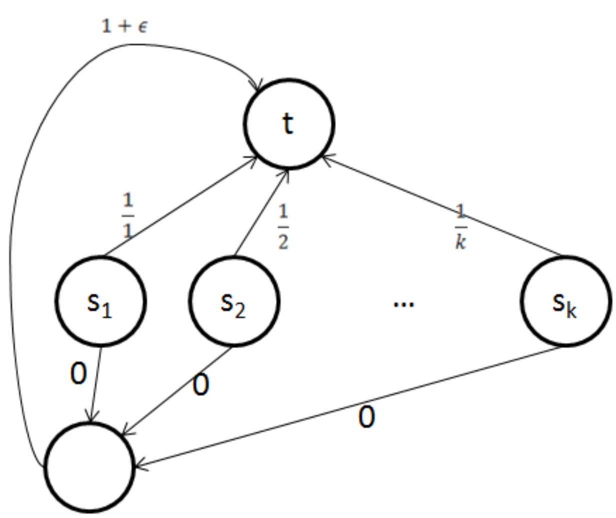

Each player i is assigned with a source node s_i and a target node t.

The optimal profile involves all players using the edge whose cost is $1+\epsilon$, which will yield the collective cost $1+\epsilon$. However, all strategies which involve the $1+\epsilon$ edge are not an equilibrium. We can prove that by contradiction: Let's assume there is an equilibrium which involves the $1+\epsilon$ edge, and let s_m be the node with maximal index which routes to this edge. There are no more than m players using this edge, which means the price of this route is lower-bounded by $\frac{1+\epsilon}{m}$. That means by changing his route, the s_m player could achieve the cost of $\frac{1}{m}$, in contradiction the profile being an equilibrium. That means there is only one equilibrium, which involves the upper edges. Hence:

$$cost(OPT) = 1+\epsilon$$

$$cost(\text{best equilibrium}) = \sum_{i=1}^{k} \frac{1}{i} = H_k \approx \log(k)$$

$$\Rightarrow \frac{cost(\text{best equilibrium})}{cost(OPT)} \approx \log(k)$$

■

Lemma 16 *In fair cost-sharing games, for each strategy profile p we have the following:*

$$cost\,(p) \leq \Phi\,(p) \leq H_k \cdot cost(p)$$

Proof: For every edge e, which is used in strategy profile p we have the following inequality:

$$1 \leq \sum_{i=1}^{p_e} \frac{1}{i} \leq \sum_{i=1}^{k} \frac{1}{i} = H_k$$

By multiplying by γ_e we get:

$$\gamma_e \leq \gamma_e \cdot \sum_{i=1}^{p_e} \frac{1}{i} \leq \gamma_e \cdot H_k$$

Now, if we sum on all edges that are used by some player in the profile we get:

$$\sum_{e \in p^i,\ p_e \geq 1} \gamma_e = cost(p) \leq \sum_{e \in p^i,\ p_e \geq 1} \gamma_e \cdot \sum_{i=1}^{p_e} \frac{1}{i} = \Phi\,(p) \leq \sum_{e \in p^i,\ p_e \geq 1} \gamma_e \cdot H_k = H_k \cdot cost(p)$$

■

Theorem 17 *In all fair cost-sharing games for k players: $PoS \leq H_k$*

Proof: Let p be a strategy profile which minimizes Φ. By Theorem 2, p is a Nash equilibrium.

According to Lemma 14: $cost(p) \leq \Phi(p)$.

By the selection of p and the optimal strategies profile p^* the following holds: $\Phi(p) \leq \Phi(p^*)$.

By the second part of the lemma: $\Phi(p^*) \leq H_k \cdot cost(p^*)$.

Combining the inequalities so far yields $\frac{cost(p)}{cost(p^*)} \leq H_k$.

Since $PoS \leq \frac{cost(p)}{cost(p^*)}$, we get $PoS \leq H_k$ ■

Definition 18 *Given a strategy profile S, a coalition deviation $S'_A \in \prod_{i \in A} s_i$ is called a* ***beneficial coalitional deviation*** *for coalition A, if*

$$\forall i \in A : c_i\left(S'_A,\ S_{-A}\right) \leq c_i(S), \text{ and}$$
$$\exists i \in A : c_i\left(S'_A,\ S_{-A}\right) < c_i(S)$$

Definition 19 *A profile S is a* ***strong equilibrium*** *if there are no beneficial coalitional deviations from it.*

Theorem 20 *The PoA of a fair cost-sharing game for k players, with respect to strong equilibrium, is bounded by H_k.*

The theorem will be proven next class.